Release Notes for IronWare Software Release 03.1.02

FastIron Edge Switch

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

Release date: February 27, 2004

Publication date: February 27, 2004

These release notes contain the enhancements and software fixes in software releases 03.1.02, 03.1.01, and 03.1.00.

This release supports the following products:

- FastIron Edge Switch 2402
- FastIron Edge Switch 4802
- FastIron Edge Switch 9604
- FastIron Edge Switch 12GCF
- FastIron Edge Switch 2402-POE
- FastIron Edge Switch 4802-POE

Software release 03.1.02 is a maintenance release. This release applies to the following Foundry Networks products:

- FastIron Edge Switch 2402
- FastIron Edge Switch 4802
- FastIron Edge Switch 9604
- FastIron Edge Switch 12GCF
- FastIron Edge Switch 2402-POE
- FastIron Edge Switch 4802-POE

These release notes contain the enhancements and software fixes in the following software releases:

- 03.1.02
- 03.1.01
- 03.1.00

NOTE: The user manuals for the FastIron Edge Switch are the same as that for IronWare software release 07.6.03. These release notes contain information that is specific to FastIron Edge Switches and describe features where FastIron Edge Switches differ from IronWare software release 07.6.03.

NOTE: Except where explicitly mentioned in these release notes, the FES2402 and FES2402-POE are similar devices, and the FES4802 and FES4802-POE are similar devices. For example, the FES2402 and FES2402-POE have similar network interfaces and port regions. The same is true of the FES4802 and FES4802-POE.

Summary of Enhancement in 03.1.02

Enhancement	Description	See Page
OSPF Syslog enhancement	You can specify which kinds of OSPF-related Syslog messages are logged.	15

Summary of Enhancement in 03.1.01

Enhancement	Description	See Page
Support for standard OSPF Version 2 MIB	Software release 03.1.01 provides support for the standard OSPF Version 2 Management Information Base (RFC 1850).	N/A

Summary of Enhancements in 03.1.00

This section lists a summary of the enhancements added in release 03.1.00.

Enhancement	Description	See Page
Ability to configure Metro Ring Protocol (MRP) Phase I	You can use Foundry's proprietary protocol, MRP, to prevent Layer 2 loops and provide fast reconvergence in Layer 2 ring topologies.	15

Enhancement	Description	See Page
Support for Private VLANs	Software release 03.1.00 provides support for Private VLANs on untagged ports. Private VLANs have the properties of standard Layer 2 port-based VLANs but also provide additional control over flooding packets on a VLAN.	15
Support for Protocol 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, thereby limiting the amount of broadcast traffic end-stations, servers, and routers need to accept.	16
Support for Subnet VLANs	For IP, and IPX you can provide more granular broadcast control by creating subnet VLANs.	17
Support for Super Aggregated VLANs (SAVs)	You can aggregate multiple VLANs within another VLAN, enabling the construction of Layer 2 paths and channels.	17
Super Aggregated VLAN 802.1Q-in-Q Support	This release supports the Super Aggregated VLAN 802.1Q-in-Q functionality, which enables interoperability of Foundry devices with other vendors' devices.	17
Support for Topology Groups	You can use topology groups to simplify configuration and enhance scalability of Layer 2 protocols. This feature enables the Foundry device to run a single instance of a Layer 2 protocol on multiple VLANs.	21
Ability to configure Uni-directional Link Detection (UDLD)	With UDLD, the device 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.	21
Support for VSRP	Virtual Switch Redundancy Protocol (VSRP) is a Foundry proprietary protocol that provides redundancy and subsecond failover in Layer 2 and Layer 3 mesh topologies.	21

Layer 3 Enhancement in 03.1.00

Enhancement	Description	See Page
Ability to apply an OSPF distribution list to an interface	Software release 03.1.00 enables you to apply an OSPF distribution list to a physical or virtual interface. In releases prior to 03.1.00, you could configure an OSPF distribution list on a global basis only.	22

System-Level Enhancements in 03.1.00

Enhancement	Description	See Page
Support for Protected Link Groups	You can configure protected link groups to minimize disruption to the network by protecting critical links from loss of data and power.	22

Enhancement	Description	See Page
Defining a port for Web management interface	A new CLI command and new MIB objects allow you to specify the TCP port that will be used to access a device's Web management interface.	24
Ability to detect POE power requirements advertised via Cisco Discovery Protocol (CDP)	In release 3.1, the FES2402-POE and FES4802-POE are compatible with Cisco's and other vendors' power consuming devices, in that they can detect and process power requirements for these devices automatically.	25
Ability to specify the amount of power to provide for a Power over Ethernet (POE) power consuming device	You can specify the amount of power that a port on the FES2402-POE or FES4802-POE should provide for a power consuming device.	25
Ability to specify the power class for a POE power consuming device	You can specify the power class for a power consuming device. The power class determines the amount of power the port should provide for a power consuming device.	26
New SNMP MIB objects for POE power consuming device	SNMP MIB objects have been added for the POE Power Consuming Device power level and power class.	27
Dynamically applying IP ACLs and MAC Filters to 802.1X ports	Foundry's 802.1X implementation supports dynamically applying an IP ACL or MAC address filter to a port, based on information received from an Authentication Server.	28

Feature Support

The FastIron Edge Switches support many of the applicable system-level, Layer 2 and Layer 3 features supported by the FastIron 4802 (FWS 4802) and BigIron Chassis devices. Configuration for most of the features is the same on the FastIron Edge Switches and on the FastIron 4802 or BigIron Chassis device.

The features that are available on the device depend on the type of software image the device is running. You can run one of the following types of software images:

- Layer 2 (supported on all models)
- Base Layer 3 (supported on all models)
- Layer 3 (full Layer 3, supported on premium models only)

Table 1 lists the software that is loaded into the device's primary and secondary flash areas at the factory. All the flash images are included on the CD-ROM shipped with the device.

Table 1: Default Software Loads

Model	Software Images	
	Primary Flash	**Secondary Flash**
FES2402 FES4802 FES9604 FES2402-POE FES4802-POE FES12GCF	Layer 2	Base Layer 3
FES2402-PREM FES4802-PREM FES9604-PREM FES2402-POE-PREM FES4802-POE-PREM FES12GCF-PREM	Full Layer 3	Layer 2

The following sections list the highlights of the features that are supported in this release.

System-Level Features Supported

All models provide the following system features:

- 802.1p prioritization
- 802.3ad link aggregation
- Auto MDI/MDIX
- Broadcast, multicast and unknown-unicast rate limiting
- DiffServ support
- Extensive management options:
  - Serial and Telnet access to industry-standard Command Line Interface (CLI)
  - Web-based GUI
  - Support for optional IronView Network Manager (standalone and HP OpenView GUI)
  - Access Control Lists (ACLs) for controlling management access
- Foundry Discovery Protocol (FDP) / Cisco Discovery Protocol (CDP)
- Jumbo frames – up to 9192 bytes (FES12GCF only)
- Multiple Syslog server logging (up to six Syslog servers)
- OSPF Version 2 MIB (RFC 1850)
- Port monitoring

- Priority mapping using ACLs
- Protected link groups
- Rate limiting (port-based input rate limiting)
- Robust security:
  - Access Control Lists (ACLs) for filtering transit traffic (applies to IP unicast traffic only)
  - Denial of Service (DoS) protection (SYN Attacks, Smurf Attacks)
  - Local passwords
  - User accounts
  - Authentication, Authorization and Accounting (AAA)
  - RADIUS, TACACS/TACACS+
  - Secure Shell (SSH) version 1.5
  - MAC port security
  - 802.1X port security
  - Address locking
  - Layer 2 MAC filtering (filtering on source and destination MAC addresses supported)
- Server trunk groups (switched IP and IPX traffic only)
- sFlow (RFC 3176) for inbound and outbound traffic
- SNMP V1, V2c, V3
- Static MAC entries with option to set priority
- Switch trunk groups

Layer 2 Features Supported

The Layer 2 software supports the following features:

- 802.1d Spanning Tree Support
  - Enhanced IronSpan support includes Fast Port Span, Fast Uplink Span, and Single-instance Span
  - Rapid Spanning Tree support allows for sub-second convergence (draft 3 supported)
  - Cisco PVST/PVST+ compatibility
- 802.1p Quality of Service (QoS)
  - Weighted Round Robin (WRR)
  - Strict Priority (SP)
- 802.1W (Rapid Spanning Tree), final IEEE standard
- Dynamic Host Configuration Protocol (DHCP) Assist
- IGMPv2 snooping (Layer 2 Multicast)
- Metro Ring Protocol 1 (MRP 1)

  In FES, the RHP received counter on non-master MRP nodes increment. This is different on other devices that support MRP 1.
- STP per-VLAN Group
- Topology groups
- Uni-directional Link Detection (UDLD) (Link keepalive)

- Virtual Switch Redundancy Protocol (VSRP)
- VLAN Support
  - 802.1Q with tagging
  - 802.1Q-in-Q Super Aggregated VLANs (SAVs)
  - Dual-mode VLANs
  - GVRP
  - Private VLANs (untagged ports only)
  - Protocol VLANs (AppleTalk, IPv4, dynamic IPv6, IPX, Decnet, NetBIOS, and other protocol types)
  - Layer 3 Subnet VLANs (IP subnet and IPX network)
  - Super Aggregated VLANs (SAVs)
  - Virtual routing interfaces
  - VLAN groups
- Wire-speed Layer 2 switching

Layer 3 Features Supported

The Layer 3 image supports all the Layer 2 features along with the following:

- AppleTalk
- Global route-only support (disabling Layer 2 switching on a global basis)
- IGMP V2
- IP
- IP multicast (DVMRP, PIM-SM, PIM-DM)
- IPX
- OSPF
- RIP V1 and V2
- VRRP and VRRPE

NOTE: The full Layer 3 image is supported on FES premium models only.

Unsupported Features

The FastIron Edge Switches do not support the following features. If required, these features are available on other Foundry devices.

System-Level Features not Supported

- Per-port route-only (global route-only is supported)
- Jumbo frames (on all models except FES12GCF)
- NetFlow
- Output rate limiting
- Server trunk groups for Layer 3 traffic (server trunking of switched AppleTalk traffic also not supported)
- Broadcast and multicast filters

- Standard BGP MIBs in RFC 1657 *Definitions of Managed Objects for the Fourth Version of the Border Gateway Protocol (BGP-4) using SMIv2*

Layer 2 Features not Supported

- SuperSpan
- VLAN-based priority

Layer 3 Features not Supported

- BGP
- Foundry Standby Router Protocol (FSRP)
- IS-IS
- Multiprotocol Border Gateway Protocol (MBGP)
- Multiprotocol Label Switching (MPLS)
- Multiprotocol Source Discovery Protocol (MSDP)
- Network Address Translation (NAT)
- Policy-Based Routing (PBR)

Feature Documentation

For feature descriptions and configuration information, see the remaining sections in these release notes and the Foundry product manuals listed in “Where To Get More Information” on page 53.

Software Image Files

To use the features in this release, you need to run the software listed in Table 2.

Table 2: Software Image Files*

Boot Image[a]	Flash Image[b]
FEB03102.bin and FEM03102.bin	FES03102.bin – Layer 2 or FER03102.bin – Layer 3 or FEL03102.bin – Base Layer 3

a.The FastIron Edge Switches have two boot images: FEBxxxxx (boot code) and FEMxxxxx (boot monitor).

The FEB03102 and FEM03102 images are backward compatible with software versions 03.*x.x.*

The FES03102, FER03102, and FEL03102 flash images require the FEB03102 and FEM03102 images. Do not use older versions of the boot and monitor images with the 03.1.02 flash images.

b.The FastIron Edge Switches have two flash areas and can thus contain two separate flash code images. However, the device can run only one image or the other at a given time.
* These images are applicable to the FastIron Edge Switches only. Also, you cannot load other images, such as B2R or B2S for BigIron and FastIron devices, on the FastIron Edge Switches.

The software is loaded at the factory. Table 1 on page 4 lists the default software loads for the device. All the software images are provided on the software CD-ROM shipped with the device. To install another software image, use the instructions in Upgrading Software.

NOTE: The software described in these notes applies only to the FastIron Edge Switches. You cannot use this software on other Foundry Stackable devices or on Foundry Chassis devices.

Upgrading Software

Use the following procedures to upgrade the software.

NOTE: This section does not describe how to upgrade a base model to a PREM model. To perform this upgrade, you need an upgrade kit. Contact Foundry Networks for information.

Upgrading the Boot and Flash Code

NOTE: Software release 03.1.02 requires 03.1.02 boot code and monitor images.

NOTE: The 03.1.02 boot code and monitor images are backward compatible with software releases 03.*x.x.*

NOTE: The FastIron Edge Switches have two boot code images, whereas other Foundry products have only one boot image. When you upgrade, make sure you install both boot images unless advised otherwise by the release notes accompanying the upgrade.

1. Place the new boot 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: `FES4802 Router#`) to copy the FEB boot code from the TFTP server into the flash memory of the management module.
   - **copy tftp flash** <ip-addr> <FEB-image-file-name> **boot**
3. Enter the following command at the Privileged EXEC level of the CLI (example: `FES4802 Router#`) to copy the FEM boot code from the TFTP server into the flash memory of the management module.
   - **copy tftp flash** <ip-addr> <FEM-image-file-name> **mon**
4. Verify that the code has been successfully copied by entering the following command at any level of the CLI:
   - **show flash**

   The line that begins with "Compressed Boot-Tftp Code size" lists the boot (FEB) code version and the line that begins with "Compressed Boot – Monitor Image size" lists the monitor (FEM) code version.

NOTE: Do not reboot. You must first upgrade the flash code as instructed in Step 5 through Step 8, below.

5. Place the new flash code on a TFTP server to which the Foundry device has access.
6. Enter the following command at the Privileged EXEC level of the CLI (example: `FES4802 Router#`) to copy the flash code from the TFTP server into the flash memory of the management module:
   - **copy tftp flash** <ip-addr> <image-file-name> **primary | secondary**
7. Verify that the flash code has been successfully copied by entering the following command at any level of the CLI:
   - **show flash**

   The line that begins "Compressed Pri Code size" lists the flash code version in the primary flash, at the end of the line. Similarly, the line that begins "Compressed Sec Code size" lists the flash code version in the secondary flash.
8. If the flash code version is correct, go to Step 9. Otherwise, go to Step 5.
9. 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: Starting with release 03.1.00, the flash image checks for compatibility with the boot and monitor images. If the images are not compatible, the FastIron Edge Switch displays a message on the console after the system boots up and before the first prompt appears. The device also logs a message to the Syslog every 15 minutes. The following shows an example message:

```
ALERT:  Currently active Boot-tftp version 01.1.00 is NOT compatible!  Please
update with version 03.1.00 or higher.
```

Managing the Device

You can manage the FastIron Edge Switch 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 – Optional standalone SNMP-based GUI application.

Logging on Through the CLI

After you configure an IP address, 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* for more information.

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:

```
FES4802 Router(config)# rooter 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
arp
boot
```

some lines omitted for brevity...

```
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 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 3: 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*.

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 login panel.

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 displays.

Figure 1 Web management interface login dialog

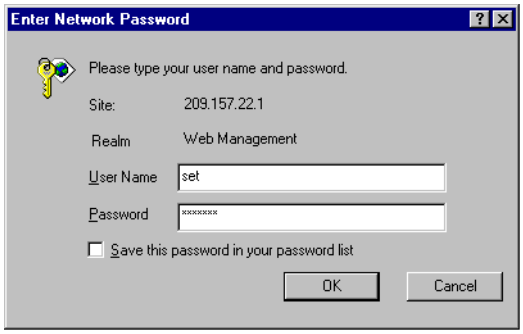

By default, you can use the user name “get” and the default read-only password “public” for read-only access. However, for read-write access, you must enter “set” for the user name, and enter a read-write community string you have configured on the device for the password.

There is no default read-write community string. You must add one using the CLI. To add an encrypted community string, enter a command such as the following:

```
FES4802 Router(config)# snmp-server community private rw
```

Syntax: snmp-server community <string> ro | rw

The <string> parameter specifies the community string name. The string can be up to 32 characters long.

The **ro | rw** parameter specifies whether the string is read-only (**ro**) or read-write (**rw**).

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*.

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 change 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:

```
FES4802 Router(config)# web-management list-menu
```

To disable the front panel frame:

```
FES4802 Router(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.

You can view a picture of the device's front panel, as shown in Figure 2, by clicking on the Monitor -> Front Panel menu option.

Figure 2 FastItron Edge Switch front panel in Web management interface

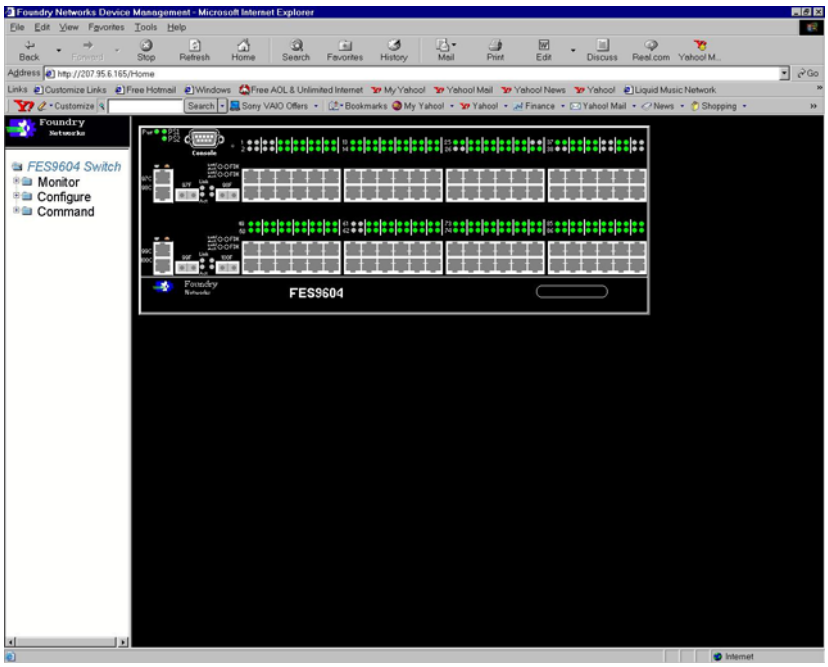

Recovering from a Lost Password

By default, the CLI does not require passwords. However, if someone has configured a password for the device but the password has been lost, you can regain super-user access to the device using the following procedure.

NOTE: Recovery from a lost password requires direct access to the serial port and a system reset.

To recover from a lost password:

1. Start a CLI session over the serial interface to the FastIron Edge Switch.
2. Reboot the device.
3. While the system is booting, before the initial system prompt appears, enter **b** to enter the boot monitor mode.
4. Enter **no password** at the prompt. (You cannot abbreviate this command.)
5. Enter **boot system flash primary** at the prompt. This command causes the device to bypass the system password check.
6. After the console prompt reappears, assign a new password.

Displaying and Saving Configuration Changes

When you make a configuration change, the change enters the device's running configuration but is not saved if you reload the software. To make a change permanent, you must save the change to the device's startup-config file.

Displaying Configuration Changes

To display the running configuration, enter the following command from any level of the CLI:

```
FES4802 Router# show running-config
```

To display the startup configuration, enter the following command from any level of the CLI:

```
FES4802 Router# show configuration
```

NOTE: You cannot display the running-config or startup-config file using the Web management interface.

Saving Configuration Changes

To permanently save a configuration change so that the change stays in effect following a software reload, use one of the following methods.

USING THE CLI

To replace the startup configuration with the running configuration, enter the following command at any Privileged EXEC or CONFIG command prompt:

```
FES4802 Router# write memory
```

USING THE WEB MANAGEMENT INTERFACE

1. Click on the plus sign next to Command in the tree view to expand the list of command options.
2. Select the Save to Flash option.
3. Select Yes when the Web management interface asks you whether you really want to save the configuration changes to flash.

Enhancement in 03.1.02

OSPF Syslog Enhancement

In release 03.1.02, you can specify which kinds of OSPF-related Syslog messages are logged. In releases prior to 03.1.02, by default all OSPF Syslog messages are logged. In configurations with a large amount of OSPF activity, this can result in the Foundry device's Syslog buffer and the Syslog server filling up with OSPF messages.

Starting with this release, by default the only OSPF messages that are logged are those indicating possible system errors. If you want other kinds of OSPF messages to be logged, you can configure the Foundry device to log them.

For example, to specify that all OSPF-related Syslog messages be logged, enter the following commands.

```
BigIron(config)# router ospf
BigIron(config-ospf-router)# log all
```

Syntax: [no] log all | adjacency | bad_packet [checksum] | database | memory | retransmit

The **log** command has the following options:

The **all** option causes all OSPF-related Syslog messages to be logged. If you later disable this option with the **no log all** command, the OSPF logging options return to their default settings.

The **adjacency** option logs essential OSPF neighbor state changes, especially on error cases. This option is disabled by default.

The **bad_packet checksum** option logs all OSPF packets that have checksum errors. This option is enabled by default.

The **bad_packet** option logs all other bad OSPF packets. This option is disabled by default.

The **database** option logs OSPF LSA-related information. This option is disabled by default.

The **memory** option logs abnormal OSPF memory usage. This option is enabled by default.

The **retransmit** option logs OSPF retransmission activities. This option is disabled by default.

Layer 2 Enhancements in 03.1.00

This section describes the Layer 2 software enhancements in release 03.1.00.

Metro Ring Protocol (MRP) Phase I

The Metro Ring Protocol Phase I is a Foundry proprietary protocol that prevents Layer 2 loops and provides fast reconvergence in Layer 2 ring topologies. MRP is an alternative to the Spanning Tree Protocol (STP) or Rapid Spanning Tree Protocol (RSTP) and is especially useful in Metropolitan Area Networks (MANs) where ring topologies are commonly used.

MRP can detect and heal a break in the ring in sub-second time. You can use MRP in combination with VSRP to provide additional design flexibility.

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).

For more information about MRP and VSRP, including configuration procedures, see the *Foundry Switch and Router Installation and Basic Configuration Guide*.

Private VLANs

Private VLANs have the properties of standard Layer 2 port-based VLANs but also provide additional control over flooding packets on a VLAN.

Software releases 03.1.00 and later include support for all three private VLAN port types (on untagged ports only), including:

- 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.
- 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.

For more information about private VLANs, including configuration rules and how to configure them, see the *Foundry Switch and Router Installation and Basic Configuration Guide*.

Configuration Notes

- When Protocol or Subnet VLANs are enabled, or if 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.
- Releases 03.1.00 and later support 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.

Protocol VLANs

Release 03.1.00 includes support for protocol-based VLANs, which provide the ability to define separate broadcast domains for several unique Layer 3 protocols within a single Layer 2 broadcast domain, thereby limiting the amount of broadcast traffic end-stations, servers, and routers need to accept.

See the *Foundry Switch and Router Installation and Basic Configuration Guide* for information on how to configure protocol VLANs.

Configuration Notes

- When Protocol or Subnet VLANs are enabled, or if 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.
- 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.

Sub-net VLANs

Protocol VLANs described in the previous section provide separate protocol broadcast domains for specific protocols. For IP, and IPX you can provide more granular broadcast control by instead creating the following types of VLANs:

- ***IP sub-net VLAN*** – An IP sub-net broadcast domain for a specific IP sub-net.
- ***IPX network VLAN*** – An IPX network broadcast domain for a specific IPX network.

Configuration Notes

- When Protocol or Subnet VLANs are enabled, or if 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.
- 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 membership mask.

See the *Foundry Switch and Router Installation and Basic Configuration Guide* for information on how to configure sub-net VLANs.

Super Aggregated VLANs (SAVs)

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 sub-net across multiple networks.

For an application example and configuration information, see the chapter “Configuring Super Aggregated VLANs” in the *Foundry Switch and Router Installation and Basic Configuration Guide*.

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.

802.1Q-in-Q Tagging

802.1Q tagging is an IEEE standard that enables 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 was sent. The tag and VLAN ID keep traffic from each VLAN segregated and private.

NOTE: This section provides details of the 802.1Q-in-Q feature only. For more details about SAV, see the section “Configuring Aggregated VLANs” in the *Foundry Switch and Router Installation and Basic Configuration Guide*.

- 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 3 shows an 802.1Q configuration example.

Figure 3 802.1Q Configuration Example

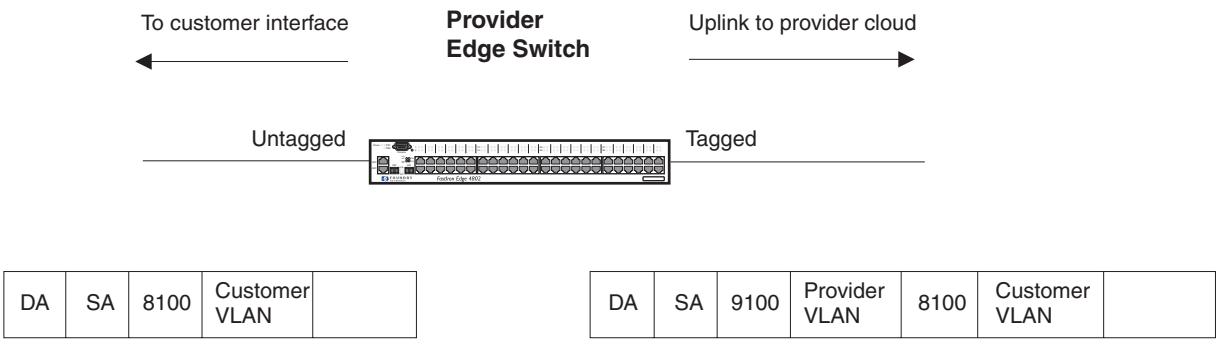

As shown in Figure 3, the ports to customer interfaces are untagged, whereas the uplink ports to the provider cloud are tagged, because multiple client VLANs share the uplink to the provider cloud. In this example, the Foundry device treats the customer's private VLAN ID and 8100 tag type as normal payload, 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 4 shows an example application of the 802.1Q-in-Q enhancement.

Figure 4 802.1Q-in-Q Configuration Example

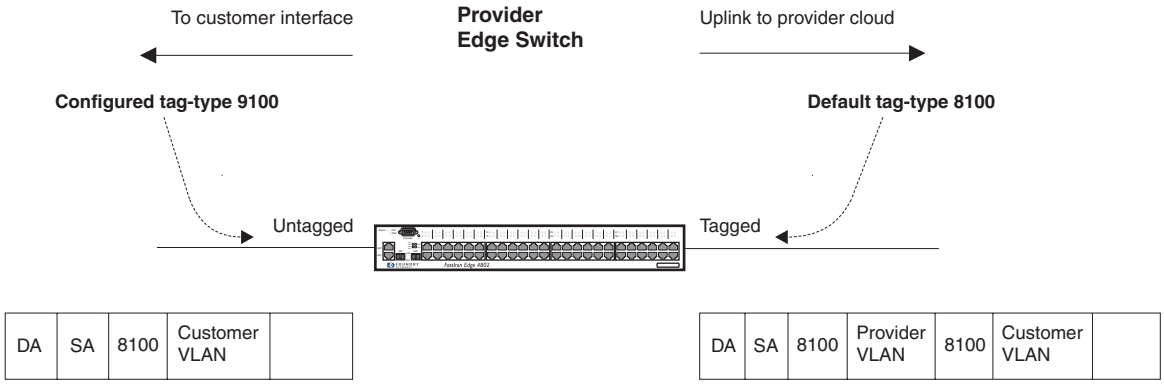

In Figure 4, the untagged ports (to customer interfaces) accept frames that have any 802.1Q tag other than the configured tag-type 9100. These packets are considered untagged on this incoming port and are re-tagged when they are sent out of the uplink towards the provider. The 802.1Q tag-type on the 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.

Configuring 802.1Q-in-Q

To enable the 802.1Q-in-Q feature, configure a 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 5, 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 5, 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 5 shows an example 802.1Q-in-Q configuration.

Figure 5 Example 802.1Q-in-Q Configuration

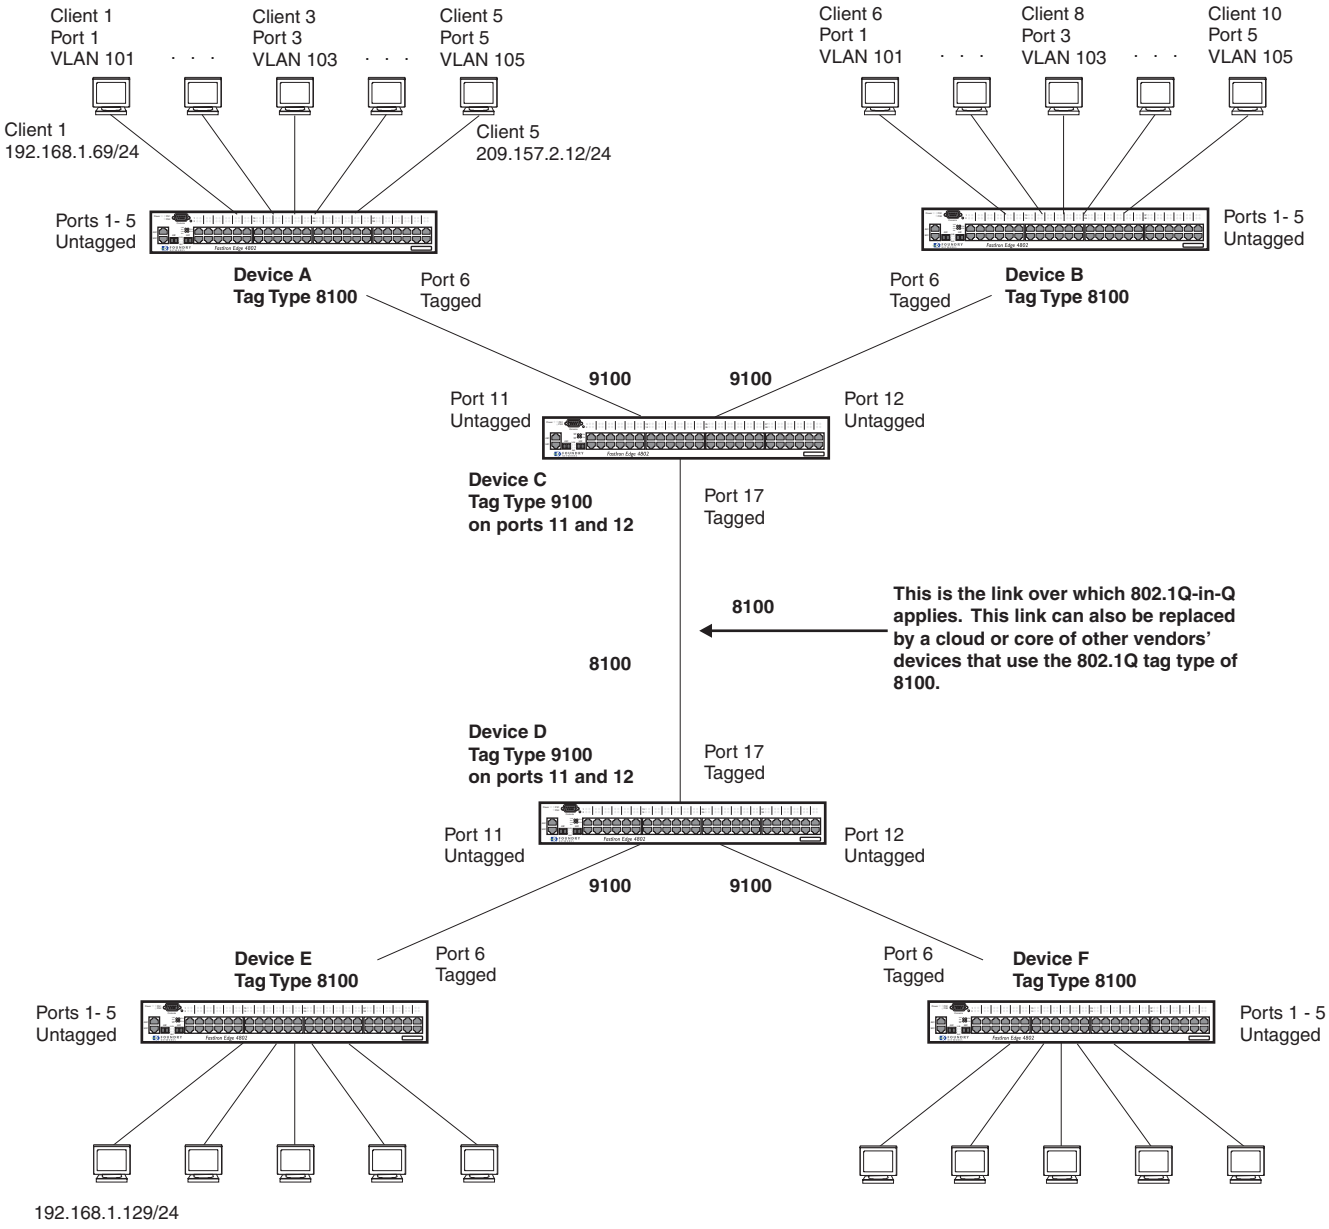

Topology Groups

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: 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. 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.

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.

For configuration information and application examples, see the *Foundry Switch and Router Installation and Basic Configuration Guide*.

VSRP

Virtual Switch Redundancy Protocol (VSRP) is a Foundry proprietary protocol that provides redundancy and subsecond 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.

NOTE: The FES devices with this release support full VSRP, as well as VSRP-awareness.

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.

You can use VSRP in combination with MRP to provide additional design flexibility.

For configuration and application examples, see the *Foundry Switch and Router Installation and Basic Configuration Guide*.

Layer 3 Enhancements in 03.1.00

This section describes the Layer 3 software enhancements in release 03.1.00.

Applying an OSPF Distribution List to an Interface

An OSPF distribution list filters specific OSPF routes from being installed in the IP route table. To configure an OSPF distribution list, you first configure a standard or extended ACL that identifies the routes to deny or permit, then configure an OSPF distribution list that applies the ACL to incoming route advertisements received on the specified interface.

In releases prior to 03.1.00, the Foundry device applies the OSPF distribution list to all incoming route updates.

In software releases 03.1.00 and later, you can optionally specify on which physical or virtual interfaces to apply an OSPF distribution list.

To apply an OSPF distribution list to an interface, enter commands such as the following:

```
FES4802(config)# router ospf
FES4802(config-ospf-router)# distribute-list acl1 in e1
```

Syntax: [no] distribute-list <acl-name> | <acl-id> in [<interface type>] [<interface number>]

This feature is disabled by default.

The <acl-name> parameter specifies the standard or extended ACL name that defines which network(s) to permit or deny.

The <acl-id> parameter specifies the standard or extended ACL number that defines which network(s) to permit or deny.

The **in** command applies the ACL to incoming route updates.

The <interface type> parameter identifies the interface type (i.e., **e** (ethernet) or **ve** (virtual)) on which to apply the ACL.

The <interface> parameter specifies the interface number on which to apply the ACL. Enter only one valid interface number. If necessary, use the **show interface brief** command to display a list of valid interfaces. If you do not specify an interface, the Foundry device applies the ACL to all incoming route updates.

System Level Enhancements in 03.1.00

Protected Link Groups

A protected link group minimizes disruption to the network by protecting critical links from loss of data and power. In a protected link group, one port in the group acts as the primary or active link, and the other ports act as secondary or standby links. The active link carries the traffic. If the active link goes down, one of the standby links takes over.

During normal operation, the active port in a protected link group is enabled and the standby ports are logically disabled. If the active port fails, the Foundry device immediately enables one of the standby ports, and switches traffic to the standby port. The standby port becomes the new, active port.

About Active Ports

When you create a protected link group, you can optionally specify which port in the protected link group is the active port. If you do not explicitly configure an active port, the Foundry device dynamically assigns one. A dynamic active port is the first port in the protected link group that comes up (usually the lowest numbered port in the group).

Static and dynamic active ports operate as follows:

- A static active port (an active port that you explicitly configured) preempts other ports in the protected link group. So, if a static active link comes back up after a failure, the Foundry device will revert to this link as the active link.

- A dynamic active port (an active port assigned by the software) is non-preemptive. Therefore, if a dynamic active link comes back up after a failure, the Foundry device does not revert to this link, but continues carrying traffic on the current active link.

Using UDLD with Protected Link Groups

You can use Uni-directional Link Detection (UDLD) with protected link groups to detect uni-directional link failures and to improve the speed at which the device detects a failure in the link. Use UDLD and protected link groups simultaneously when the FastIron Edge Switch is connected to a device with slower link detection times.

Configuration Notes

- You can configure a maximum of 32 protected link groups.
- There is no restriction on the number of ports in a protected link group.
- Each port can belong to one protected link group at a time.
- There is no restriction on the type of ports in a protected link group. A protected link group can consist of Gigabit fiber ports, 10/100/1000 copper ports, and 10/100 ports, or any combination thereof.
- There is no restriction on the properties of ports in a protected link group. For example, member ports can be in the same VLAN or in different VLANs.

Creating a Protected Link Group and Assigning an Active Port

To create a protected link group:

1. Specify the member ports in the protected link group. Enter a command such as the following:

   ```
   FastIron 4802(config)# protected-link-group 10 e 1 to 4
   ```

2. Optionally specify which port will be the active port for the protected link group. Enter a command such as the following:

   ```
   FastIron 4802(config)# protected-link-group 10 active-port e 1
   ```

   NOTE: If you do not explicitly configure an active port, the Foundry device automatically assigns one as the first port in the protected link group to come up.

These commands configure port e1 as the active port and ports e2 – e4 as standby ports. If port 1 goes down, the Foundry device enables the first available standby port, and switches the traffic to that port. Since the above configuration consists of a statically configured active port, the active port preempts other ports in the protected link group. See “About Active Ports” on page 22.

Syntax: [no] protected-link-group <group-ID> ethernet <portnum> to <portnum>

The <group-ID> parameter specifies the protected link group number. Enter a number from 1 – 32.

Each **ethernet** parameter introduces a port group.

The <portnum> **to** <portnum> parameters specify the ports in the protected link group.

Syntax: [no] protected-link-group <group-ID> active-port ethernet <portnum>

The <group-ID> parameter specifies the protected link group number. Enter a number from 1 – 32.

The **active-port ethernet** <portnum> parameter defines the active port.

Viewing Information about Protected Link Groups

You can view protected link group information using the **show protected-link-group** command. The following shows an example output.

```
FES4802# show protected-link-group

Group ID: 1
Member Port(s): ethe 1 to 7
Configured Active Port: 7
Current Active Port: 7
Standby Port(s): ethe 5

Total Number of Protected Link Groups: 1
```

Syntax: show protected-link-group [group-ID]

Table 1: CLI Display of Protected Link Group Information

This Field...	Displays...
Group ID	The ID number of the protected link group.
Member Port(s)	The ports that are members of the protected link group.
Configured Active Port	The statically configured active port. If you do not statically configure an active port, this value will be "None".
Current Active Port	The current active port for the protected link group. If all member ports are down, this value will be "None".
Standby Port(s)	The member ports that are on standby.

Specifying a Port for Web Management Access

The following command allows you to specify the TCP port that will be used to access a device's Web management interface.

```
FES4802(config)# web-management tcp-port 168
```

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

The **tcp-port** <port number> parameter 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.

MIB for Web Management on a TCP Port

Name, OID, and Syntax	Access	Description
snAgWebMgmtServerTcpPort 1.3.6.1.4.1.1991.1.1.2.1.63 Syntax: Integer	Read-write	This object allows you to specify which TCP port will be used for the Web management interface. Also, Element Manager of IronView Network Manager will query the device for this port number before it sends HTTP packets to the device. Enter a number from 1 – 65535.

Detecting POE Power Requirements Advertised via CDP

Many power consuming devices, such as Cisco's VOIP phones and other vendors' devices, use CDP to advertise their power requirements to power sourcing devices, such as Foundry's FES2402-POE and FES4802-POE.

- In releases prior to 03.1.00, the FES2402-POE and FES4802-POE are unable to detect and process power requirements for power consuming devices that use CDP to advertise their power requirements.
- In releases 03.1.00 and later, the FES2402-POE and FES4802-POE are compatible with Cisco's and other vendors' power consuming devices, in that they can detect and process power requirements for these devices automatically.

Configuration Considerations

- If you configure a port with a power level for a power consuming device, the power level takes precedence over the CDP power requirement. Therefore, if you want the device to adhere to the CDP power requirement, do not configure a power level on the port.
- If you configure a port with a power class and the Foundry device receives a CDP power requirement from a power consuming device, it adjusts the power on the port to comply with the CDP power requirement.
- The FES2402-POE and FES4802-POE adjust a port's power only if there are available power resources on the device.

Configuring the POE Device to Detect CDP Power Requirements

To enable the FES2402-POE and FES4802-POE to detect CDP power requirements, enter the following commands:

```
FES4802# config t
FES4802(config)# cdp run
```

Syntax: [no]cdp run

Specifying the Power Level for a POE Power Consuming Device

By default, each port on the FES2402-POE and FES4802-POE provides 15.4 watts of power to each POE power consuming device connected to the switch. In releases prior to 03.1.00, you cannot change the default value.

In software releases 03.1.00 and later, you can configure the amount of power that a port will provide to a power consuming device. You can specify from 1 to 15.4 watts of power for each device connected to the switch.

NOTE: There are two ways to configure the power level for a POE power consuming device. The first method is discussed in this section. The other method is provided in the section "Specifying the Power Class for a POE Power Consuming Device" on page 26. For each POE port, you can configure either a power level or a power class. You cannot configure both. You can, however, configure a power level on one port and power class on another port.

The commands for this feature differ depending on whether or not the device is 802.3af-compliant. Refer to the appropriate section, below.

802.3af-Compliant Device

To configure the power level for an 802.3af-compliant device, enter commands such as the following:

```
FES4802 Router# config t
FES4802 Router(config)# interface e 1
FES4802 Router(config-if-e100-1)# inline power configurepower 14000
```

These commands enable in-line power on interface e 1 and set the POE power level to 14,000 milliwatts (14 watts).

Syntax: inline power configurepower <power level>

where <power level> is the number of milliwatts, between 1000 and 15400.

802.3af Non-Compliant Device

To configure the power level for a device that is not 802.3af-compliant (for example, a legacy device such as a Cisco VOIP phone), enter commands such as the following:

```
FES4802 Router# config t
FES4802 Router(config)# interface e 1
FES4802 Router(config-if-e100-1)# inline power legacy-powerdevice configurepower
7000
```

These commands enable in-line power on interface e 1 and set the POE power level to 7000 milliwatts (7 watts).

Syntax: inline power legacy-powerdevice configurepower <power level>

where <power level> is the number of milliwatts, between 1000 and 15400.

Specifying the Power Class for a POE Power Consuming Device

A power class specifies the amount of power that an FES2402-POE or FES4802-POE will supply to a power consuming device. Table 2 shows the different power classes and their respective power consumption needs.

Table 2: Power Classes for Powered Devices

Class	Power (Watts)
0	15.4 (default)
1	4
2	7
3	15.4

By default, the power class for all power consuming devices is zero (0). As shown in Table 2, a power consuming device with a class of 0 receives 15.4 watts of power.

- In releases prior to 03.1.00, you cannot change the default power class.
- In releases 03.1.00 and later, you can configure the power class for a power consuming device.

NOTE: The power class sets the power level for a power consuming device. You can also set the power level as instructed in the section “Specifying the Power Level for a POE Power Consuming Device” on page 25. For each POE port, you can configure either a power class or a power level. You cannot configure both. You can, however, configure a power level on one port and power class on another port.

The commands for this feature differ depending on whether or not the device is 802.3af-compliant. Refer to the appropriate section, below.

802.3af-Compliant Devices

To configure the power class for an 802.3af-compliant device, enter commands such as the following:

```
FES4802 Router# config t
FES4802 Router(config)# interface e 1
FES4802 Router(config-if-e100-1)# inline power class 2
```

These commands enable in-line power on interface e 1 and set the power class to 2.

Syntax: inline power class <class value>

where <class value> is the power class. This number can be from 0 – 3, as shown in Table 2.

802.3af Non-Compliant Devices

To configure the power class for a device that is not 802.3af-compliant (for example, a legacy device such as a Cisco VOIP phone), enter commands such as the following:

```
FES4802 Router# config t
FES4802 Router(config)# interface e 1
FES4802 Router(config-if-e100-1)# inline power legacy-powerdevice class 0
```

These commands enable in-line power on interface e 1 and set the power class to 0.

Syntax: inline power legacy-powerdevice class <class value>

where <class value> is the power class. This number can be from 0 – 3, as shown in Table 2.

SNMP MIB Object for Power Level and Power Class for a POE Power Consuming Device

The following MIB objects have been added for the POE Power Consuming Device power level and power class.

Name, OID, and Syntax	Access	Description
snSwPortInLinePowerControl 1.3.6.1.4.1.1991.1.1.3.3.1.1.50 Syntax: Integer	Read-write	Controls inline power on/off to a port. Valid values: • other(1) – The port does not have inline power capability, • disable(2) – The device is a 802.3af-compliant device and the inline power capability on this port is disabled. • enable(3) – The device is a 802.3af-compliant device and the inline power capability on this port is enabled. • enableLegacyDevice(4) – This device is non-802.3af-compliant and the inline power capability on this port is enabled.
snSwPortInLinePowerWattage 1.3.6.1.4.1.1991.1.1.3.3.1.1.59 Syntax: Integer	Read-write	Adjust the inline power wattage. Each unit is milliwatts. This object can only be set aftersnSwPortInLinePowerControl object has been set to enable(3) or enableLegacyDevice(4). If a port does not have inline power capability, reading this object returns undefined value. Valid values: 1000 – 15400 milliwatts
snSwPortInLinePowerClass 1.3.6.1.4.1.1991.1.1.3.3.1.1.60 Syntax: Integer	Read-write	Adjust the inline power class. This object can only be set after snSwPortInLinePowerControl has been set to 'enable(3)'or 'enableLegacyDevice(4)'. If a port does not have inline power capability, reading this object returns undefined value. Valid values: • 0 – 15.4 • 1 – 4 • 2 – 7 • 3 – 15.4 Default: 0

Dynamically Applying IP ACLs and MAC Filters to 802.1X Ports

In releases 03.1.00 and later, Foundry's 802.1X implementation supports dynamically applying an IP ACL or MAC address filter to a port, based on information received from an Authentication Server.

When a client/supplicant successfully completes the Extensible Authentication Protocol (EAP) authentication process, the Authentication Server (the RADIUS server) sends the Authenticator (the Foundry device) a RADIUS Access-Accept message that grants the client access to the network. The RADIUS Access-Accept message contains attributes set for the user in the user's access profile on the RADIUS server.

If the Access-Accept message contains Filter-ID (type 11) and/or Vendor-Specific (type 26) attributes, the Foundry device can use information in these attributes to apply an IP ACL or MAC address filter to the authenticated port. This IP ACL or MAC address filter applies to the port for as long as the client is connected to the network. When the client disconnects from the network, the IP ACL or MAC address filter is no longer applied to the port. If an IP ACL or MAC address filter had been applied to the port prior to 802.1X authentication, it is then re-applied to the port.

The Foundry device uses information in the Filter ID and Vendor-Specific attributes as follows:

- The Filter-ID attribute can specify the number of an existing IP ACL or MAC address filter configured on the Foundry device. In this case, the IP ACL or MAC address filter with the specified number is applied to the port.
- The Vendor-Specific attribute can specify actual syntax for a Foundry IP ACL or MAC address filter, which is then applied to the authenticated port. Configuring a Vendor-Specific attribute in this way allows you to create IP ACLs and MAC filters that apply to individual users; that is, ***per-user*** IP ACLs or MAC address filters.

Configuration Considerations

The following restrictions apply to dynamic IP ACLs or MAC address filters:

- A maximum of one IP ACL can be configured in the inbound direction on an interface
- A maximum of one IP ACL can be configured in the outbound direction on an interface
- MAC address filters cannot be configured in the outbound direction on an interface

Disabling and Enabling Strict Security Mode for Dynamic Filter Assignment

By default, 802.1X dynamic filter assignment operates in ***strict security mode***. When strict security mode is enabled, 802.1X authentication for a port fails if the Filter-ID attribute contains invalid information, or if insufficient system resources are available to implement the per-user IP ACLs or MAC address filters specified in the Vendor-Specific attribute.

When strict security mode is enabled:

- If the Filter-ID attribute in the Access-Accept message contains a value that does not refer to an existing filter (that is, a MAC address filter or IP ACL configured on the device), then the port will not be authenticated, regardless of any other information in the message (for example, if the Tunnel-Private-Group-ID attribute specifies a VLAN to which to assign the port).
- If the Vendor-Specific attribute specifies the syntax for a filter, but there are insufficient system resources to implement the filter, then the port will not be authenticated.
- If the device does not have the system resources available to dynamically apply a filter to a port, then the port will not be authenticated.

NOTE: If the Access-Accept message contains values for both the Filter-ID and Vendor-Specific attributes, then the value in the Vendor-Specific attribute (the per-user filter) takes precedence.

Also, if authentication for a port fails because the Filter-ID attribute referred to a non-existent filter, or there were insufficient system resources to implement the filter, then a Syslog message is generated. See Table 3 on page 33 for a description of this message.

When strict security mode is disabled:

- If the Filter-ID attribute in the Access-Accept message contains a value that does not refer to an existing filter (that is, a MAC address filter or IP ACL configured on the device), then the port is still authenticated, but no filter is dynamically applied to it.
- If the Vendor-Specific attribute specifies the syntax for a filter, but there are insufficient system resources to implement the filter, then the port is still authenticated, but the filter specified in the Vendor-Specific attribute is not applied to the port.

By default, strict security mode is enabled for all 802.1X-enabled interfaces, but you can manually disable or enable it, either globally or for specific interfaces.

To disable strict security mode globally, enter the following commands:

```
FES4802(config)# dot1x-enable
FES4802(config-dot1x)# no global-filter-strict-security
```

After you have globally disabled strict security mode on the device, you can re-enable it by entering the following command:

```
FES4802(config-dot1x)# global-filter-strict-security
```

Syntax: [no] global-filter-strict-security

To disable strict security mode for a specific interface, enter commands such as the following:

```
FES4802(config)# interface e 1
FES4802(config-if-1)# no dot1x filter-strict-security
```

To re-enable strict security mode for an interface, enter the following command:

```
FES4802(config-if-1)# dot1x filter-strict-security
```

Syntax: [no] dot1x filter-strict-security

The output of the **show dot1x** and **show dot1x config** commands has been enhanced to indicate whether strict security mode is enabled or disabled globally and on an interface.

To display the status of strict security mode globally on the device, enter the following command:

```
FES4802# show dot1x
PAE Capability:    Authenticator Only
system-auth-control: Enable
re-authentication: Disable
global-filter-strict-security: Enable
quiet-period:    60 Seconds
tx-period:    30 Seconds
supptimeout:    30 Seconds
servertimeout:     30 Seconds
maxreq:    2
re-authperiod:    3600 Seconds
security-hold-time: 60 Seconds
Protocol Version:    1
```

Syntax: show dot1x

To display the status of strict security mode on an interface, enter a command such as the following:

```
FES4802# show dot1x config e 3

Port 3 Configuration:
Authenticator PAE state:    AUTHENTICATED
Backend Authentication state:     IDLE
AdminControlledDirections:    BOTH
OperControlledDirections:    BOTH
AuthControlledPortControl:     Auto
```

```
AuthControlledPortStatus:    authorized
quiet-period:    60 Seconds
tx-period:    30 Seconds
supptimeout:    30 Seconds
servertimeout:     30 Seconds
maxreq:    2
re-authperiod:    3600 Seconds
security-hold-time: 60 Seconds
re-authentication: Disable
multiple-hosts: Disable
filter-strict-security: Enable
Protocol Version:    1
```

Syntax: show dot1x config <portnum>

Dynamically Applying Existing ACLs or MAC Address Filters

When a port is authenticated using 802.1X security, an IP ACL or MAC address filter that exists in the running-config on the Foundry device can be dynamically applied to the port. To do this, you configure the Filter-ID (type 11) attribute on the RADIUS server. The Filter-ID attribute specifies the name or number of the Foundry IP ACL or MAC address filter, as well as whether to apply it in the inbound or outbound direction.

The following is the syntax for configuring the Filter-ID attribute to refer to a Foundry IP ACL or MAC address filter:

Value	Description
ip.<number>.in	Applies the specified numbered ACL to the 802.1X authenticated port in the inbound direction.
ip.<name>.in	Applies the specified named ACL to the 802.1X authenticated port in the inbound direction.
ip.<number>.out	Applies the specified numbered ACL to the 802.1X authenticated port in the outbound direction.
ip.<name>.out	Applies the specified named ACL to the 802.1X authenticated port in the outbound direction.
mac.<number>.in	Applies the specified numbered MAC address filter to the 802.1X authenticated port in the inbound direction.

The following table lists examples of values you can assign to the Filter-ID attribute on the RADIUS server to refer to IP ACLs and MAC address filters configured on a Foundry device.

Possible Values for the Filter ID attribute on the RADIUS server	ACL or MAC address filter configured on the Foundry device
ip.2.in ip.2.out	access-list 2 permit host 36.48.0.3 access-list 2 permit 36.0.0.0 0.255.255.255
ip.102.in ip.102.out	access-list 102 permit ip 36.0.0.0 0.255.255.255 any
ip.fdry_filter.in ip.fdry_filter.out	ip access-list standard fdry_filter permit host 36.48.0.3
mac.2.in	mac filter 2 permit 3333.3333.3333 ffff.ffff.ffff any etype eq 0800

Possible Values for the Filter ID attribute on the RADIUS server	ACL or MAC address filter configured on the Foundry device
mac.2.in	mac filter 2 permit 3333.3333.3333 ffff.ffff.ffff any etype eq 0800
mac.3.in	mac filter 3 permit 2222.2222.2222 ffff.ffff.ffff any etype eq 0800

Notes

- The <name> in the Filter ID attribute is case-sensitive.
- You can specify only numbered MAC address filters in the Filter ID attribute. Named MAC address filters are not supported.
- MAC address filters are supported only for the inbound direction. Outbound MAC address filters are not supported.
- Dynamically assigned IP ACLs and MAC address filters are subject to the same configuration restrictions as non-dynamically assigned IP ACLs and MAC address filters. See the *Foundry Enterprise Configuration and Management Guide* for more information.
- Multiple IP ACLs and MAC address filters can be specified in the Filter ID attribute, allowing multiple filters to be simultaneously applied to an 802.1X authenticated port. Use commas, semicolons, or carriage returns to separate the filters (for example: ip.3.out,mac.2.in).

Restrictions on Configuring Both MAC Address Filters and IP ACLs on the FastIron Edge Switch

When 802.1X dynamic filter/ACL assignment is configured on the FES, the following restrictions apply:

- If a MAC address filter is applied on any interface of the FES statically (using the CLI) or dynamically (through 802.1X), an IP ACL cannot be statically or dynamically applied in the outbound direction on any interface on the FES.
- If an IP ACL is applied in the outbound direction on any interface on the FES statically or dynamically, no MAC address filters can be applied statically or dynamically on any interface of the FES.
- If a MAC address filter is applied on an interface statically or dynamically, then neither an inbound nor outbound IP ACL can be applied statically or dynamically on the same interface.
- If an inbound or outbound IP ACL is applied to an interface statically or dynamically, then no MAC address filters can be statically or dynamically applied on the same interface.

Configuring Per-User IP ACLs or MAC Address Filters

Per-user IP ACLs and MAC address filters make use of the Vendor-Specific (type 26) attribute to dynamically apply filters to ports. Defined in the Vendor-Specific attribute are Foundry ACL or MAC address filter statements. When the RADIUS server returns the Access-Accept message granting a client access to the network, the Foundry device reads the statements in the Vendor-Specific attribute and applies these IP ACLs or MAC address filters to the client's port. When the client disconnects from the network, the dynamically applied filters are no longer applied to the port. If any filters had been applied to the port previous to the client connecting, then those filters are reapplied to the port.

The following is the syntax for configuring the Foundry Vendor-Specific attribute with ACL or MAC address filter statements:

Value	Description
ipacl.e.in=<extended-acl-entries>	Applies the specified extended ACL entries to the 802.1X authenticated port in the inbound direction.
ipacl.e.out=<extended-acl-entries>	Applies the specified extended ACL entries to the 802.1X authenticated port in the outbound direction.

Value	Description
macfilter.in=<mac-filter-entries>	Applies the specified MAC address filter entries to the 802.1X authenticated port in the inbound direction.

The following table shows examples of IP ACLs and MAC address filters configured in the Foundry Vendor-Specific attribute on a RADIUS server. These IP ACLs and MAC address filters follow the same syntax as other Foundry ACLs and MAC address filters. See the *Foundry Enterprise Configuration and Management Guide* for information on syntax.

ACL or MAC address filter	Vendor-Specific attribute on RADIUS server
Extended ACL with one entry (outbound direction)	ipacl.e.out=deny ip 1.1.1.1 0.0.255.255 20.20.20.20 255.255.0.0
MAC address filter with one entry	macfilter.in= deny any any
MAC address filter with two entries	macfilter.in= permit 0000.0000.3333 ffff.ffff.0000 any, macfilter.in= permit 0000.0000.4444 ffff.ffff.0000 any

The RADIUS server allows one instance of the Vendor-Specific attribute to be sent in an Access-Accept message. However, the Vendor-Specific attribute can specify multiple IP ACLs or MAC address filters. You can use commas, semicolons, or carriage returns to separate the filters (for example: ipacl.e.in= permit ip any any,ipacl.e.in = deny ip any any).

Displaying Information About Dynamically Applied MAC Filters and IP ACLs

You can display information about the user-defined and dynamically applied MAC filters and IP ACLs currently active on the device.

Displaying User-Defined MAC Filters and IP ACLs

To display the user-defined MAC filters active on the device, enter the following command:

```
FES4802# show dot1x mac-address-filter

Port 3 (User defined MAC Address Filter):
    mac filter 1 permit any any
```

Syntax: show dot1x mac-address-filter

To display the user-defined IP ACLs active on the device, enter the following command:

```
FES4802# show dot1x ip-acl

Port 3 (User defined IP ACLs):

Extended IP access list Port_3_E_IN
permit udp any any

Extended IP access list Port_3_E_OUT
permit udp any any
```

Syntax: show dot1x ip-acl

Displaying Dynamically Applied MAC Filters and IP ACLs

To display the dynamically applied MAC address filters active on an interface, enter a command such as the following:

```
FES4802# show dot1x mac-address-filter e 3

Port 3 MAC Address Filter information:
  802.1X Dynamic MAC Address Filter :
     mac filter-group 2
  Port default MAC Address Filter:
     No mac address filter is set
```

Syntax: show dot1x mac-address-filter <portnum> | all

The **all** keyword displays all dynamically applied MAC address filters active on the device.

To display the dynamically applied IP ACLs active on an interface, enter a command such as the following:

```
FES4802# show dot1x ip-acl e 3

Port 3 IP ACL information:
  802.1X dynamic IP ACL (user defined) in:
    ip access-list extended Port_3_E_IN in
  Port default IP ACL in:
    No inbound ip access-list is set
  802.1X dynamic IP ACL (user defined) out:
    ip access-list extended Port_3_E_OUT out
  Port default IP ACL out:
    No outbound ip access-list is set
```

Syntax: show dot1x ip-acl <portnum> | all

The **all** keyword displays all dynamically applied IP ACLs active on the device.

New Syslog Message for 802.1X Dynamically Applied ACLs and MAC Filters

The following 802.1X-related Syslog message was added in release 03.1.00.

Table 3: Syslog Message for 802.1X

Message Level	Message	Explanation
Information	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: • 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)

Configuration Notes

For a few features, configuration procedures or defaults on the FastIron Edge Switches differ from those on BigIron Chassis devices or the FastIron 4802 because of the FastIron Edge Switches' hardware architecture. The following sections describe the differences.

Port Numbers

Many commands and displays use port numbers. The ports are labeled on the front panel.

The CLI examples in the manuals listed in "Where To Get More Information" on page 53 use chassis-based port numbering. When you enter commands on a FastIron Edge Switch, just specify the port number. The slot numbers used in the BigIron CLI examples apply only to Chassis devices.

Here is an example. The following commands change the CLI from the global CONFIG level to the configuration level for the first port on the device.

- BigIron commands:

```
BigIron(config)# interface e 1/1
BigIron(config-if-e100-1/1)#
```

- FastIron Edge Switch commands:

```
FES4802 Router(config)# interface e 1
FES4802 Router(config-if-e100-1)#
```

Enabling Power over Ethernet

To enable a port to provide in-line power for power consuming devices, you must configure the port using one of the following CLI commands.

To enable PoE for 802.3af-compliant devices, enter commands such as the following:

```
FES4802 Router# config t
FES4802 Router(config)# interface e 1
FES4802 Router(config-if-e100-1)# inline power
```

To enable PoE for devices that are not 802.3af-compliant (for example, legacy devices such as Cisco VOIP phones), enter commands such as the following:

```
FES4802 Router# config t
FES4802 Router(config)# interface e 1
FES4802 Router(config-if-e100-1)# inline power legacy-powerdevice
```

Syntax: inline power [legacy-powerdevice]

where [legacy-powerdevice] is required for devices other than 802.3af-compliant devices.

The following legacy devices are currently supported on the FES POE devices. Other legacy devices may have been tested with the FES2402 and FES4802 after the release of this document. Contact your Foundry account representative about installing legacy devices that are not included in this list.

Table 4: Legacy Devices Supported on the FES2402-POE and FES4802-POE*

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

Table 4: Legacy Devices Supported on the FES2402-POE and FES4802-POE*

Legacy Device	Firmware

* 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.

Displaying Power over Ethernet Information

The **show inline power** command displays information about Power over Ethernet. This command also shows the total power supply capacity and how much power is available for power-consuming devices.

To display the statistics, enter the following command:

```
FES4802 Router# show inline power

Power Supply : total capacity is 480000 of which 206200 is currently available
power has been successfully allocated 20 times and rejected 0 times

Port         Detection          Class         Power Enable      Power
-----------------------------------------------------------------------
01           NO-PD              Unknown          OFF            No
02           802.3AF-PD         Class1           ON             Good
03           802.3AF-PD         Class0           ON             Good
04           802.3AF-PD         Class0           ON             Good
05           802.3AF-PD         Class0           ON             Good
06           802.3AF-PD         Class0           ON             Good
07           LEGACY             Class1           ON             Good
08           LEGACY             Unknown          ON             Good
09           802.3AF-PD         Class1           ON             Good
10           802.3AF-PD         Class0           ON             Good
11           NO-PD              Unknown          OFF            No
12           NO-PD              Unknown          OFF            No
13           NO-PD              Unknown          OFF            No
14           802.3AF-PD         Class0           ON             Good
15           LEGACY             Unknown          ON             Good
```

Syntax: show inline power

Table 5 provides definitions for the statistics.

Table 5: Power over Ethernet Statistics

This Column...	Displays...
Port	The port number.

Table 5: Power over Ethernet Statistics (Continued)

This Column...	Displays...
Detection	The type of powered device connected to the port. This value can be one of the following: • 802.3AF-PD – The powered device connected to this port is 802.3af-compliant. • LEGACY – The powered device connected to this port is a legacy product (not 802.3af-compliant). • Open – Power over Ethernet is configured on this port; however, there is no device connected to this port. • NO-PD – Power over Ethernet is configured on this port, and the device connected to this port is a non-powered device, or there is no device connected to this port.
Class	Determines the maximum amount of power a powered device receives. This value can be one of the following: • 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 – The device attached to the port cannot advertise its class.
Power Enable	Shows whether or not Power over Ethernet is enabled on the port. This value can be one of the following: • ON – This port has been configured to provide inline power. • OFF – This port has not been configured to provide inline power.
Power	Shows the status of the power provided to the powered device. This value can be one of the following: • NO – The port is not providing inline power. • Good – The port is providing inline power and is functioning normally.

Layer 2 MAC Filtering Differences

The Layer 2 MAC filtering on the FastIron Edge Switch is performed in hardware. Layer 2 MAC filtering on the FastIron Edge Switch differs from other Foundry devices in that you can only filter on source and destination MAC addresses. Other Foundry devices allow you to also filter on the encapsulation type and frame type.

Dynamic Link Aggregation Differences

The dynamic link aggregation (802.3ad) implementation on the FastIron Edge Switch allows any number of ports up to eight to be aggregated into a link. The feature does not require the aggregate link to consist of exactly 2, 4, or 8 ports.

NOTE: The trunks that will be formed by link aggregation will strictly adhere to the static trunking rules on the FastIron Edge Switch. Be careful in selecting keys if you are manually configuring link aggregation keys. Make sure that the possible trunks that you expect to be formed conform to the static trunking rules.

Trunking Differences

FastIron Edge Switches support switch trunk groups and server trunk groups.

- A switch trunk group is a logical link made of multiple physical links between two networking switches, such as two FastIron Edge Switches or a FastIron Edge Switch and a BigIron Chassis device.
- A server trunk group is a logical link on which traffic is load balanced across the ports in the link. Server trunking load balances Layer 2 switched IP and IPX traffic among the ports in a trunk group. The device load balances the traffic by evenly distributing the traffic flows (source IP or IPX address plus destination IP or IPX address) across the trunk ports.

  NOTE: Layer 2 and Layer 3 AppleTalk traffic is not load-balanced. Layer 3 routed IP or IPX traffic also is not load balanced. These traffic types will however still be forwarded on the trunk ports.

Table 6 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 6: 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

Table 7 lists the valid trunk group configurations.

Table 7: FastIron Edge Switch Trunk Group Support

Model	Maximum Number of Groups		Valid Number of Ports in a Group		Port Ranges	
	10/100	Gigabit	10/100	Gigabit	10/100	Gigabit
9604	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

Table 7: FastIron Edge Switch Trunk Group Support (Continued)

Model	Maximum Number of Groups		Valid Number of Ports in a Group		Port Ranges	
	10/100	Gigabit	10/100	Gigabit	10/100	Gigabit
4802 and 4802-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
2402 and 2402-POE	3	1	2, 3, 4, 5, 6, 7, or 8	2	1 – 8, 9 – 16, 17 – 24	25 – 26
12GCF	N/A	6	N/A	2 – 8	N/A	N/A

Configuration Rules

- You cannot combine 10/100 ports and Gigabit ports in the same trunk group.
- All ports in a trunk group must be in the same port range. See Table 7. Note that this restriction does not apply to the FES12GCF, which does not have port ranges.
- With one exception, 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 can select any port within a range to be the first member of the trunk group.
- The ports must be consecutive. The group's port range cannot contain a gap. For example, if you select port 4 as the first port, the next port must be 5, followed by 6, and so on.
- A 10/100 trunk group can contain 2, 3, 4, 5, 6, 7, or 8 ports. The device does not restrict the number of ports to 2, 4, or 8. (Some other Foundry devices do restrict trunk groups to these numbers of ports.) Gigabit trunk groups can contain 2 ports on a 2402 or 4802, and 2, 3, or 4 ports on a 9604. Gigabit trunk groups on the FES12GCF can contain 2, 3, 4, 5, 6, 7, or 8 ports.
- 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 and the other end is a different type of switch, make sure the other switch can support a five-port trunk group.
- Server trunk groups load balance Layer 2 switched IP and IPX traffic. Other traffic is forwarded the same as on a switch trunk.

Configuration Syntax

To configure a trunk group, enter commands such as the following:

```
FES4802 Router(config)# trunk ethernet 1 to 8
FES4802 Router(config)# trunk deploy
```

The first command configures a switch trunk group consisting of 10/100 ports 1 – 8. The **trunk deploy** command activates the new trunk group. The trunk group enters the configuration but doesn't become active until you enter the **trunk deploy** command or until you save the configuration (write memory), then reload the software.

Syntax: [no] trunk [server | switch] ethernet <primary-portnum> to <portnum>

NOTE: The **server** option is supported only in release 02.0.00 and later.

Syntax: trunk deploy

The following example configures another 10/100 trunk group:

```
FES4802 Router(config)# trunk ethernet 12 to 16
FES4802 Router(config)# trunk deploy
```


The **trunk** command in this example configures a group that starts with port 12 and contains ports 12, 13, 14, 15, and 16.

The following example shows how to configure a trunk group consisting of all four Gigabit ports on a 9604:

```
FES4802 Router(config)# trunk ethernet 97 to 100
FES4802 Router(config)# trunk deploy
```

To configure a server trunk group containing two Gigabit ports on a 9604, enter a command such as the following:

```
FES4802 Router(config)# trunk server ethernet 98 to 99
FES4802 Router(config)# trunk deploy
```

Displaying Trunk Group Information

To display trunk group information, enter the following command:

Syntax: show trunk [ethernet <portnum> to <portnum>]

See the "Displaying Trunk Group Configuration Information" section in the "Configuring Trunk Groups and Dynamic Link Aggregation" chapter of the *Foundry Switch and Router Installation and Basic Configuration Guide*.

Fixed Rate Limiting Differences

NOTE: The fixed rate limiting feature on the FastIron Edge Switch is completely different from the rate limiting features on the BigIron Chassis device or FastIron 4802. Do not attempt to configure the feature using the information in the "Configuring IronClad Rate Limiting (IronCore)" chapter or "JetCore Adaptive Rate Limiting" chapter of the *Foundry Enterprise Configuration and Management Guide*.

Software releases 03.1.00 and later support fixed rate limiting for inbound traffic, on individual ports. The fixed rate limiting is at line rate and occurs in hardware. Fixed rate limiting allows you to specify the maximum number of bits per second (bps) a port can receive. The port drops all traffic that exceeds the specified bps within a given one-second interval. Fixed rate limiting applies to all traffic on the rate limited port.

The rate you specify applies to each one-second interval. All traffic that exceeds the specified rate within a one-second interval is dropped. Unused bandwidth is not carried over from one interval to the next.

Configuring Rate Limiting

To configure rate limiting on a port, enter commands such as the following:

```
FES4802 Router(config)# interface ethernet 48
FES4802 Router(config-if-e100-48)# rate input fixed 10000000
```

These commands limit the average rate for inbound traffic on port 48 to 10,000,000 bps.

Syntax: [no] rate input fixed <average-rate> [payload-only]

The <average-rate> parameter specifies the number of bits per second (bps) the port can receive. The minimum rate that can be configured is 240,000 bps.

By default, rate limiting is optimized for packets that are 256 bytes in size. This packet size includes 14 bytes of Layer 2 header (Ethernet II untagged) and 4 bytes of Layer 2 CRC.

To optimize rate limiting for all packet sizes, use the **payload-only** parameter. If this parameter is specified, then the system excludes Layer 2 header and Layer 2 checksum (CRC) from the calculations, and the rate is accurate for all packet sizes and Layer 2 overhead (Layer 2 header + CRC). Layer 2 overhead for different encapsulations is as follows:

- Untagged Ethernet-II – 18 bytes
- Tagged Ethernet-II – 22 bytes
- LLC over Untagged Ethernet-II – 21 bytes
- LLC over Tagged Ethernet-II – 25 bytes
- LLC/SNAP over Untagged Ethernet-II – 26 bytes

- LLC/SNAP over Tagged Ethernet-II – 30 bytes

Displaying the Fixed Rate Limiting Configuration

To display the fixed rate limiting configuration on the device, enter the following command:

```
FES4802 Switch(config-if-e100-21)#show rate-limit fixed

Total rate-limited interface count: 11.

 Port    Configured Input Rate        Actual Input Rate             Mode
    1                  1000000                  1000000      Payload-Only
    3                 10000000                 10005000           Default
    7                 10000000                 10000000      Payload-Only
    9                  7500000                  7502000      Payload-Only
   11                  8000000                  7999000           Default
   12                  8000000                  7999000           Default
   13                  8000000                  7999000           Default
   14                  8000000                  7999000           Default
   15                  8000000                  7999000           Default
   21                  8000000                  8000000      Payload-Only
   25                  7500000                  7502000           Default
```

Syntax: show rate-limit fixed

The command lists the number of ports on which fixed rate limiting is configured, then provides the following information for each of the ports:

- The **Configured Input Rate** is the rate requested in the configuration.
- The **Actual Input Rate** is the rate provided by the hardware for the request.
- **Mode** will indicate if the payload-only option has been specified.

Broadcast, Unknown-Unicast, and Multicast Rate Limiting Differences

Configuration of the broadcast/unknown-unicast and multicast rate limiting on the FastIron Edge Switch is different from configuration of the comparable feature on the FastIron 4802 and BigIron Chassis device. On the FastIron Edge Switch, you can configure a single maximum rate for broadcasts and unknown-unicasts combined. You also can add multicasts to the combined maximum rate. On the FastIron 4802 and BigIron Chassis device, rate limiting of broadcasts, unknown-unicasts, and multicasts is enabled separately for each packet type, and each packet type has a separate rate. You can configure global rates and individual port rates. On the FastIron Edge Switch, rates are configurable on individual ports and groups of ports, but not globally.

The implementation on the FastIron 4802 and BigIron Chassis device requires IGMP snooping to be disabled in order for the limiting to work. The FastIron Edge Switch does not have this requirement. You can use broadcast/unknown-unicast and multicast limiting and IGMP at the same time.

VLAN Differences

FastIron Edge Switches support Layer 2 port-based VLANs. In release 02.0.00 and later, dual-mode ports, STP per VLAN group, GARP VLAN Registration Protocol (GVRP), and 802.1W also are supported.

Support and configuration for the VLAN features on the FastIron Edge Switches are the same as described in the *Foundry Switch and Router Installation and Basic Configuration Guide*. The following section describes the support for Layer 2 port-based VLANs.

Configuration Rules for Layer 2 Port-Based VLANs

- You can configure up to 4063 port-based VLANs on a Layer 2 Switch or 4061 port-based VLANs on a Layer 3 Switch. Each port-based VLAN can contain either tagged or untagged ports. A port cannot be a member of more than one port-based VLAN unless the port is tagged. On both device types, valid VLAN IDs are 1 – 4095. You can configure up to the maximum number of VLANs within that ID range.

NOTE: VLAN ID 4094 is reserved for use by Single STP. VLAN IDs 4091 and 4092 are reserved for use in the Layer 3 Switch and Base Layer 3 images. You can configure these VLAN IDs in the Layer 2 Switch image.

For additional configuration rules and syntax information, see the "Configuring Virtual LANs (VLANs)" chapter of the *Foundry Switch and Router Installation and Basic Configuration Guide*.

MAC Aging Differences

By default, learned MAC entries do not age out until they are unused for 300 – 600 seconds. You can change the MAC age time by entering the following command:

```
FES4802 Router(config)# mac-age-timer 60
```

Syntax: [no] mac-age-timer <secs>

You can configure 0 or a value from 60 – 600 (seconds), in 60-second intervals. If you set the MAC age time to 0, aging is disabled.

NOTE: The actual age time is from one to two times the configured value. For example, if you set the MAC age time to 60 seconds, learned MAC entries age out after remaining unused for between 60 – 120 seconds.

To display the MAC table, enter the following command:

```
FES4802 Router(config)# show mac-address
Total active entries from all ports = 3
Total static entries from all ports = 1
  MAC-Address    Port     Type   VLAN
1234.1234.1234     15   Static      1
0004.8038.2f24     14  Dynamic      1
0004.8038.2f00     13  Dynamic      1
0010.5a86.b159     10  Dynamic      1
```

In addition, the output of the **show mac-address** command is changed. The Age column has been replaced by the Type column. The Type column indicates whether the MAC entry is static or dynamic. A static entry is one you create using the **static-mac-address** command. A dynamic entry is one that is learned by the software from network traffic.

ACL Differences

As with most Foundry devices, the FastIron Edge Switches use flow-based ACL's (not rule-based ACL's). JetCore devices running software release 07.6.01 or later use hardware-based ACL's. For information about flow-based ACL's, see the "IP Access Control Lists (ACLs)" chapter of the *Foundry Enterprise Configuration and Management Guide*.

NOTE: Inbound and outbound ACLs on the FastIron Edge Switches are processed in hardware.

FastIron Edge Switches support both standard and extended IP ACLs. You can use ACLs for filtering transit traffic, for securing management access to the device, and as input to other features that use ACLs.

- Standard ACLs provide Layer 3 filtering based on source IP address.
- Extended ACLs provide Layer 3 and Layer 4 filtering based on source and destination IP addresses and IP protocol. For the TCP or UDP IP protocol, extended ACLs also filter based on TCP or UDP application port (HTTP, SNMP, and so on).

NOTE: ACLs are available in Layer 2 as well as Layer 3 and Base Layer 3 software images.

Support for up to 4000 ACL Entries

The FastIron Edge Switches support up to 4000 ACL entries.

How Flow-Based ACLs Work on the FastIron Edge Switch

ACLs on the FastIron Edge Switch are flow-based. The first time the device receives a packet for a given flow (source IP address plus destination IP address), the device compares the packet against the ACLs applied to the inbound traffic direction on the interface that received the packet.

- If an ACL denies the packet, the device drops the packet.
- If an ACL permits the packet, the device programs a forwarding entry for the packet in hardware and uses the entry to forward subsequent traffic in the same flow.

Flow entries are created in hardware for permitted TCP and UDP flows. All other flows, non TCP/UDP, and all denied flows are handled by the CPU. Furthermore, there is no difference in this regard between standard and extended ACLs.

Configuration Considerations

- You cannot use ACLs and rate limiting on the same port. If you configure both features on a port, the feature that was configured first is active. The other feature does not take effect.

Configuration Syntax

See the "IP Access Control Lists (ACLs)" chapter in the *Foundry Switch and Router Installation and Basic Configuration Guide*.

Limiting Broadcast, Multicast, and Unknown Unicast Traffic

FastIron Edge Switches 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 FastIron Edge Switch forwards each second using the following methods.

You can configure limits on individual ports or groups of ports. The valid range is 1 – 4294967295 packets per second. If you specify 0, limiting is disabled. Limiting is disabled by default.

When you enable broadcast limiting, unknown unicasts also are limited. The total number of broadcast and unknown-unicast packets sent on the port will not exceed the number you specify. To also limit multicast packets, enable multicast limiting after you enable broadcast limiting. In this case, the total number of broadcast, unknown-unicast, and multicast packets sent on the port will not exceed the number you specify.

You cannot enable multicast limiting unless broadcast/unknown-unicast limiting is already enabled.

To enable broadcast/unknown-unicast limiting on a group of ports, enter commands such as the following:

```
FES4802 Router(config)# interface ethernet 1 to 8
FES4802 Router(config-mif-e1000-1-8)# broadcast limit 10000
```

These commands configure broadcast and unknown-unicast limiting on ports 1 – 8. On each port, the total combined number of broadcasts and unknown unicasts cannot exceed 10,000.

To include multicasts in the 10,000 packets per second limit on each of the ports, enter the following command after enabling broadcast/unknown-unicast limiting:

```
FES4802 Router(config-mif-e1000-1-8)# multicast limit
```

Syntax: [no] broadcast limit <num>

Syntax: [no] multicast limit

The <num> parameter specifies the maximum number of packets per second and can be from 1 – 4294967295. The limit you specify applies to broadcast and unknown-unicast packets. If you enter the **multicast limit** command, multicast packets are included in the limit you specify. If you specify 0, limiting is disabled. Limiting is disabled by default.

Additional Information

For more information about this feature, see the following:

- "Broadcast, Unknown-Unicast, and Multicast Rate Limiting Differences" on page 40 (This section describes the differences between the FastIron Edge Switch broadcast/unknown-unicast and multicast limiting feature and the feature's implementation on the BigIron Chassis device and FastIron 4802.)

IP Load Sharing Differences

FastIron Edge Switches support load sharing among equal-cost paths to the same route destination. In a case where the IP route table has multiple equal-cost paths, the device load shares based on the destination IP address.

FastIron Edge Switches use host-based IP load sharing. This is the only type of IP load sharing supported on FastIron Edge Switches.

The host-based load sharing method uses a simple round-robin mechanism to select an equal-cost path for traffic to a destination host. When the device receives traffic for a destination host and the IP route table has multiple equal-cost paths to the host, the device checks the IP forwarding cache for a forwarding entry to the destination.

- If the IP forwarding cache contains a forwarding entry for the destination, the device uses the entry to forward the traffic.
- If the IP forwarding cache does not contain a forwarding entry for the destination, the software selects the next path in the rotation (the path after the one the software used for the previous load sharing selection). The software then creates an IP forwarding cache entry that associates the destination host IP address with the selected path (next-hop IP address).

A cache entry for host-based IP load sharing has an age time of ten minutes. If a cache entry is not used before the age time expires, the device deletes the cache entry. The age time for IP load sharing cache entries is not configurable.

sFlow Differences

FastIron Edge Switches support sFlow packet sampling for inbound and outbound traffic on sFlow-enabled ports. Sampling only of inbound traffic is supported on other Foundry devices. Byte and packet count statistics for both traffic directions are supported on FastIron Edge Switches and on other Foundry devices.

Quality of Service Differences

FastIron Edge Switches provide the following QoS features:

- Configurable queuing mechanisms
- Automatic mapping of 802.1p priorities to QoS queues
- 802.1Q support for features such as Voice over IP (VoIP)
- Configurable reassignment of traffic from one queue to another
- Prioritization through mapping of DSCP values to hardware forwarding queues (ToS-based QoS)

Queuing Mechanisms

FastIron Edge Switches support the following queueing mechanisms:

- Strict Priority (SP)
- Weighted Round Robin (WRR)

The default is WRR.

NOTE: Software releases 02.0.00 and later support the SP and WRR mechanisms. The default in 02.0.00 and later is WRR. The default in previous releases is SP.

The following table lists the queues.

Queue	Description
qosp3	Premium (highest-priority) queue
qosp2	High priority queue
qosp1	Normal priority queue
qosp0	Best-effort queue

Strict Priority (SP)

SP ensures service for high priority traffic. To do so, SP assigns the maximum valid weight to each queue, to cause the queuing mechanism to serve as many packets in one queue as possible before moving to a lower queue. This method biases the queuing mechanism to favor the higher queues over the lower queues. SP processes as many packets as possible in qosp3 before processing any packets in qosp2, then processes as many packets as possible in qosp2 before processing any packets in qosp1, and so on.

Weighted Round Robin (WRR)

WRR ensures that all four queues are serviced during each cycle. WRR rotates service among all the four queues, forwarding a specific number of bytes in one queue before moving on to the next one in a round-robin fashion. This process avoids starvation of the queues.

NOTE: The FastIron Edge Switch queue cycles are based on bytes. The device services a given number of bytes (based on the weight) in each queue cycle. BigIron queue cycles are based on packets.

The bytes-based scheme is more accurate compared to a packets-based scheme if there is a large variation in the size of the packets.

Changing the Bandwidth Allocations

To change the bandwidth percentages for the queues, enter a command such as the following. Note that this example uses the default queue names.

```
FES4802 Router(config)# qos profile qosp3 65 qosp2 20 qosp1 8 qosp0 7
Profile qosp3     : PREMIUM     bandwidth requested  65% calculated  65%
Profile qosp2     : HIGH        bandwidth requested  20% calculated  20%
Profile qosp1     : NORMAL      bandwidth requested  8% calculated   8%
Profile qosp0     : BEST-EFFORT bandwidth requested   7% calculated   7%
```

Syntax: [no] qos profile <queue> <percentage> <queue> <percentage> <queue> <percentage> <queue> <percentage>

Each <queue> parameter specifies the name of a queue. You can specify the queues in any order on the command line, but you must specify each queue.

The <percentage> parameter specifies a number for the percentage of the device's outbound bandwidth that is allocated to the queue.

NOTE: The percentages you enter must equal 100.

NOTE: The FastIron Edge Switch does not adjust the bandwidth percentages you enter. BigIron QoS does adjust the bandwidth percentages to ensure that each queue has at least its required minimum bandwidth percentage. The FastIron Edge Switch queues do not have a minimum required bandwidth percentage, so adjustment is unnecessary. For example, queue qosp3 on a BigIron device must have at least 50% of the bandwidth. There is no such requirement on the FastIron Edge Switch.

802.1p Support

The FastIron Edge Switch maps the 802.1p priority value of each packet to one of the device's four QoS queues. The following table shows the default incoming mapping.

Priority Level	Queue
6, 7	qosp3
4, 5	qosp2
2, 3	qosp1
0, 1	qosp0

By default, all traffic has priority 0. You can assign a higher priority to traffic based on the following:

- Incoming port (sometimes called the ingress port)
- Static MAC entry
- Layer 3 and Layer 4 information (IP and TCP/UDP source and destination information)
- AppleTalk socket

See Assigning QoS Priorities to Traffic.

In addition, in release 02.0.00 and later the device automatically maps a packet's DSCP value to a hardware forwarding queue. See "ToS-Based QoS" on page 46.

802.1Q Marking

If a packet enters the device on a tagged port, the device prioritizes the packet by mapping its 802.1Q value to a hardware forwarding queue. If a packet enters the device on an untagged port but is forwarded on a tagged port, the device tags the packet and adds an 802.1Q value. The 802.1Q value is based on the priority assigned to the packet as it travels through the device.

The following table shows the default outgoing mapping.

Queue	VLAN Priority Tag
qosp3	6
qosp2	4
qosp1	2
qosp0	0

Renaming the Queues

The default queue names are qosp3, qosp2, qosp1, and qosp0. You can change one or more of the names if desired.

To rename queue qosp3 (the premium queue) to "91-octane", enter the following command:

```
FES4802 Router(config)# qos name qosp3 91-octane
```

Syntax: qos name <old-name> <new-name>

The <old-name> parameter specifies the name of the queue before the change.

The <new-name> parameter specifies the new name of the queue. You can specify an alphanumeric string up to 32 characters long.

Assigning QoS Priorities to Traffic

All traffic is in the best-effort queue (qosp0) by default. You can assign traffic to a higher queue based on the following:

- Incoming port (sometimes called the ingress port)
- Static MAC entry
- Layer 3 and Layer 4 information (IP and TCP/UDP source and destination information)
- AppleTalk socket

Changing a Port's Priority

To change the QoS priority of port 1 to the premium queue (qosp3), enter the following commands:

```
FES4802 Router(config)# interface ethernet 1
FES4802 Router(config-if-1)# priority 7
```

The device will assign priority 7 to untagged switched traffic received on port 1.

Syntax: [no] priority <num>

The <num> parameter can be from 0 – 7 and specifies the IEEE 802.1 equivalent to one of the four QoS queues.

ToS-Based QoS

Software releases 03.0.xx and 03.1.xx support basic ToS-based QoS. Basic ToS-based QoS provides prioritization to a packet being forwarded out the device, by mapping the packet's DSCP value to an internal forwarding priority, which is mapped to a hardware forwarding queue.

This support is enabled by default.

NOTE: This feature applies to routed packets only. This feature does not apply to switched packets.

Software releases 03.0.xx and 03.1.xx also support marking of the DSCP value. However, it does not support other advanced ToS-based QoS features described in the "JetCore Type of Service (ToS) Based QoS" chapter of the *Foundry Enterprise Configuration and Management Guide*.

Port Monitoring Differences

FastIron Edge Switches support monitoring of both inbound and outbound traffic on individual ports. To configure port monitoring, specify the ***mirror port***, then enable monitoring on the ***monitored port***.

- The mirror port is the port to which the monitored traffic is copied. Attach your protocol analyzer to the mirror port.
- The monitored port is the port whose traffic you want to monitor.

NOTE: Disabling Layer 2 switching (entering the **route only** command) at any level of the CLI may prohibit the Foundry device from monitoring unknown unicast, multicast, and broadcast traffic on a mirror port.

Configuration Rules

Refer to the following rules when configuring port mirroring and monitoring [57806]

- The FastIron Edge Switch can have only one mirror port.
- The same port cannot be both a monitored port and the mirror port.
- Port monitoring can be active on only one port within a port region at a time.
- The monitored port and its mirror port must be in the same port-based VLAN.

- There is no restriction on which regions a mirror port and monitor port can be in. They can be in the same region or in separate regions.
- More than one monitored port can be assigned to the same mirror port.

NOTE: Release 01.0.00 required a one-to-one mapping of monitored ports to mirror ports. This is not required in release 02.0.00 and later.

- The mirror port cannot be a trunk port.
- You can monitor 10/100 ports and Gigabit ports.
- You can enable individual trunk ports to be monitored but you cannot enable an entire trunk to be monitored. Only a single port in a 10/100 trunk can be monitored. Multiple ports in a Gigabit trunk can be monitored.

Command Syntax

To configure port monitoring, enter commands such as the following:

```
FES4802 Router(config)# mirror-port ethernet 4
FES4802 Router(config)# interface ethernet 11
FES4802 Router(config-if-11)# monitor ethernet 4 both
```

Syntax: [no] mirror-port ethernet <portnum>

Syntax: [no] monitor ethernet <portnum> both

The <portnum> parameter specifies the mirror port, to which the monitored port's traffic will be copied.

NOTE: You must specify **both**, for both traffic directions. The **in** and **out** parameters are not supported on FastIron Edge Switches.

To display the port monitoring configuration, enter the **show monitor** command.

Port Statistics Differences

As with other Foundry devices, the FastIron Edge Switches support port-level Ethernet statistics. However, the meaning of the statistics listed by the **show statistics** command differs for the FastIron Edge Switches. The differences are listed in this section.

To display the statistics, enter a command such as the following:

```
FES4802 Router(config)# show statistics ethernet 3
Port  Link State     Dupl Speed Trunk Tag Priori MAC            Name
3     Up   Forward   Half 100M  None  No  level0 00e0.5200.0102

 Port 3 Counters:
        InOctets                3200          OutOctets                256
          InPkts                  50            OutPkts                  4
  InBroadcastPkts                  0    OutBroadcastPkts                 3
  InMulticastPkts                 48    OutMulticastPkts                 0
    InUnicastPkts                  2      OutUnicastPkts                 1
  InGoodFragments                  0
   InBadFragments                  0
       InDiscards                  0
         InErrors                  0
       Collisions                  0      LateCollisions                 0
              CRC                  0         MACRxErrors                 0
        GiantPkts                  0           ShortPkts                 0
           Jabber                  0
     InBitsPerSec                264       OutBitsPerSec                16
     InPktsPerSec                  0       OutPktsPerSec                 0
    InUtilization              0.00%      OutUtilization             0.00%
```

Syntax: show statistics [ethernet <portnum>]

The meaning of some of the statistics listed by the **show statistics** command is different from the meaning on other Foundry devices. Table 8 lists the statistics displayed for the FastIron Edge Switch.

Table 8: Port Statistics

This Line...	Displays...
Port Configuration	
Port	The port number.
Link	The link state.
State	The STP state.
Dupl	The mode (full-duplex or half-duplex).
Speed	The port speed (10M, 100M, or 1000M).
Trunk	The trunk group number, if the port is a member of a trunk group.
Tag	Whether the port is a tagged member of a VLAN.
Priori	The QoS forwarding priority of the port (level0, level1, level2, or level3).
MAC	The MAC address of the port.
Name	The name of the port, if you assigned a name.
Statistics	
InOctets	The total number of good octets and bad octets received.
OutOctets	The total number of good octets and bad octets sent.

Table 8: Port Statistics (Continued)

This Line...	Displays...
InPkts	The total number of packets received. The count includes rejected and local packets that are not sent to the switching core for transmission.
OutPkts	The total number of good packets sent. The count includes unicast, multicast, and broadcast packets.
InBroadcastPkts	The total number of good broadcast packets received.
OutBroadcastPkts	The total number of good broadcast packets sent.
InMulticastPkts	The total number of good multicast packets received.
OutMulticastPkts	The total number of good multicast packets sent.
InUnicastPkts	The total number of good unicast packets received.
OutUnicastPkts	The total number of good unicast packets sent.
InGoodFragments	The total number of packets received for which the following was true: • The length was less than 64 bytes • No Collision or Late Collision was detected. • The CRC was valid.
InBadFragments	The total number of packets received for which the following was true: • The length was less than 64 bytes • No Collision or Late Collision was detected. • The CRC was invalid.
inDiscards	The total number of packets that were received and then dropped due to one of the following conditions: • Lack of receive buffers • Overload on the address recognition machine
InErrors	The total number of packets received that contained one of the following errors: • CRC error – applies to regularly sized packets between 64 bytes and the maximum allowable frame size. • Oversize – applies to packets longer than the maximum allowable frame size but with a valid CRC. • Jabber – applies to packets longer than the maximum allowable frame size and with an invalid CRC. • Fragment – applies to packets shorter than 64 bytes and with an invalid CRC. • Runt – applies to packets shorter than 64 bytes but with a valid CRC, received on a full-duplex port.
Collisions	The total number of packets received in which a Collision event was detected.

Table 8: Port Statistics (Continued)

This Line...	Displays...
LateCollisions	The total number of packets received in which a Collision event was detected, but for which a receive error (Rx Error) event was not detected.
CRC	The total number of packets received for which the following was true: • The data length was between 64 bytes and the maximum allowable frame size. • No Collision or Late Collision was detected. • The CRC was invalid.
MACRxErrors	The total number of packets either received or sent in which one of the following receive errors (Rx Error) was detected: • CRC error • Oversize frame • Fragment • Jabber • Collision • Late Collision **Note**: This statistic applies only to Gigabit Ethernet ports.
GiantPkts	The total number of packets for which the following was true: • The data length was longer than the maximum allowable frame size. • No Rx Error was detected. **Note**: Packets are counted for this statistic regardless of whether the CRC is valid or invalid.
ShortPkts	The total number of packets received for which the following was true: • The data length was less than 64 bytes. • No Rx Error was detected. • No Collision or Late Collision was detected. **Note**: Packets are counted for this statistic regardless of whether the CRC is valid or invalid.
Jabber	The total number of packets received for which the following was true: • The data length was longer than the maximum allowable frame size. • No Rx Error was detected. • The CRC was invalid.
InBitsPerSec	The number of bits received per second.
OutBitsPerSec	The number of bits sent per second.
InPktsPerSec	The number of packets received per second.

Table 8: Port Statistics (Continued)

This Line...	Displays...
OutPktsPerSec	The number of packets sent per second.
InUtilization	The percentage of the port's bandwidth used by received traffic.
OutUtilization	The percentage of the port's bandwidth used by sent traffic.

Address Locking Differences

On the FastIron Edge Switch, the lock-address option is not supported on static trunk ports. Also, link-aggregation configured ports are not supported.

MAC Port Security Differences

The following summarizes the differences in the MAC port security implementation between the FastIron Edge Switch and other Foundry devices:

- The MAC port security feature is not supported for ports that are static trunk group members or ports that are configured for link aggregation.
- For MAC port security to work on the FastIron Edge Switch, a VLAN ID must be reserved for the feature. The default VLAN ID reserved for the MAC port security feature is 4090. To use a different VLAN ID for the MAC port security feature, use the **reserved-vlan-id <num>** command, which is available at the MAC port security configuration level. If MAC port security is already enabled on any port, then follow the **reserved-vlan-id <num>** command with a **write memory** command and reboot the device.
- The following new port security commands have been added on the FastIron Edge Switch:
  - The command **show port security ethernet <port_num> restricted-macs** displays a list of restricted MAC addresses on the port.
  - The command **clear port security restricted-macs [all | ethernet <port_num>]** clears all restricted MAC addresses from the port.
  - The command **clear port security statistics [all | ethernet <port_num>]** clears violation statistics for the port.
- You can specify a number of minutes that the device drops packets from a violating address. To do this, use the **violation restrict <age>** command. The <age> can be from 0 – 1440 minutes. The default is 5 minutes. Specifying 0 drops packets from the violating address permanently.

  Aging for restricted MAC addresses is done in software. There can be a worst case inaccuracy of one minute from the specified time.

  The restricted MAC addresses are denied in hardware in the FES.
- When the **restrict** option is used, the maximum number of MAC addresses that can be restricted is 128. If the number of violating MAC addresses exceeds this number, the port is shut down. An SNMP trap and the following Syslog message are generated: "Port Security violation restrict limit 128 exceeded on interface ethernet <port_id>". This is followed by a shutdown Syslog message and trap.
- The SNMP trap generated for restricted MAC addresses has been enhanced to indicate the VLAN ID associated with the MAC address, as well as the port number and MAC address.
- When specifying a secure MAC address on a tagged port, you must specify the VLAN ID as well. To do this, use the command **secure-mac-address <mac-address> <vlan-id>**.

NOTE: If MAC port security is enabled on a port, and you dynamically change the VLAN membership of the port, make sure that you also change the VLAN ID specified in the **secure-mac-address** configuration statement for the port.

In generation of the configuration, the **vlan-id** is generated for both tagged and untagged ports. When you display the configuration, you will see an entry for the secure MAC addresses `secure-mac-address <address> <vlan>`. For example, you may see the following line:

```
secure-mac-address 0000.1111.2222 10
```

This line means that MAC address 0000.1111.2222 on VLAN 10 is a secure MAC address.

DHCP Assist Differences

The following summarizes the difference in the DHCP Assist feature between the FastIron Edge Switch and other Foundry devices:

- When DHCP Assist is enabled on any port on the FastIron Edge Switch, then Layer 2 broadcast packets are forwarded by the CPU. Unknown unicast and multicast packets are still forwarded in hardware. When DHCP Assist is not enabled on the FastIron Edge Switch, Layer 2 broadcast packets are forwarded in hardware.
- In software release 2.0, using DHCP Assist on the FastIron Edge Switch required a VLAN ID for trap broadcasts. You could configure the VLAN ID using the **broadcast-trap-vlan-id** command. Starting with software release 3.0, you no longer need to reserve a VLAN ID for this feature. Consequently, the **broadcast-trap-vlan-id** command has been removed from the CLI.

SNMP MIBs

The following Simple Network Management Protocol (SNMP) Management Information Base (MIB) objects are not supported in the FastIron Edge Switch:

- The standard MIBs in RFCs 1657 (BGP4) and 1850 (OSPF)
- The Foundry MIBs for BGP

All other MIBs documented in the *Foundry Management Information Base Reference* are supported.

Base Layer 3

For information on how to configure static IP and RIP in the Base Layer 3 software image see the "Configuring Base Layer 3" chapter of the *Foundry FastIron Edge Switch Installation and Basic Configuration Guide*.

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*.

Where To Get More Information

These release notes provide basic setup information. For more advanced configuration information, see the May 2003 or later versions of the documents listed in Table 9.

Table 9: Feature Documentation

Title	Contents
Foundry FastIron Edge Switch Installation and Basic Configuration Guide	• Product Overview • Installation • Configuring Basic Features • Updating Software • Hardware Specifications • Supported RFCs
Foundry Switch and Router Installation and Basic Configuration Guide	• Link Aggregation • Spanning Tree Protocol • Virtual LANs • Layer 2 Multicast
Foundry Security Guide	• Security (passwords, user accounts, AAA, RADIUS, and TACACS/TACACS+) • Secure Shell (SSH) • Denial of Service Protection
Foundry Enterprise Configuration and Management Guide	• ACLs • IP • RIP • IP Multicast • OSPF • VRRP and VRRPE • IPX • AppleTalk
Foundry Switch and Router Command Line Interface Reference	Syntax information for all CLI commands. See the "Command List" chapter for a complete list of the CLI commands and page references to syntax information.
Foundry Management Information Base Reference	Simple Network Management Protocol (SNMP) Management Information Base (MIB) objects. NOTE: Standard MIB RFCs 1657 (BGP4) and 1850 (OSPF), and the MIBs for BGP are not supported in FES Software Release 03.1.00.
Foundry Diagnostic Guide	Diagnostic commands available on Foundry devices

Software Fixes

This section lists the software fixes in software releases 03.1.02, 03.1.01, and 03.1.00.

The software fixes are sorted by category, then by priority.

The **P** column indicates the priority of the software fix, as follows:

- 0 = Critical
- 1 = Major
- 2 = Medium
- 3 = Minor

Software Fixes in 03.1.02

This section lists the software fixes in release 03.1.02.

Software release 03.1.02 contains all the software fixes from patch release 03.1.01a.

Table 10: Software Fixes in Release 03.1.02

Category	P	Description	Bug ID #
CLI	2	Module: FES 4802-POE Symptom: The **ip default-network** command is missing from the CLI Resolution: Fixed in 03.1.02	25130
Diagnostics	1	Module: N/A Symptom: A diagnostic test fails while attempting to verify an EEPROM socket that does not have an EEPROM installed. Resolution: In software release 03.1.02, the diagnostic test first verifies that there is an EEPROM installed in the socket. If an EEPROM does not exist, the test is skipped.	20474
IPX Stack	2	Module: FES Symptom: When IPX network VLANs are configured on the FastIron Edge Switch, it does not transmit IPX RIP packets. Resolution: Fixed in 03.1.02	23074
Other	1	Module: N/A Symptom: After a cold start, the FES device reloads the software and does not provide complete data regarding the reload. Resolution: Fixed in 03.1.02	23538
Port Security	0	Module: FES Symptom: Disabling the Spanning Tree Protocol (STP) at the Interface Port Security level of the CLI causes the device to globally disable STP. Resolution: Fixed in 03.1.02	23400
SNMP Management	0	Module: FES Symptom: Reading the snFdpCachedAddrValue MIB variable causes some memory corruption, which may lead to a software reload after some delay. Resolution: Fixed in 03.1.02	25339

Table 10: Software Fixes in Release 03.1.02

Category	P	Description	Bug ID #
SNMP Management	1	Module: N/A Symptom: The FES device reports a false power supply failure. Resolution: In software release 03.1.02, the software performs more thorough checks of the power supply, effectively eliminating any false power supply failure reports.	25115
SNMP Management	2	Module: FES Symptom: The following MIB objects are not available: • snFdpGlobalRun • snFdpGlobalMessageInterval • snFdpGlobalHoldTime • snFdpGlobalCdpRun • snFdpInterfaceTable • snFdpCacheTable • snFdpCachedAddressTable Resolution: Fixed in 03.1.02	23342
System	0	Module: FES Symptom: A software reload may occur if a Layer 2 spanning tree loop occurs and BPDU packets are spinning in the loop. Resolution: Fixed in 03.1.02	23785
VRRP	2	Module: FES Symptom: When configuring VRRP and VRRP-E, the CLI prompt always displays VRID 0. For example, if you are configuring VRID 1, the CLI prompt displays: `FES2402 Router(config-if-e100-1-vrid-0)#` Also, the device truncates log messages for VRRP and VRRP-E interface state changes. As a result, the "master" or "standby" state change keyword does not appear in the VRRP log report. Resolution: Fixed in 03.1.02	22979

Software Fixes in 03.1.01

The following table lists the software fixes in release 03.1.01.

Software release 03.1.01 contains all the software fixes from patch releases 03.1.00a – 03.1.00g.

Table 11: Software Fixes in Release 03.1.01

Category	P	Description	Bug ID #
AAA	0	Module: FES Symptom: If you cut and paste a set of ACL commands where the commands are indented, the Foundry device may remove the first character in the **permit** or **deny** clause, thereby making the command invalid. The Foundry device does not apply the rest of the ACL commands, as the device interprets this as an error in the configuration. Resolution: Fixed in 03.1.01	22875
AAA	2	Module: FES Symptom: The Foundry device rejects pre-authorization of "AUTHOR_STATUS_PASS_REPL 0x02" (status = 2). Resolution: Fixed in 03.1.01	16275
Appletalk	1	Symptom: Appletalk clients are not able to view Appletalk zones for other clients connected to the same FastIron Edge Switch. Resolution: Fixed in 03.1.01	21520
FDP/CDP	1	Module: FES Symptom: When FDP is enabled on the Foundry device, and a directly connected device has a VLAN ID greater than 4000, a software reload occurs when the **show fdp neighbor detail** command is executed from the Foundry device. Resolution: Fixed in 03.1.01	22017
Gigabit negotiation mode	1	Module: FES Symptom: The Foundry device does not apply the default gigabit negotiation setting (**gig-default** command) if it is configured at the global CONFIG level of the CLI. However, the device applies the default gigabit negotiation setting if it is configured at the Interface level of the CLI. Resolution: Fixed in 03.1.01	21704
OSPF	1	Module: FES Symptom: The output of the **show ip ospf int** command repeats the same information several times. Resolution: Fixed in 03.1.01	20273
OSPF	2	Module: FES Symptom: The default route appears in the routing table even though **distribute- list** is used to filter it out. Resolution: Fixed in 03.1.01	18566

Table 11: Software Fixes in Release 03.1.01

Category	P	Description	Bug ID #
Other	1	Module: FES Symptom: The Foundry device intermittently performs a software reload. Resolution: Fixed in 03.1.01	21700
SNMP Management	0	Module: FES Symptom: When MPLS is configured on the Foundry device, the Management module reloads during an SNMP walk of snRtIpPortIfConfigTable and snRtIpRipPortIfConfigTable. Resolution: Fixed in 03.1.01	21173
SNMP Management	1	Module: FES Symptom: If you configure more than one community string, the software applies only the last configured community string. For example, if you configure both a read-only and read-write community string, the software does not allow SNMP access to the device using the read-only community string. It does allow access to the device using the read-write community string, since this is the last configured community string. Resolution: Fixed in 03.1.01	22330
SNMP Management	1	Module: FES Symptom: The SNMP Get requests for the following objects return incorrect values: • snChasGen.snChasActualTemperature • snChasGen.snChasWarningTemperature • snChasGen.snChasShutdownTemperature Resolution: Release 03.1.01 reverses the software fix from bug ID #10436 (from a previous release), which incorrectly changed each temperature MIB object unit from 0.5 to 1.0 degrees Celsius.	18155
SNMP Management	2	Module: FES Symptom: If an STP group is configured on the device, the device does not respond to Get Next or Get Bulk requests once the requests come to any object in the snPortStpTable. Resolution: Fixed in 03.1.01	11414
SSH	0	Module: FES Symptom: The Foundry device prompts for a password but fails to authenticate it, even though Secure Copy (SCP) is configured on the device. Resolution: Fixed in 03.1.01	13712

Table 11: Software Fixes in Release 03.1.01

Category	P	Description	Bug ID #
SSH	3	Module: FES Symptom: In an SSH session, if you insert a word with the help of the left arrow key, the Foundry device misinterprets the left arrow key as a backspace. For example, if you enter **show server cache1**, then use the arrow key to insert **http** so that the command becomes **show server http cache1**, the Foundry device processes the command as **show server http**. Note that this does not occur with Telnet sessions. Resolution: Fixed in 03.1.01	22242
Web Management	2	Module: FES Symptom: The Web management interface always displays a value of zero for the temperature. Resolution: Fixed in 03.1.01	22258

Software Fixes in 03.1.00

The following table lists the software fixes from patch releases 03.0.01d and 03.0.01e.

Table 12: Software Fixes in Release 03.1.00

Category	P	Description	Bug ID #
Other	0	Module: FES Symptom: The POE device fails to provide inline power to POE ports when a power supply comes back up. Resolution: Fixed in 03.1.00	21560
System	1	Module: FES Symptom: A POE device logs power supply failures in the Syslog, even though the power supplies are operating normally. The false reports cause unwarranted power flaps on POE ports. Resolution: Fixed in 03.1.00	21522
System	1	Module: FES Symptom: A POE port might not turn OFF the inline power when a powered device is physically disconnected from the port. This is evident by the power LED which will remain ON even after disconnecting the cable from the port. Under normal operations, a POE port will disable the inline power and turn OFF the LED when a powered device is physically disconnected from the port. Resolution: Fixed in 03.1.00	21118

Known Issues in 03.1.02

The following table lists the known issues in release 03.1.02.

The **P** column indicates the priority of the known issue, as follows:

- 0 = Critical
- 1 = Major
- 2 = Medium
- 3 = Minor

The known issues are sorted by category, then by priority.

Table 13: Known Issues in Release 03.1.02

Category	P	Description	Bug ID #
Appletalk	1	Module: FES Symptom: Appletalk clients do not correctly detect network addresses if **clear cache** and **clear route** commands are issued several times. Workaround: Clear the Appletalk configuration then reconfigure.	23593
Web Management	2	Module: FES Symptom: In Web Management, the Configuration –> VLAN –> Static Station window displays **router-type** and **host-type** as configurable options, although neither are not supported on the FastIron Edge Switches. Workaround: Do not select the **router-type** or **host-type** radio buttons when configuring static MAC entries.	21377

