Brocade Fabric OS v5.0.5b

Release Notes v1.0

November 03, 2006

Document History

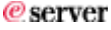| Document Title	Summary of Changes	Publication Date
Brocade Fabric OS v5.0.5b Release Notes v1.0	First release	November 03, 2006

CONTENTS

Quick Look

If you are already using the most recent version of the Fabric OS v5.0.5a Release Notes, here are the changes between that version and this version.

- Changes (italics and boldface) to the wording in the second paragraph on page 22, under the heading "N-Port ID Virtualization."

- Listing of new fixes for Fabric OS v5.0.5b at the end of this document.

Overview

Fabric OS v5.0.5b contains fixes to defects found since the release of Fabric OS v5.0.5. Aside from these fixes, this release includes the same feature set as Fabric OS v5.0.5.

Fabric OS v5.0.5 includes support for three new platforms: SilkWorm 4016, SilkWorm 4018 and SilkWorm 4024.

Fabric OS has been updated with two new features, tsTimeZone command and Dynamic Ports On Demand (Dynamic POD), both described below. The Dynamic POD functionality is currently supported on the SilkWorm 4020 and SilkWorm 4024 platforms only.

Here is some general information about previous releases:

- Brocade Fabric OS v5.0.3 and Fabric OS v5.0.4 were maintenance releases.

- Brocade Fabric OS v5.0.2 supported one additional platform: SilkWorm 4020.

- Brocade Fabric OS v5.0.1 supported two additional platforms: SilkWorm 48000 256-port director and SilkWorm 200E standalone switch.

tsTimeZone command

The *Energy Policy Act of 2005* changed the time change dates for Daylight Saving Time (DST) in the US. Beginning in 2007, DST will begin on the second Sunday of March and end the first Sunday of November. A new time was added to the zone format to allow you to set up time zones by region; DST is then automatically adjusted.

The **tsTimeZone** command is used to display or set the time zone as follows:

- `tstimezone [hourOffset [, minuteOffset]]`
  hourOffset (optional) – Specify the number of hours relative to GMT; must be an integer. Valid values are -12 through 12.
  minuteOffset (optional) – Specify the number of minutes relative to the hourOffset value; must be an integer. Valid values are -30, 0, or 30. The hourOffset can be specified without a minuteOffset, which then defaults to zero.

- `tstimezone –interactive`

- `tstimezone <timezonename>`
  zonename – The name of the zoneinfo time zone; will vary as additional time zones are added.

For example, to display the current time zone setup and then change them to GMT -3:30:

```
switch:admin> tsTimeZone
Time Zone Hour Offset: 0
Time Zone Minute Offset: 0

switch:admin> tsTimeZone -3, -30
Updating Time Zone configuration...done.
System Time Zone change will take effect at next reboot.

switch:admin> tsTimeZone
Time Zone Hour Offset: -3
Time Zone Minute Offset: -30
```

Dynamic Ports On Demand

SilkWorm 4018, SilkWorm 4020 and SilkWorm 4024 only: Dynamic Ports On Demand (Dynamic POD) automatically enables ports as they are physically cabled or in the case of a server blade when the server is powered on. Existing Ports On Demand (POD) functionality is static, that is, a preset group of ports are enabled with each POD license. In contrast, the new Dynamic POD functionality does not require a predefined assignment of ports. With Dynamic POD, the determining factor is the total number of ports in use and the number of purchased dynamic ports.

Dynamic POD:

- Automatically detects HBA connected server ports or cabled ports
- Assigns a POD license to each of these ports

This dynamic allocation of POD licenses is controlled with the **licensePort** command (see below). When the Dynamic POD mechanism detects a server blade in an online state, it automatically assigns the blade a POD license.

The Dynamic POD mechanism must detect which ports have active links and make license assignments to these ports. If the initial assignment of Dynamic POD ports does not exhaust the purchased POD number, then new connections can be made at a later time not to exceed the total number of purchased dynamic ports. A port is considered connected if it can be brought to the online state.

NOTE: Dynamic POD does not consider disabled ports as candidates for license assignments. You can persistently disable an otherwise viable port to cause it not to come online and preserve a license assignment to be used by some other port.

When Dynamic POD is used in server blades that are embedded in a blade server enclosure, such as the SilkWorm 4024, a port can be detected and assigned a POD license only if the server blade is installed with an HBA present. A server blade that does not have a functioning HBA will not be treated as an active link for the purpose of initial POD port assignment.

The **licensePort** command is used to manage Dynamic POD assignments with the following options:

- `licensePort  -reserve portnum`
  Sets a port as reserved, that is, it will receive a POD license even if the port is not capable of being online.

- `licensePort  -release portnum`
  Deassigns a license that has already been assigned, but does not block future assignments. To block a specific port from future assignments, run **portCfgPersistentDisable** on the port.

- `licensePort  -show`
  Displays an overview of the POD license status and port assignments.

SilkWorm 48000

The SilkWorm 48000 director is based on the proven technology and innovation that has earned Brocade clear market share leadership in the SAN market. Brocade is the only vendor that offers the full spectrum of products from entry to enterprise, for multi-protocol routing and SAN solutions for bladed servers. Brocade leads in every category in the industry with the first 4 Gbit/sec products. Brocade's Advanced Fabric Services, which are delivered across the product family, are also extended with the introduction of the SilkWorm 48000.

- **High-end performance:** The SilkWorm 48000 is the industry's first Fibre Channel director that supports 4 Gbit/sec port speeds. It delivers exceptional performance and scalability with up to 256 ports in a single domain. The high-performance architecture provides auto speed negotiation to support legacy 1 and 2 Gbit/sec server and storage devices as well as new and forthcoming 4 Gbit/sec devices. The new 4 Gbit/sec technology also provides the ability to aggregate up to eight 4 Gbit/sec ports to create an Inter-Switch Link (ISL) trunk at up to an unprecedented 32 Gbit/sec of

bandwidth between directors. High-end performance also applies to the extension of Fibre Channel over distance, supporting distances up to 500 kilometers. Trunking in the SilkWorm 48000 can also be extended over distance, enabling new levels of performance between data centers.

- **Investment protection:** Fully compatible with existing Brocade storage network offerings, the highly flexible blade format of the SilkWorm 48000 provides "pay-as-you-grow" scalability and support for multiple protocols and transports. Routing with Logical Private SAN (LSAN) enables secure selective sharing of resources between isolated SANs. FICON and CUP support enables an intermix of mainframe and open systems in a consolidated SAN.

- **Mission-critical availability, scalability, and flexibility:** The SilkWorm 48000 is designed for continuous operation. It supports "five-nines" availability with built-in redundancy; FRUs capable of hot-swap install/uninstall, and hardware and software upgrades concurrent with operation. The SilkWorm 48000 provides 256 ports per system and 768 ports per rack to help maximize valuable data center real estate. The leading network scalability of the Brocade SilkWorm family of products is extended with the SilkWorm 48000, which provides the largest building block for creating the largest storage networks.

- **Lower Total Cost of Ownership (TCO):** The SilkWorm 48000 lowers the overall costs of deploying and operating SAN infrastructures. With twice the port density of previous directors, the SilkWorm 48000 delivers more efficient use of expensive data center floor space. Lower power consumption per port represents significant cost savings in electricity and cooling expenses for the data center, as much as $10,000 per year per system. More ports per director also means fewer devices to manage in large fabrics, improving administrative efficiencies for IT departments.

Fabric OS v5.0.1 includes all basic switch and fabric support software, as well as optionally licensed software enabled via license keys. It comprises two major software components: firmware, which initializes and manages the switch hardware, and diagnostics.

Optionally licensed products include:

- Brocade Extended Fabrics—Provides up to 500 km of switched fabric connectivity at full bandwidth over long distances.

- Brocade ISL Trunking Over Extended Fabrics—ISL Trunking has been enhanced to enable trunking over long-distance links of up to 250 km via a new command.

- Brocade Web Tools—Enables administration, configuration, and maintenance of fabric switches and SANs.

- Brocade Fabric Manager—Enables administration, configuration, and maintenance of fabric switches and SANs with host-based software.

- Brocade Advanced Performance Monitoring—Enables performance monitoring of networked storage resources.

- Brocade Fabric Watch—Monitors mission-critical switch operations.

  Included in every switch:

- Brocade Advanced Zoning—Segments a fabric into virtual private SANs.

NOTE: Brocade software release policy is to carry forward all fixes in patches to subsequent maintenance and feature releases of Fabric OS.

SilkWorm 200E

As the latest addition to the SilkWorm family of fabric switches and directors, the SilkWorm 200E provides small- to medium-size businesses deploying their first SAN or expanding their current SAN with low-cost access to easy-to-manage SAN technology. The SilkWorm 200E provides the lowest-cost 8- to

16-port SAN switch available for those who want the benefits of SAN solutions with the option to scale to larger fabrics on a "pay-as-you-grow" basis.

Brocade further simplifies the process of implementing SAN solutions with the SilkWorm 200E. The simplicity and ease-of-use features of the SilkWorm 200E help increase administrator productivity and lower the cost of management, which can benefit organizations with limited IT expertise. In addition, the SilkWorm 200E leverages industry-leading 4 Gbit/sec Fibre Channel technology to provide extremely high performance.

Delivering 8, 12, or 16 ports in a 1U form factor, the SilkWorm 200E enables substantial cost savings— from capital and operating expenses to overall management. It extends the Brocade modular building block approach to the development of storage networks. This approach has been widely adopted by storage networking vendors and is the de facto standard in the storage networking industry. The SilkWorm 200E stands up to any mission-critical test and offers significant business and performance advantages to small- to medium-size businesses as they develop and grow.

Supported Switches and Directors

Fabric OS v.5.04 supports the SilkWorm 200E, 4020, 48000, 200E, 3014, 3016, 4012, 3250, 3850, 3900, 4100, 12000, 24000, and 48000.

Firmware Upgrades

The recommended procedure for upgrading Fabric OS firmware levels is to limit the release levels to two or fewer releases. For example, upgrading a switch from v4.1.0 to v5.0.1 requires a two-step process: first upgrading to v4.4.0 and then upgrading to v5.0.1.

Technical Support

Contact your switch supplier for hardware, firmware, and software support, including product repairs and part ordering. To expedite your call, have the following information immediately available:

1. **General Information**

   - Technical Support contract number, if applicable
   - Switch model
   - Switch operating system version
   - Error numbers and messages received
   - **supportSave** command output
   - Detailed description of the problem and specific questions
   - Description of any troubleshooting steps already performed and results

2. **Switch Serial Number**

   The switch serial number and corresponding bar code are provided on the serial number label, as shown here.

   ```
   *FT00X0054E9
      FT00X0054E9
   ```

   The serial number label is located as follows:
   - SilkWorm 4024 – Bottom of the switch module
   - SilkWorm 4016 – Top of the switch module
   - SilkWorm 3016 and 4012—Side of switch module
   - SilkWorm 200E—Nonport side of the chassis
   - SilkWorm 3250, 3850, and 3900—Bottom of the chassis

- SilkWorm 4100—On the switch ID pull-out tab located on the port side and on the inside of the chassis, near power supply 1 (on the right when looking at the nonport side)
- SilkWorm 12000, 24000, and 48000 directors—Inside front of the chassis, on the wall to the left of the ports
- SilkWorm Multiprotocol Router Model AP7420—On the bottom of the chassis and on the back of the chassis

3. **World Wide Name (WWN)**

- SilkWorm 200E, 3016, 3250, 3850, 3900, 4024, 4016, 4012, and 4100 switches and SilkWorm 12000, 24000, and 48000 directors—Provide the license ID. Use the **licenseIDShow** command to display the license ID.

- SilkWorm Multiprotocol Router Model AP7420—Provide the switch WWN. Use the **switchShow** command to display the switch WWN.

- All other SilkWorm switches—Provide the switch WWN. Use the **wwn** command to display the switch WWN.

Standards Compliance

This software conforms to the following Fibre Channel Standards in a manner consistent with accepted engineering practices and procedures. In certain cases, Brocade might add proprietary supplemental functions to those specified in the standards. Brocade verifies conformance with Fibre Channels Standards by subjecting its switches to SANmark Conformance Tests developed by the Fibre Channel Industry Association. Brocade switches have earned the SANmark logo, indicating such conformance. SANmark is a limited testing program and does not test all standards or all aspects of standards. For a list of standards conformance, visit this Brocade Web site:

http://www.brocade.com/sanstandards

Fabric OS Requirements

The following table lists the earliest versions of Brocade software supported in this release, that is, the *earliest* software versions that interoperate. Brocade recommends using the *latest* software versions to get the greatest benefit from the SAN.

For a list of the effective end-of-life dates for all versions of Fabric OS, visit the following Brocade Web site:

http://www.brocade.com/support/end_of_life.jsp

	Earliest Compatible Version	With Secure Fabric OS Enabled	Recommended Software Versions
SilkWorm 4024	v5.0.5	v5.0.5 or later	Latest 5.0.x
SilkWorm 4016	v5.0.4	v5.0.4 or later	Latest 5.0.x
SilkWorm 4020	v5.0.2 [1]	v5.0.2 or later [1]	Latest v5.0.x
SilkWorm 4012	v5.0.0	v5.0.x	Latest v5.0.x
SilkWorm 2000 series	v2.6.1	v2.6.1 or later	v2.6.2d
SilkWorm 3200 and 3800	v3.1.0	v3.1.2 or later	v3.2.0a
SilkWorm 12000	v4.1.0	v4.2.0 or later	Latest v5.0.x
SilkWorm 3014	v4.4.1	v4.4.1 or later	Latest v5.0.x
SilkWorm 3016	v4.2.1	V4.2.1 or later	Latest v5.0.x
SilkWorm 3250, 3850, 3900, and 24000	v4.1.0	v4.2.0 or later	v5.1.0 or later
SilkWorm 200E and 48000	v5.0.1 [2]	v5.0.1 or later [2]	v5.1.0 or later [2]
SilkWorm 4100	v4.4.0c	v4.4.0c or later	v5.1.0 or later
Fabric Manager	See the Fabric Manager Release Notes		

1. Fabric OS 5.0.2 is supported only on the SilkWorm 4020. Attempts to load this software version on Brocade switches other than the SilkWorm 4020 will result in File-Not-Found errors.

2. If you are using the FR4-18i switch module in a SilkWorm 48000 director, you must use Fabric OS v5.1.0 or later.

Important Notes

This section lists information that you should consider when running this Fabric OS release.

As of May 15, 2005, Brocade no longer includes a PKI Certificate as part of the installed Secure Fabric OS. If you wish to activate Secure Fabric OS on a supported director or switch, you must contact Brocade to obtain a PKI certificate.

Refer to the *Secure Fabric OS Administrator's Guide*, Chapter 2, "Adding Secure Fabric OS to the Fabric," for a description of how to obtain certificates from the Brocade Certificate Authority.

Upgrading and Installing FC4-16 and FC4-32 Blades

If you are planning to install FC4-16 or FC4-32 blades, you must upgrade firmware to Fabric OS v5.0.1 on both CPs *before* you install the blades.

General

The major new features incorporated in Fabric OS v5.0.1 are summarized below.

SilkWorm 48000 Platform Support

The SilkWorm 48000 places Condor-ASIC-based port and CP blades into the same core-edge infrastructure that was provided by the SilkWorm 24000 product. When a system is fully populated, it supports 256 ports in a single domain.

System/Blade Identification

Two CP blades that have different processors and slightly different hardware characteristics can co-exist in an active/standby relationship in the same SilkWorm 48000 chassis.

The platform identifiers for the two blades differ – largely to support the proper selection of platform-specific RPMs for the two different blades. When **switchShow** is issued from an active SilkWorm 24000 CP blade, and the same command is issued when the SilkWorm 48000 CP blade is the active CP blade, **switchShow** shows two different switchType values. The switch type, however, tracks with the active CP blade – just as it does for the SilkWorm 24000 and 12000.

The condition of heterogeneous CP blades in a single chassis is designed to be transient. It should exist only until you have upgraded the system to homogeneous blades. However, Fabric OS does not distinguish between the cases in which the mixed configuration exists for a short period of time or a longer period of time.

Chassis Configuration Options

With the Fabric OS v4.4.0 release, a new command, **aptPolicy**, allowed you to configure which egress port is selected for a frame, based on a particular policy:

- Port-based path selection (paths are chosen based on ingress port and destination only). This also includes user-configured paths. (Required for FICON see below)

- Device-based path selection (paths are chosen based on SID and DID). Note: Device based routing is currently not used and should not be activated at any time.

- Exchange-based path selection (paths are chosen based on SID, DID, and OXID). This is the default routing policy for Open Systems environments.

- For the SilkWorm 48000, the **aptPolicy** command is not available unless the chassis has been configured to run using option 5 described in the table below.

With the introduction of Fabric OS v5.0.1b and FICON support for the SilkWorm 4100 and 48000, the **aptPolicy** routing policy for FICON must be configured for port-based path selection on any director or switch with FICON devices attached. Other switches that reside in the fabric with Open Systems devices exclusively can remain configured as exchange-based routing. Any Brocade-supported FICON platforms can be cascaded without issue with FICON devices attached for backwards compatibility (that is, you can connect any of the SilkWorm 3900, 4100, 12000, 24000, and 48000 together).

For all other chassis configurations modes (1-4), the default routing policy is port-based path selection (paths are chosen based on ingress port and destination only). This cannot be changed. This also includes user-configured paths.

NOTE: Chassis configuration mode 1 is supported for FICON. Chassis configuration modes 2-4 are not supported for FICON configurations.

Table 1 SilkWorm Chassis Option Descriptions

Option	Number of Domains: Domains	Routing Module	Supported CPs	Supported Port Blades	Implications/Notes
1	1: 128	CER	CP2 or CP4	FC2-16, FC4-16	CP4 will be faulted if inserted into a D2 chassis
2	2: 64/64	CER/CER	CP2 only	FC2-16 only	
3	2: 64/64	CER/XYR	CP2 only	Left side: FC2-16 Right side: 12K	Same support as Fabric OS v4.4
4	2: 64/64	XYR/CER	CP2 only	Left side: 12K Right side: FC2-16	Same support as Fabric OS v4.4
5	1: 256	RTE	CP4 only	FC4-16, FC4-32	CP4 will be faulted if inserted into a D2 chassis

Key

CER = Core Edge Routing. Port-based routing scheme, same as routing option supported in v4.2 and v4.4

XYR = X-Y Linear Routing. Routing scheme used on SW12000 switches

RTE = Advanced Routing. Exchange-based (default) or device-based routing scheme

CP2 = SilkWorm 24000 CP blade

CP4 = SilkWorm 48000 CP blade

FC2-16 = 2G, 16-port blade

FC4-16 = 4G, 16-port blade

FC4-32 = 4G, 32-port blade

12K = SilkWorm 12000-port blade (2G, 16-port)

SilkWorm 200E Platform Support

The SilkWorm 200E is a 16-port pizza-box Fibre Channel switch using the Brocade Goldeneye ASIC. The Goldeneye ASIC implements a large subset of Brocade Condor ASIC functionality. Fabric OS v5.0.1 supports this platform including the SilkWorm 200E ports-on-demand (POD) features, which delivers 8, 12, or 16 ports in a 1U form factor.

Reliability

This release of Fabric OS features RSCN suppression: the ability to control RSCNs originating from hosts on a port-by-port basis.

Enhanced RAS Log Messages

New with Fabric OS v5.0.1 are Zoning Audit messages. These messages record information about the type of zoning change made (including such tasks as **cfgenable** and **cfgdisable**) and the role level and user name making the changes. The messages are recorded in the RASlog whether change was made through the CLI or Web Tools. Note that occasional redundant entries are possible due to an extra HTTP entry when zoning changes are performed through the CLI.

Scalability

Scaling the SAN is addressed under two topics:

- Single switch scalability, that is, the ability to handle up to 256 switch ports with some number of directly attached Nx_Port types or the ability to effectively operate in a multiple-switch environment
- Fabric scalability, that is, the maximum number of ports and domains available fabric-wide

Single Switch Scalability

The parameters specified here describe the number and types of ports that can be directly attached to a single-switch SW48000 domain:

- Maximum 256 user ports active on a single domain
- Maximum 252 initiators attached to a SilkWorm 48000 (with the balance of the ports connected to target ports, or ISLs)
- Maximum 14 switch ports connected to loops (for example, JBODs) of up to 24 devices (with the balance of the switch ports connected to N-Ports)

The rate of N-port connections is metered to ensure that devices sensitive to timeouts (such as in FICON environments) are not adversely affected. This mechanism delays N-port connections until all members of the fabric become reachable (allows routing and fabric formation to proceed without competition for CPU from N-port related service loads). When all these conditions are met, all ports that have been disabled for this reason are re-enabled.

Fabric Scalability

Fabric OS v5.0.1 supports the same fabric scalability as Fabric OS v4.4.0: 2,650 ports with 50 domains.

FICON

FICON is now supported for the SilkWorm 48000 and 4100 platforms.

The FICON protocol is now supported on the following SilkWorm models and Fabric OS releases:

- SilkWorm 48000 and Fabric OS v5.0.1b
  A single-domain configuration is supported with a mix of 16-port and 32-port SilkWorm 48000 port blades in a SilkWorm 48000 director. Dual-domain configurations are not supported on the SilkWorm 48000 director. Mixed port blade configurations with SilkWorm 24000 and SilkWorm 48000 port blades in the same director are not supported in a FICON environment.

- SilkWorm 4100 and Fabric OS v5.0.1b

Both the SilkWorm 48000 and 4100 require the port-based routing policy either in a single-switch configuration or a cascaded-switch configuration only on those switches in the fabric that contain FICON devices (option 1 of the **aptPolicy** command). Other switches in the fabric may exist with the default exchange-based routing option
 (option 3 of the **aptPolicy** command) if only Open Systems devices are attached to those switches.

CUP is supported on SilkWorm 4100 and 48000 running Fabric OS 5.0.1b.

Fabric OS provides standard support for FICON single-switch operation. Multiple-switch cascaded FICON operation (double-byte addressing) requires a Brocade Secure Fabric OS license.

Control Unit Port (CUP) operation requires a Brocade FICON CUP license.

Although there are no specific zoning rules related to FICON environments, it is recommended that you follow standard FCP zoning practices. For management purposes, put FCP devices in one zone and FICON devices in another zone when operating in a mixed environment.

Any Brocade-supported FICON platforms can be cascaded without issue with FICON devices attached for backwards compatibility (that is, you can connect any of the SilkWorm 3900, 4100, 12000, 24000, and 48000 together).

For further information on FICON, refer to the latest Brocade technical documentation for the Fabric OS v 5.0.1 and Fabric Manager 5.0.0.

Problem Determination

Fabric OS v5.0.1 features the **fcPing** command, which provides the ability to check Fibre Channel connectivity between any two nodes in a fabric.

Security-Related Enhancement

A new role-based access control role, switch administrator, allows an administrator to control a switch but not modify any fabric-wide configuration, that is, security, zoning, or user configuration (see the **userConfig** command).

Merging Zones

Before linking two switches together, it is important to know the zone database limit of adjacent switches. For details, refer to the section "Merging Zones," in the *Fabric OS Administrator's Guide* documentation update on of this document.

Web Tools

For instructions on installing Mozilla 1.6 on Solaris 2.8 and Solaris 2.9, refer to the following Web site:

http://www.mozilla.org/releases/mozilla1.6/installation.html

Issue: The Mozilla browser does not support the Switch Admin module properly in Fabric OS v2.6.x. In Fabric OS v2.6.2, a warning message is displayed. For other v2.6.x versions, no warning message is displayed.

Workaround: Use Netscape 4.7.7 or later.

The added supported browsers, operating systems, and Java Plug-ins introduce the following limitations when using mixed OS versions in Web Tools v5.0.1, as identified in the following table.

Web Tools Compatibility Limitations

Launch Switch Environment	Problems
Firmware: Fabric OS v3.1.0+, v4.1.0+, or v5.0.1+ **Operating System**: Any supported operating system (with supported browser) **Browser**: Any supported browser (on supported operating system	**Issue**: When viewing the topology from Web Tools, if your initial login was a v3.1.0+, v4.1.0+, or v5.0.1+ switch and you view the topology from a switch with a previous version of the Fabric OS, there is no print function available in the Fabric Topology window. Web Tools v3.1.0+, v4.1.0+, and v5.0.1+ include a Print button in the Fabric Topology window; earlier versions do not. **Workaround**: If the Fabric Topology window does not include a Print button, right-click anywhere inside the window and select **Print** from the popup menu.

Launch Switch Environment	Problems
Firmware: Fabric OS v2.6.x **Operating System:** Solaris **Browser:** Mozilla	**Issue**: The Switch Admin does not launch correctly. • If you try to launch Switch Admin using Fabric OS v2.6.2 on a Solaris operating system with a Mozilla browser, a warning message is displayed, telling you to use the Netscape browser. • If you try to launch Switch Admin using Fabric OS v2.6.1 or earlier on a Solaris operating system with a Mozilla browser, the Switch Admin fails and no warning is displayed. **Workaround**: Although the Netscape browser is not supported by Web Tools for switches running Fabric OS v2.6.2, v3.1.2, or v4.2.0 or later, if you must access the Switch Admin on a switch running Fabric OS v2.6.x from a Solaris operating system, use the Netscape 4.77 browser.
Firmware: Version *prior* to Fabric OS v2.6.2, v3.1.2, or v4.2.0 with secure mode enabled **Operating System:** Solaris **Browser:** Mozilla	**Issue**: If you try to launch Switch Admin, Zoning, Fabric Watch, or High Availability Admin using firmware versions prior to v2.6.2, v3.1.2, or v4.2.0 on a Solaris operating system with a Mozilla browser, the browser might crash due to a buffer overflow problem with Mozilla. **Workaround**: Although the Netscape browser is not supported by Web Tools for switches running Fabric OS v2.6.2, v3.1.2, or v4.2.0 or later, if you must access the Switch Admin, Zoning, Fabric Watch, or High Availability Admin on a switch running firmware versions prior to v2.6.2, v3.1.2, or v4.2.0 or later from a Solaris operating system, use the Netscape 4.77 browser.
Firmware: Version *prior* to Fabric OS v2.6.2, v3.1.2, or v4.2.0a **Operating System:** Any supported operating system (with supported browser) **Browser:** Any supported browser (on supported operating system)	**Issue**: When trying to access a switch running firmware versions prior to Fabric OS v2.6.2, v3.1.2, or v4.2.0 from the launch switch, Switch Explorer will display a null pointer exception, and the SwitchInfo applet will not display; Switch Explorer does not work properly with switches running the latest firmware. **Workaround**: Use a launch switch running Fabric OS v2.6.2, v3.1.2, or v4.2.0 or later to access the switch.
Firmware: Version *prior* to Fabric OS v4.4.0 **Operating System:** Any supported operating system (with supported browser) **Browser**: Any supported browser (on supported operating system)	**Issue:** When trying to perform end-to-end monitoring (Brocade Advanced Performance Monitoring) on a local switch with a Fabric OS prior to v4.4.0, the SilkWorm 4100 is displayed as a 16-port switch. **Workaround:** For a SilkWorm 4100, use a launch switch running Fabric OS v4.4.0 or later to perform end-to-end monitoring on the switch.

Launch Switch Environment	Problems
Firmware: Version *prior* to Fabric OS v4.4.0 **Operating System:** Any supported operating system (with supported browser) **Browser:** Any supported browser (on supported operating system)	**Issue:** When trying to perform zoning on a local switch with a Fabric OS version prior to v4.4.0, the SilkWorm 4100 is displayed as a 16-port switch. **Workaround:** If you are running Brocade Secure Fabric OS, select a switch running Fabric OS v4.4.0 or later as the primary FCS switch. If you are not running Brocade Secure Fabric OS, use a launch switch running Fabric OS v4.4.0 or later to perform zoning on the switch.
Firmware: Version *prior* to Fabric OS v2.6.2, v3.1.2, or v4.2.0 **Operating System:** Solaris **Browser:** Netscape	**Issue**: Any switches running Fabric OS v2.6.2, v3.1.2, or v4.2.0 or later are unsupported through Netscape. **Workaround**: The Netscape browser is not supported by Web Tools for switches running Fabric OS v2.6.2, v3.1.2, or v4.2.0 or later. Use the Mozilla browser v1.6 to manage all of your switches from a Solaris operating system.
Firmware: Version *prior* to Fabric OS v2.6.1, v3.0.x, or v4.0.x **Operating System:** Windows **Browser:** Internet Explorer	**Issue**: When you are trying to run Fabric View with a large fabric, the browser might crash. **Workaround**: Use a launch switch that runs Fabric OS v2.6.1, v3.0.x, or v4.0.x or later so that you can use Switch Explorer (not Fabric View). Use a launch switch with v.2.6.2, v3.1.x, or v4.1.x or later.
Firmware: Fabric OS v5.0.1+ **Operating System**: Any supported operating system (with supported browser) **Browser**: Internet Explorer and Mozilla	**Issue:** If you upgrade from Fabric OS v4.x to v5.x, you must upgrade your Java plug-in version to v1.4.2_06 from any prior version installed on your system. **Workaround:** For Internet Explorer, before launching Web Tools, check your Java plug-in version. If you have a version lower than 1.4.2_06, then you must uninstall it. When you launch Web Tools and you see a warning about a missing plug-in, follow the prompts. This procedure will make sure that the correct plug-in version is actually installed. For Mozilla, follow the Mozilla Java plug-in installation instructions to install Java v1.4.2_06.

RFEs

The Fabric OS v5.0.3 release included the following Requests for Enhancement (RFEs):

RFE #	Description
3545	Log the history of firmware versions whenever there is a firmware downgrade or upgrade performed on the switch
2365	Change the message when issuing the HAfailover command so it queries the user about to "fail the control processor."
3273	Change the message for a successful login message in the event log so it shows the IP address of the station logging in.
3073	Change portlog event settings to be persistent across reboots.
2887	Allow the Switch Names to support the '-' character.
3474	Allow general users to determine the MAC address of an ethernet port.

Other Notes

The tables below list other important information about the SilkWorm 4012, SilkWorm 48000, and Fabric OS v5.0.x.

SilkWorm 4012	Description
Chassis	Early versions of the Silkworm 4012 (including units used for beta) have the potential to interfere with a Cisco GbE3 switch being removed or installed into the adjacent slot when the SilkWorm 4012 is present.

SilkWorm 48000	Description
Fan insertion for the Silkworm 4100	If a fan is marked as faulty (amber flashing LED on fan assembly) within a few seconds after insertion, it may be a false failure indication due to a momentary disconnection caused by uneven insertion (contact bounce). Restore the fan to an operational status as follows: 1) Pull the fan assembly out half way. 2) Reinsert the fan at a moderate pace with a steady application of moderate force until the fan assembly is seated securely. At this point, the fan should power up and the fan LED should indicate a functioning fan (green light). If the fan continues to indicate a fault (amber LED), then remove fan assembly and repeat procedure with a replacement fan assembly.
FDMI host name support	If you have HBAs that support FDMI exposure of host names in a fabric you will need Fabric OS v3.2.0a and v4.4.0d to ensure that the host names are properly propagated to v5.0.1 switches
PID 2 support for the SilkWorm 48000	The additional ports (128-258) on a SilkWorm 48000 require updates to certain Fabric OS releases in a special circumstance, that is, running PID-2 Format with a SilkWorm 48000 in the fabric. Minimum Fabric OS version required: 2.6.2d, 3.2.0a, 4.4.0d

SilkWorm 48000	Description
Power cycling	A minimum of 15 seconds between power cycles is required.
Proxy switches	If you are using a Fabric OS v4.x switch as an API or SMI-S proxy to manage a v5.0.1 switch, you will need Fabric OS v4.4.0d.
Secure Fabric OS support for the SilkWorm 48000	The additional ports (128-258) on a 48000 require updates to certain Fabric OS releases in special circumstances; that is, Secure Fabric OS fabric with a SilkWorm 48000 in the fabric and port numbers higher than 127 specified in DCC policies. Minimum Fabric OS version required: 2.6.2d, 3.2.0a, 4.4.0d
SilkWorm 48000 hardware updates	The cable management comb, located on the lower portion of the port side of the SilkWorm 48000 director, has been updated with a slightly modified design prior to final release. The new design includes a reduced length of comb lower deck by 25mm. Replacement of the cable comb is a simple process, requiring the removal and replacement of two screws. The lower deck is now at a 4.5-degree angle. These changes are required to support improved part manufacturability.

Fabric OS	Description
Advanced Performance Monitor	Adding Advanced Performance Monitor (**perfAddUserMonitor**) without zoning enabled at the same time will stop all frame traffic. The only frames that can go through are those that match the definitions in the **perfAddUserMonitor** command, in most cases, a very narrow definition. The result is that almost all traffic is blocked. Add Advanced Performance Monitor only when zoning is also enabled.
Diagnostics	Both backport and spinsilk tests are not supported for the Saturn platform, including any "mixed-bladed" platforms that include a Saturn blade type (FC4-16, FC4-32 or CP4). Instead, the user can run minicycle test. If minicycle is run from the burnin script and both lb_mode 1 and lb_mode 7 are selected, the user will get the same port frame passing coverage as spinsilk and backporttest.
Nondefault operands	**IMPORTANT:** The use of nondefault operands for diagnostic commands is recommended for advanced users and technical support only.
SNMP	Starting with the FOS 4.4.0 release, Brocade added the ability to enable traps on a more granular level. After an upgrade, the **snmpMibCapSet** command should be run from the CLI to update the settings. This allows additional flexibility in controlling SNMP traps. The default setting is for all traps to be disabled.
Upgrade	Fabric OS v5.0.0 is superceded by version v5.0.1, and you are strongly encouraged to upgrade to v5.0.1.

Fabric OS	Description
Upgrading / downgrading	When considering an upgrade to a later Fabric OS release the user should save the zone database configuration immediately following the upgrade. Changes to the zoning database can then be conducted. If you are considering downgrading to the prior Fabric OS release, remember to clear the zoning database then restore the saved zoning database configuration prior to the downgrade.
Upgrading to Fabric OS v4.2.0 to v5.0.1	The SilkWorm and FA traps in pre-Fabric OS v4.4.0 code were turned on and off as a group; and it wasn't possible to set individual SilkWorm or FA traps. In v4.4.0 the ability to turn traps on and off individually was added. That means, that individual traps need to be turned on explicitly after the corresponding trap group is turned on.

After the upgrade from Fabric OS v4.2.0 firmware, individual traps are turned off by default even if the corresponding trap group was turned on before the upgrade. Therefore if you have been previously monitoring these traps, you need to use either **snmpMibCapSet** or the newer snmpconfig command to turn the desired traps on individually. |
| FM Firmware Download | For an embedded switch running Fabric OS 5.0.2 and 5.0.3 and using the small firmware file format, Fabric Manager (FM) version 5.0 supports firmware download using the external ftp server to download the firmware to the switch. In FM's options dialog, set the FTP server setting to the external option. In the firmware download dialog, select the external ftp server to download. FM does not support importing the image into the FM firmware repository. |
| Zoning | With AUDIT logging enabled, while performing zoning changes via CLI, an additional audit log from HTTP may also appear along with the audit logs from zoning. This message does not always appear, and when it does, it represents redundant reporting by the CAL layer. |
| Loss of sync between Emulex HBA and Brocade 4Gb/sec switch | **Issue:** If there is a loss of sync forcing a link to be re-esatablished, it is possible that links between Emulex HBAs and 4Gb Brocade switches may not automatically be re-established. This issue could occur after an error that has forced the switch and HBA to re-establish link initialization such as a cold switch reboot.

Workaround: Use the command *portCfgGPort* to congfigure the switch port in point-to-point only mode, also known as G port mode. To configure the HBA to point-to-point mode, please refer to Emulex HBAnyware™ documentation. To re-establish the link on the affected port without traffic disruption on other ports, issue the comands *portDisable* and *portEnable* commands on the affected port. |

Documentation Updates

This section provides information on additions and corrections to the documentation.

This release of Fabric OS is supported by the v5.0.1 documentation set. The most recent Fabric OS v5.0.1 documentation is available on the Brocade Connect Web site:

http://www.brocadeconnect.com/

Fabric OS Administrator's Guide
(Publication number 53-0000518-07)

On page B-2, in the section "Supported Brocade Features," add the following text to the bullet statement:

- Brocade translative mode

  Registers private storage target devices into the fabric, it can be used in a heterogeneous fabric if the devices are connected directly to Brocade switches. The devices will be accessible from any port on the fabric.

  Note
  Switches with a Condor ASIC do not support translative mode.

On page 3-17, in the section "To enable or disable RADIUS service," add the following:

  Warning
  When you issue **aaaConfig --radius on**, all sessions in which you are logged on are logged off immediately, and local authentication is disabled.

On page 4-8, in the section "Considerations for Downgrading Firmware," add the following:

- Do not attempt to perform a firmware downgrade from v5.0.1 to v4.2.2 when you have a zone configuration larger than 128K.

On page 4-11, at the end of the section "Upgrading SilkWorm Directors," add the following:

  Caution
  To successfully download firmware to a director you must have an active Ethernet connection on *both* CPs.

On page 5-5, in Table 5-1: SilkWorm Director Terminology and Abbreviations, remove the following rows:

Term	Abbreviation	Blade ID	Description
D1 Chassis	n/a	n/a	The first generation chassis. These chassis have a manufacture date prior to January 1, 2004. In Fabric OS 5.0.1, use the chassisShow command to view the backplane revision number for this chassis, 0x1F.
D2 Chassis	n/a	n/a	The second generation chassis. These chassis have a manufacture date from to January 1, 2004 to May 1, 2005. In Fabric OS 5.0, use the chassisShow command to view the backplane revision number for this chassis, 0x1D.
D3 Chassis	n/a	n/a	The third generation chassis. These chassis have a manufacture date from to May 1, 2005 to the present. In Fabric OS 5.0, use the chassisShow command to view the backplane revision number for this chassis, 0x1B.

On page 6-2, in the section "Specifying the Routing Policy," add the following text to the bullet statement on port-based path selection:

- Port-based path selection

Default on SilkWorm 3016, 3250, 3850, 3900, 12000, 24000, and 48000 (using configuration option 1). These switches support the port-based policy only; you cannot change the routing policy for these switches. SilkWorm 200E, 4012, and 4100 switches can also use port-based routing. The default (and only) routing policy used in FICON environments is port-based routing.

In Chapter 7, "Administering FICON Fabrics," add the following:

N-Port ID Virtualization

N-Port ID Virtualization (NPIV) requires an N_Port ID Virtualization license on the switch *for Fabric OS v5.0.x. The NPIV solution for this release is supported for the FCP traffic connecting to the mainframe environment.* The NPIV license must be installed before NPIV functionality can be enabled on any port. For Bloom-based switches and port blades supporting FICON (SilkWorm 3900, 12000 and 24000), the default behavior is that NPIV is disabled for every port. For Condor-based switches and port blades (SilkWorm 4100 and 48000), the default behavior is that NPIV is enabled for every port.

The following example shows the license required for NPIV:

```
switch:admin> licenseshow
R9cRceRSdSEdSdn:
N_Port ID Virtualization license
```

Use the **portCfgNPIV** command to enable or disable NPIV on a port-by-port basis.

The following example shows NPIV being enabled on port 10 on a SilkWorm 4100:

```
switch:admin> portCfgNPIVPort 10, 1
```

The **portCfgShow** command shows the NPIV capability of switch ports. The following example shows whether or not a port is configured for NPIV:

```
switch:admin> portcfgshow
Ports of Slot 0    0  1  2  3    4  5  6  7    8  9 10 11   12 13 14 15
-----------------+--+--+--+--+----+--+--+--+----+--+--+--+----+--+--+--
Speed             AN AN AN AN   AN AN AN AN   AN AN AN AN   AN AN AN AN
Trunk Port        ON ON ON ON   ON ON ON ON   ON ON ON ON   ON ON ON ON
Long Distance     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
VC Link Init      .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Locked L_Port     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Locked G_Port     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Disabled E_Port   .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
ISL R_RDY Mode    .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
RSCN Suppressed   .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Persistent Disable.. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
NPIV capability   ON ON ON ON   ON ON ON ON   ON ON ON ON   ON ON ON ON
```

The output of the commands **switchShow** and portShow show NPIV information for a given port. If a port is an F_Port, and you enter the **switchShow** command, then the port WWN of the N_Port is returned. For an NPIV F_Port, there are multiple N_Ports, each with a different port WWN. The **switchShow** command output indicates whether or not it is an NPIV F_Port, and identifies the number of virtual N_Ports behind it. Following is sample output from the **switchShow** command:

```
    switch: admin> switchshow
switchName:      swd77
switchType:      32.0
switchState: Online
switchMode:      Native
switchRole:      Principal
switchDomain: 99
switchId: fffc63
switchWwn:       10:00:00:05:1e:35:37:40
zoning:          OFF
```

```
switchBeacon:    OFF

Area Port Media Speed State
==============================
  0   0   id    N2   Online    F-Port  50:05:07:64:01:20:73:b8
  1   1   id    N2   Online    F-Port  50:05:07:64:01:60:73:b8
  2   2   id    N2   Online    F-Port  65 NPIV public
  3   3   id    N2   Online    F-Port  50:05:07:64:01:e0:73:b8
  4   4   id    N2   Online    F-Port  50:05:07:64:01:20:73:b5
  ...

  <output truncated>
```

The **portShow** command shows the NPIV attributes and lists all the N_Port (physical and virtual) port WWNs under "*portWwn of device(s) connected.*" Use the **portLoginShow** command to display the login information for the virtual PIDs of a port. Following is sample output for **portShow** and **portLoginShow**:

```
switch:admin> portshow 2
portName: 02
portHealth: HEALTHY

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x24b03      PRESENT ACTIVE F_PORT G_PORT NPIV LOGICAL_ONLINE LOGIN
NOELP LED ACCEPT
portType:  10.0
portState: 1    Online
portPhys:  6    In_Sync
portScn:   32   F_Port
port generation number:    148
portId:    630200
portIfId:   43020005
portWwn:   20:02:00:05:1e:35:37:40
portWwn of device(s) connected:
   c0:50:76:ff:fb:00:16:fc
   c0:50:76:ff:fb:00:16:f8
   ...

        <output truncated>
     ...
   c0:50:76:ff:fb:00:16:80
   50:05:07:64:01:a0:73:b8
Distance:  normal
portSpeed: N2Gbps

Interrupts:        0         Link_failure: 16         Frjt:        0
Unknown:           0         Loss_of_sync: 422        Fbsy:        0
Lli:               294803    Loss_of_sig:  808
Proc_rqrd:         0         Protocol_err: 0
Timed_out:         0         Invalid_word: 0
Rx_flushed:        0         Invalid_crc:  0
Tx_unavail:        0         Delim_err:    0
Free_buffer:       0         Address_err:  1458
Overrun:           0         Lr_in:        15
Suspended:         0         Lr_out:       17
Parity_err:        0         Ols_in:       16
2_parity_err:      0         Ols_out:      15
CMI_bus_err:       0

switch:admin> portloginshow 2
Type  PID      World Wide Name        credit df_sz cos
=====================================================
  fe  630240  c0:50:76:ff:fb:00:16:fc   101   2048   c   scr=3
```

```
fe  63023f  c0:50:76:ff:fb:00:16:f8   101   2048   c   scr=3
fe  63023e  c0:50:76:ff:fb:00:17:ec   101   2048   c   scr=3
...
<output truncated>
...
ff  630202  c0:50:76:ff:fb:00:17:70   192   2048   c   d_id=FFFFFC
ff  630201  c0:50:76:ff:fb:00:16:80   192   2048   c   d_id=FFFFFC
```

Also note the following behaviors relating to NPIV:

- The LIRR database is not updated when a virtual port requests a logout (LOGO). The LIRR database is updated when the base port logs out or the port goes offline.
- An NPIV port can have a maximum of 255 virtual PIDs per port. By default, 126 virtual PIDs are set for each port.
- There is no limit on the maximum number of virtual PIDs a switch can support. The maximum number of virtual PIDs per port is the limiting factor.
- Each NPIV device is transparent to the user, and has its own device PID, Port WWN, and Node WWN, and should act the same as all other physical devices in the fabric. The same zoning rules apply to NPIV devices as non-NPIV devices. Zones can be defined by (domain, port) and/or by WWN zoning.
  Note:
  To perform zoning to the granularity of the virtual N_Port IDs, you must use WWN-based zoning.

On page 7-1, in the section "FICON Overview," replace the IBM Redbook Reference with the following:

Refer to the IBM Redbook, *FICON® Implementation Guide* (SG24-6497-00)

On page 7-1, in the section "FICON Overview," add the following:

NPIV (N-port ID Virtualization) operation requires a Brocade N_Port ID Virtualization license.

On page 7-2, in the section "FICON Overview," add the following:

FICON is supported for the SilkWorm 4100 and 48000 platforms.

The FICON protocol is supported on the following SilkWorm models and Fabric OS releases:

- SilkWorm 48000, Fabric OS v5.0.1b or later. A single-domain configuration is supported with a mix of 16-port and 32-port SilkWorm 48000 port blades in a SilkWorm 48000. Dual-domain configurations are not supported on the SilkWorm 48000. Mixed port blade configurations of SilkWorm 24000 and SilkWorm 48000 port blades (FC2-16, FC4-16 or FC4-32) in the same director is not supported in a FICON environment.

- SilkWorm 4100, Fabric OS v5.0.1b or later.

Both the SilkWorm 48000 and 4100 require the port-based routing policy either in a single-switch configuration or a cascaded-switch configuration on those switches in the fabric that have FICON devices attached (option 1 of the **aptPolicy** command). Other switches in the fabric may use the default exchange-based routing policy (option 3 of the **aptPolicy** command) only when Open Systems devices are attached to those switches.

```
CUP is supported on the SilkWorm 4100 and 48000 running Fabric OS v5.0.1b or
later.
```

On page 7-4, in the section "Configuring Switches," add the following to the recommended FICON environment configuration settings:

The port-based routing policy is recommended for the SilkWorm 4100 and 48000 on any switch that has FICON devices attached. Other switches in the fabric with Open Systems devices exclusively can still use exchange-based routing.

Some 1-Gbit/sec storage devices cannot auto-negotiate speed with the SilkWorm 48000 or 4100 ports. For these types of devices, configure ports that are connected to 1-Gbit/sec storage devices for fixed 1-Gbit/sec speed.

On page 7-4, in the section "Preparing a Switch," add the following to step 2:

- **pkiShow** to determine the existence of PKI objects, such as switch private key, private key passphrase, CSR, root certificate, and switch certificate. If none of these objects exist, refer to the *Secure Fabric OS Administrator's Guide* for information about creating the PKI objects and obtaining the digital certificate file.

On page 7-4, in the section "Preparing a Switch," add a new step after step 3:

4. Change the routing policy on the switch from the default exchange-based policy to the required port-based policy for those switches with FICON devices directly attached. For the SilkWorm 4100, refer to the *Fabric OS Command Reference Manual* for details about the **aptPolicy** command. For the SilkWorm 48000, refer to Chapter 12 of the *WebTools Administrator's Guide.*

On page 7-9, in the section "Setup Summary," add the following to step 2:

For SilkWorm 48000 only: Use the **portDisable** command to disable (block) port 126.

Port 126 is not supported in a CUP environment. After **fmsmode** has been successfully enabled, port 126 remains disabled. It cannot be used either as an F_Port or an E_Port. Because port 126 is not available after enabling **fmsmode**, you should first move any fiber connected to port 126 to another free port.

In Table 7-2, on page 7-12, replace the HCP default setting from "1 (on)" to "0 (off)".

On page 7-19, in the section "Sample IOCP Configuration File for SilkWorm 3900, 12000, and 24000 Switches," replace the IBM Redbook reference with the following:

For more information, refer to the IBM Redbook publication FICON® Implementation Guide (SG24-6497-00) section 2.7.1 on switch numbering.

On page 7-20, in the section "Sample IOCP Configuration File for SilkWorm 3900, 12000, and 24000 Switches," add the following switches to the list:

- SilkWorm 4100

- SilkWorm 48000

On page 11-3, replace the section "Choosing an Extended ISL Mode" with the following text:

Choosing an Extended ISL Mode

Table 11-1 lists the extended ISL modes for switches that have a Bloom ASIC. You can configure extended ISL modes with the **portCfgLongDistance** command when the Extended Fabrics license is activated.

Table 11-1 Extended ISL Modes: Switches with Bloom ASIC

Mode	Description	Buffer Allocation		Distance @ 1 Gbit/sec	Distance @ 2 Gbit/sec	Earliest Fabric OS Release	Extended Fabrics License Required?
		1 Gbit/ sec	2 Gbit/ sec				
L0	Level 0 static mode, the default	5 (26)[b]	5 (26)	10 km	5 km	All	No
LE	Level E static mode, supports links beyond 5 km	13	19	n/a	10 km	v3.x, v4.x	No
L0.5	Level 0.5 static mode (designated LM when listed with the **portcfgshow** command)	19	34	25 km	25 km	v3.1.0, v4.1.0, v4.x, v5.x	Yes
L1	Level 1 static mode	27	54	50 km	50 km	All	Yes
L2	Level 2 static mode	60	64	100 km	60 km	All	Yes
LD[a]	Dynamic mode uses automatic distance detection for a user-specified distance	Auto	Auto	Auto	Auto	V3.1.0, v4.1.0, v4.4.0, v5.x (depending on the model)	Yes

a. The dynamic long-distance mode (LD) automatically configures the number of buffer credits required, based on the actual link distance.
b. For each data channel (in this case, there are 4) there are 5 credits, plus 6 extra credits.

Table 11-2 lists the extended ISL modes for switches that have a Goldeneye ASIC.

Table 11-2 Extended ISL Modes: Switches with Goldeneye ASIC (SilkWorm 200E only)

Mode[a]	Buffer Allocation			Distance @ 1 Gbit/ sec	Distance @ 2 Gbit/ sec	Distance @ 4 Gbit/ sec	Earliest Fabric OS Release	Extended Fabrics License Required?
	1 Gbit/ sec	2 Gbit/ sec	4 Gbit/ sec					
L0	3 (17)[b]	3 (17)	3 (17)	6 km	3 km	1 km	All	No
LE	10	15	25	n/a	10 km	10 km	v3.x, v4.x	No

a. No LD mode appears because the distances for SilkWorm 3016 and 4012 vary depending on the number of internal and external ports.

b. For each data channel (in this case, there are 4) there are 5 credits, plus 6 extra credits.

Table 11-3 lists the extended ISL modes for switches that have a Condor ASIC.

Table 11-3 Extended ISL Modes: Switches with Condor ASIC.

Mode	Buffer Allocation			Distance @ 1 Gbit/sec	Distance @ 2 Gbit/sec	Distance @ 4 Gbit/sec	Earliest Fabric OS Release	Extended Fabrics License Required?
	1 Gbit/sec	2 Gbit/sec	4 Gbit/sec					
L0	5 (26)[b]	5 (26)	5 (26)	10 km	5 km	2 km	All	No
LE	11	16	26	n/a	10 km	10 km	V3.x, v4.x	No
L0.5	18	31	56	25 km	25 km	25 km	v3.1.0, v4.1.0, v4.x, v5.x	Yes
L1	31	56	106	50 km	50 km	50 km	All	Yes
L12	56	106	206	100 km	100 km	100 km	All	Yes
LD[a]	Auto	Auto	Auto	Auto	Auto	Auto	v3.1.0, v4.1.0, v4.4.0, v5.x (depending on the model)	Yes

a. The dynamic long-distance mode (LD) automatically configures the number of buffer credits required, based on the actual link distance.

b. For each data channel (in this case, there are 4) there are 5 credits, plus 6 extra credits.

For dynamic long distance links, you can approximate the number of buffer credits using the following formula:

*Buffer credits = [(distance in km) * (data rate) * 1000] / 2112*

The data rate is 1.0625 for 1 Gbit/sec, 2.125 for 2 Gbit/sec, and 4.25 for 4 Gbit/sec and Fibre Channel. This formula provides the minimum number of credits that will be allocated to a given port; the actual number will likely be higher.

On page 12-9, in the section "Trunking Over Extended Fabrics," add the following section:

Trunking Distances

Enhanced trunking support for SilkWorm 200E (which supports LE mode only) is summarized in Table 12-1.

Table 12-1 Trunking Support for SilkWorm 200E (Goldeneye ASIC)

Mode	Distance	Number of 2Gbit/sec ports	Number of 4 Gbit/sec ports
LE	10 km	16 (four 4-port trunks)	8 (two 4-port trunks)

Enhanced trunking support for the SilkWorm 4012 (Goldeneye ASIC) is summarized in Table 12-2.

Table 12-2 Trunking Support for the SilkWorm 4012 (Goldeneye ASIC)

Mode	Distance	Number of 2Gbit/sec ports	Number of 4 Gbit/sec ports
LE	10 km	4 (one 4-port trunk)	4 (one 4-port trunk)
L0.5	25 km	4 (one 4-port trunk)	3 (one 3-port trunk)
L1	50 km	3 (one 3-port trunk)	1 (one 1-port trunk)
L2	100 km	1 (one 1-port trunk)	0
LD	200 km	0	0
LD	250 km	0	0
LD	500 km	0	0

Enhanced trunking support for the SilkWorm 3900 (Bloom and Bloom2 ASICs) is summarized in Table 12-3.

Table 12-3 Trunking Support for the SilkWorm 3900 (Bloom and Bloom2 ASICs)

Mode	Distance	Number of 2 Gbit/sec ports
LE	10 km	4 (one 4-port trunk)
L0.5	25 km	3 (one 3-port trunk)
L1	50 km	1 (one 2-port trunk)
L2	100 km	0
LD	200 km	0
LD	250 km	0
LD	500 km	0

Enhanced trunking support for the SilkWorm 4100 (Condor ASIC) is summarized in Table 12-4.

Table 12-4 Trunking Support for the SilkWorm 4100 (Condor ASIC)

Mode	Distance	Number of 2Gbit/sec ports	Number of 4 Gbit/sec ports
LE	10 km	32 (four 8-port trunks)	32 (four 8-port trunks)
L0.5	25 km	32 (four 8-port trunks)	15 (one 8-port trunk)
L1	50 km	15 (one 2-port trunk)	7 (one 7-port trunk)
L2	100 km	7 (one 7-port trunk)	3 (one 3-port trunk)
LD	200 km	3 (one 3-port trunk)	0
LD	250 km	3 (one 3-port trunk)	0
LD	500 km	0	0

Enhanced trunking support for the SilkWorm 48000 is summarized in Table 12-5 and 12-6.

Table 12-5 Trunking Support for FC4-16 port blades (SilkWorm 48000)

Mode	Distance	Number of 2Gbit/sec ports	Number of 4 Gbit/sec ports
LE	10 km	16 (two 8-port trunks)	16 (two 8-port trunks)
L0.5	25 km	16 (two 8-port trunks)	12 (one 8-port trunk, one 4-port trunk)
L1	50 km	12 (one 8-port trunk, one 4-port trunk)	5 (one 5-port trunk)
L2	100 km	5 (one 5-port trunk)	2 (one 2-port trunk)
LD	200 km	2 (one 2-port trunk)	0
LD	250 km	2 (one 2-port trunk)	0
LD	500 km	0	0

Table 12-6 Trunking Support for FC4-32 port blades (SilkWorm 48000)

Mode	Distance	Number of 2Gbit/sec ports	Number of 4 Gbit/sec ports
LE	10 km	32 (four 8-port trunks)	32 (four 8-port trunks)
L0.5	25 km	32 (four 8-port trunks)	26 (two 8-port trunks, two 5-port trunks)
L1	50 km	26 (two 8-port trunks, two 5-port trunks)	12 (two 6-port trunks)
L2	100 km	12 (two 6-port trunks)	6 (two 3-port trunks)
LD	200 km	6 (one 2-port trunk)	0
LD	250 km	4 (two 2-port trunks)	0
LD	500 km	0	0

On page 13-15, at the end of the section "Creating and Maintaining Zones," add the following text:

Merging Zones

Before linking two switches together, it is important that you know the zone database limit of adjacent switches. For example, when switches running Fabric OS v3.2, v4.4.0, or v5.x discover that the zone merge database is larger than its pre-determined zone database size limit, they issue a reject notification before symmetrically segmenting their own ends of the ISL, thereby preventing the new switch from joining the fabric.

Symmetrical segmentation occurs when both ends of an ISL are shut down. Subsequently, no frames are exchanged between those two switches.

Asymmetrical segmentation not only prevents frames from being exchanged between switches, but also causes routing inconsistencies.

The best way to avoid either type of segmentation is to know the zone database size limit of adjacent switches. The following tables provide the expected behavior based on different database sizes after a zone merge is specified.

Table 1 Resulting Database Size: 0 to 96K

Receiver / Initiator	FOS v2.6	FOS v3.1	FOS v3.2	FOS v4.0/ v4.1/v4.2	FOS v4.3/ v4.4.0	FOS v5.0.0/ v5.0.1	Fibre Channel Router	XPath v7.3
FOS v2.6/v3.1	Join	Join	Join	Join	Join	Join	Join	Join
FOS v3.2	Join	Join	Join	Join	Join	Join	Join	Join
FOS v4.0/v4.1/ v4.2	Join	Join	Join	Join	Join	Join	Join	Join
FOS v4.3/v4.4.0	Join	Join	Join	Join	Join	Join	Join	Join
FOS v5.0.0/v5.0.1	Join	Join	Join	Join	Join	Join	Join	Join
Fibre Channel Router	Join	Join	Join	Join	Join	Join	Join	Join
XPath v7.3	Join	Join	Join	Join	Join	Join	Join	Join

Table 2   Resulting Database Size: 96K to 128K

Receiver / Initiator	FOS v2.6	FOS v3.1	FOS v3.2	FOS v4.0/ v4.1/v4.2	FOS v4.3/ v4.4.0	FOS v5.0.0/ v5.0.1	Fibre Channel Router	XPath v7.3
FOS v2.6/v3.1	Segment	Segment	Segment	Segment	Segment	Segment	Join	Segment
FOS v3.2	Segment	Segment	Join	Join	Join	Join	Join	Join
FOS v4.0/v4.1/ v4.2	Segment	Segment	Segment	Join	Join	Join	Join	Join
FOS v4.3/v4.4.0	Segment	Segment	Join	Join	Join	Join	Join	Join
FOS v5.0.0/v5.0.1	Segment	Segment	Join	Join	Join	Join	Join	Join
Fibre Channel Router	Join	Join	Join	Join	Join	Join	Join	Join
XPath v7.3	Segment	Segment	Segment	Join	Join	Join	Join	Join

Table 3   Resulting Database Size: 128K to 256K

Receiver / Initiator	FOS v2.6	FOS v3.1	FOS v3.2	FOS v4.0/ v4.1/v4.2	FOS v4.3/ v4.4.0	FOS v5.0.0/ v5.0.1	Fibre Channel Router	XPath v7.3
FOS v2.6/v3.1	Segment	Segment	Segment	Segment	Segment	Segment	Join	Segment
FOS v3.2	Segment	Segment	Join	Segment	Join	Join	Join	Segment
FOS v4.0/v4.1/ v4.2	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS v4.3/v4.4.0	Segment	Segment	Join	Segment	Join	Join	Join	Segment
FOS v5.0.0/v5.0.1	Segment	Segment	Join	Segment	Join	Join	Join	Segment
Fibre Channel Router	Join	Join	Join	Segment	Join	Join	Join	Segment
XPath v7.3	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment

Table 4 Resulting Database Size: 256K to 1M

Receiver / Initiator	FOS v2.6	FOS v3.1	FOS v3.2	FOS v4.0/ v4.1/v4.2	FOS v4.3/ v4.4.0	FOS v5.0.0/ v5.0.1	Fibre Channel Router	XPath v7.3
FOS v2.6/v3.1	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS v3.2	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS v4.0/v4.1/ v4.2	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS v4.3/v4.4.0	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS v5.0.0/v5.0.1	Segment	Segment	Segment	Asymmetrical Segment	Segment	Join	Join	Segment
Fibre Channel Router	Segment	Segment	Segment	Segment	Segment	Join	Join	Segment
XPath v7.3	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment

Fabric OS Command Reference Manual
(Publication number 53-0000519-10)

For **perfClearAlpaCrc**, in the "Description" section, add the following note:

> This command is not supported on the SilkWorm 200E, 4100, and 48000 platforms.

For all Performance Monitor (**perf***) commands, remove "user" from the "Synopsis" sections. These commands support only admin and switchAdmin access.

In Fabric OS v5.0.1b or later, the **portCfgNPIVPort** command is supported.

For **haFailover**, in the "Description" section, remove the following sentence:

> Because **haFailover** results in an active CP reboot, a warning message and confirmation are displayed. If the user confirms, the failover takes place.

Also for **haFailover**, in the "Example" section, replace the existing example with the following:

> To force the failover of the active CP to the standby CP in the switch:

```
switch:admin> hafailover
Local CP (Slot 6, CP1): Active, Warm Recovered
Remote CP (Slot 5, CP0): Standby, Healthy
HA enabled, Heartbeat Up, HA State synchronized

Warning: This command is being run on a redundant control
processor(CP) system. If the above status does not indicate
'HA State synchronized', then the CPs are not synchronized
and this operation will cause the active CP to reset. This
will cause disruption to devices attached to both switch 0
and switch 1 and will require that existing telnet sessions
be restarted. To reboot a single logical switch on this
system, use the switchreboot command while logged in to
that logical switch.

Are you sure you want to fail over to the standby CP [y/n]?
```

In Chapter 2, "Fabric OS Commands," remove the following commands:

- **diagEsdPorts**

- **portCfgMcastLoopback**

On page 2-21, add the following note to the **aptPolicy** "Description" section:

"Note: This command is supported only on SilkWorm 200E, 4012, 4100, and 48000 platforms."

"Note: On the SilkWorm 4100 all three options can be changed from the Command Line interface. On SilkWorm 48000 platforms in chassis mode 5, only options 2 and 3 can be changed at the Command Line."

On page 2-36, 2-107, and 2-108, for the **burninErrClear**, **diagSetBurnin**, and **diagSetCycle** commands, respectively, add the following note to the "Description" sections:

"It is advisable to run the **burninErrClear** command prior to running **diagSetBurnin** and **diagSetCycle**."

On page 2-76, remove the reference to "fabric.ops.mode.vcEncode: 0" from the **configShow** output in the "Example" section.

On page 2-85, remove the HTTP and RPCd content from Table 2-7 for the **configure** command.

On page 2-102, add the following note to the **diagHelp** "Description" section:

"Use default operands when running diagnostics commands. Nondefault settings require detailed knowledge of the underlying hardware and are intended for support personnel only. Contact support if you want to use these operands."

On page 2-176, in the **firmwareDownloadStatus** "Example" section, change the two instances of "It may take up to 10 minutes." to "This step will take up to 30 minutes."

On page 2-224, in the **ifModeShow** "Example" section, replace the existing output with the following:

```
switch:admin> ifmodeshow eth0
Link mode: negotiated 100baseTx-HD, link ok
MAC Address: 00:60:69:D0:24:40
```

On page 2-348, for **portCfgLPort**, add the following to the end of the "Description" section:

For Bloom ASIC-based platforms, a device port configured as an L_Port supports E_Ports; however, this configuration is not recommended. In later ASIC designs, this configuration is no longer supported.

On page 2-496, add the following note to the **slTest** "Description" section:

"NOTE: This command is supported only on SilkWorm 200E, 4012, 4100, 4900, and 48000 platforms"

On page 2-544, in the **switchName** "Operands" section, add the hyphen character to the list of supported characters for switch names.

On page 2-318, 2-324, and 2-327, change the availability for **perfMonitorShow, perfShowEEMonitor,** and **perfShowFilterMonitor** from all "all users" to "admin".

On page 2-345, replace the **portCfgLongDistance** "Description" section with the following:

"Use this command to allocate enough full-size frame buffers on a particular port to support a long-distance link up to 500 km. The port can be used as an F/FL/E_Port. F/FL_Ports can be configured only for long distance using LE, L0.5, L1, or L2 modes. Changes made by this command are persistent across switch reboots or power cycles.

The value of *distance_level* can be one of the following (the numerical value representing each *distance_level* is shown in parentheses):

L0 (0) Reconfigure the port to be a regular switch port. A total of 26 full-size frame buffers are reserved for data traffic, regardless of the port's operating speed.

L0.5 Level 0.5 (**portCfgShow** displays the two-letter code as LM) long distance, up to 25 km.

L1 (1) Level 1 long distance, up to 50 km.

L2 (2) Level 2 long distance, up to 100 km. For previously released switches (Bloom1-based), the number of frames buffers is limited to 63.

LE (3) Level E mode is for E_Ports for distances beyond 5 km and up to 10 km. LE does not require an Extended Fabrics license.

LD Automatic long-distance configuration. The buffer credits for the given E_Port are automatically configured, based on the actual link distance. Up to a total of 250 full-size frame buffers are reserved, depending upon the distance measured during E_Port initialization. If the desired distance is provided, it is used as the upper limit to the measured distance. For Bloom1-based systems, the number of frame buffers is limited to 63.

A long-distance link also can be configured to be part of a trunk group (refer to **portCfgTrunkPort**). Two or more long-distance links in a port group forms a trunk group when they are configured for the same speed, the same distance level, and their link distances are nearly equal.

Note: For details about buffer allocation at specific speeds and distances, refer to the "Administering Extended Fabrics" chapter of the *Fabric OS Administrator' s Guide*.

The *vc_translation_link_init* option is used to enable the long-distance link initialization sequence.

desired_distance is a required parameter to configure a port as an LD-mode link. The desired distance is used as the upper limit of the link distance to calculate buffer availability for other ports in the same port group. When the measured distance is more than *desired_distance*, the *desired_distance* is used to allocate the buffers. In this case, the port operates in degraded mode instead being disabled due to insufficient buffers.

Pressing **Ctrl-D** cancels the configuration update.

When a port is configured to be a long-distance port, the output of **portShow** and **switchShow** displays the long-distance level. In the **portShow** output, the long-distance level is indicated as follows:

- L0　　　normal
- LE　　　standard <= 10 km
- LM　　　medium long <= 25 km
- L1　　　long <= 50 km
- L2　　　super long <= 100 km
- LD　　　auto

In the **switchShow** output, the long distance mode displays as L*x*, where *x* is the second letter in two-letter distance-level code described earlier; however, L0.5 mode displays LM.

Note: The **portCfgISLMode** and **portCfgLongDistance** mode cannot both be enabled at the same time; otherwise, fabric segmentation occurs.

If a port is configured as a long distance port, the remaining ports of that port group could be disabled, fail to initialize, or move to "buffer limited" mode due to a lack of frame buffers. SilkWorm 3014, 3016, 3250, 3850, and 3900 switches and 12000 and 24000 directors do not support "buffer limited" mode and can have up to four ports per port group. SilkWorm 200E, 4100, and 4012 switches support "buffer limited" mode and can have up

to eight ports per port group. On SilkWorm 48000 directors, the FC4-16 and FC4-32 port blades support "buffer limited" mode and can have up to four ports per port group."

On page 2-409 and 2-541, change the availability for **portSwapShow** and **supportShowCfgShow** from "admin" to "all users".

On page 2-496, add the following note to the "Description" section:

"**Note:** This command is supported only on SilkWorm 200E, 4012, 4100, 4900, and 48000 platforms."

In Chapter 5, "MUA-Based Roles," add the following to table 5-1:

Command	Description
burninLevel	Sets the diagnostics burn-in level.
burninStatus	Displays the diagnostics burn-in level.
configDownload	Downloads a switch configuration file from a host file, omitting zoning and security configurations.
errModuleShow	Displays all the defined error log modules.
fabricLog	Displays or manipulates the fabric log.
fabStateResize	Changes the number of state entries.
historyMode	Displays the mode of the history log.
minisPropShow	Displays ASIC pair properties.
portCfg	Sets a port's configuration to be disabled or enabled.
setEsdMode	Enables or disables ESD mode.
setGbicMode	Enables or disables media mode.
setMediaMode	Enables or disables media mode.
setModem	Enables or disables modem dial-in to a control processor (CP).
setSfpMode	Enables or disables media mode.
supportShowCfgDisable	Disables a group of commands under the **supportShow** command.
supportShowCfgEnable	Enables a group of commands under the **supportShow** command.
supportShowCfgShow	Displays the groups of commands enabled for display by the **supportShow** command.
traceDump	Displays, initiates, or removes a trace dump.
traceFtp	Displays, enables, or disables the trace auto-FTP or FTPs a trace dump file to the customer FTP server.
traceTrig	Sets, removes, or displays trace triggers.
voltShow	Displays current level of the voltage sensors on a system.

In table 5-1, remove the following commands:

- backplaneTest
- backport
- bladeBeacon
- bladeDisable
- bladeEnable
- camTest
- centralMemoryTest
- crossPortTest
- fanDisable
- fanEnable
- ficonHelp
- filterTest
- haDisable
- haDump
- haEnable
- haFailover
- haShow
- haSyncStart
- haSyncStop
- itemList
- loopPortTest
- miniCycle
- powerOffListSet
- powerOffListShow
- spinFab
- spinJitter
- spinSilk
- statsClear
- statsTest
- switchReboot
- switchShutdown
- switchStart
- turboRamTest
- txdPath
- userRename

The following table lists platform support for legacy and new diagnostic commands.

Diagnostic Command	Supported SilkWorm Platforms
backplaneTest	3014, 3016, 3250, 3850, 3900, 12000, 24000
camTest	3014, 3016, 3250, 3850, 3900, 12000, 24000
centralMemoryTest	3014, 3016, 3250, 3850, 3900, 12000, 24000

Add the following paragraph to the **switchShow** on page 2-551:

> **"Note**:
> For all Bloom or Bloom2 based switches with Fabric OS v5.0.1 firmware, private device targets are displayed in **switchShow**. For Condor or Goldeneye based switches, private device targets are not displayed in **switchShow**."

Fabric OS MIB Reference Manual
(Publication number 53_0000521_09)

Add the following section at the end of Chapter 1.

Firmware Upgrades and Enabled Traps

Prior to Fabric OS v4.4, traps were turned on and off as a group (for example, the SW-Trap, or FA-Trap). In these versions of the Fabric OS it was not possible to set individual traps (such as, swSensorStatusChangeTrap, swTrackChangesTrap, or connUnitEventTrap).

In Fabric OS v4.4 or above you can to turn on and off traps individually within a trap group. The individual traps need to be enabled explicitly after the corresponding trap group is enabled.

Because the pre- Fabric OS v4.4 firmware only has trap group level settings, when you upgrade to the Fabric OS v4.4 firmware or above, individual traps are turned off by default even if the corresponding trap group was enabled before upgrading. When moving from a downlevel version to Fabric OS v4.4 or above you must use either **snmpmibcapset** or **snmpconfig** command to turn on explicitly the individual traps within each trap group.

Add the following note to page 1-1 under System Message Log (RASlog) section:

> **Note**:
> When the fabric is formatted in PID format 2, the error messages do not reflect the change. Port numbers in all error messages reflect the PID mode 1 port-numbering scheme.

Add the following note to page 1-1 under System Message Log (RASlog) section:

> **Note**:
> When the fabric is formatted in PID format 2, the error messages do not reflect the change. Port numbers in all error messages reflect the PID mode 1 port-numbering scheme.

On Page 1-7, Heading "Before Loading Mibs" replace the v4.2.0, v4.4.0, and v5.0.1 entries Table 1-1 with the following;

Fabric OS v4.2.0 and previous	Yes	No 2	No
Fabric OS v4.4.0	Yes	No 2	Yes 3
Fabric OS v5.0.1	Yes	No 2	Yes 3

Note:
1. The corresponding Fabric OS has SNMPv2 capabilities, but it is not officially supported by Brocade.
2. The Structure of Management Information version 2 (SMIv2) framework is used in defining the MIBs.
3. Fabric OS v4.4.0 and v5.0.1 support SNMPv3-USM (snmpUsmMIB) MIB, which is available as RFC 3414.

Add the following descriptions on page 3-6 in Table 3-2:

Display string	Represents textual information taken from the NVT ASCII character set, as defined in pages 4, 10-11 of RFC 854.
Milliseconds	Represents time unit value in milliseconds.
Microseconds	Represents time unit value in microseconds.

Add the following descriptions on page 3-32 in Table 3-3:

Display string	Represents textual information taken from the NVT ASCII character set, as defined in pages 4, 10-11 of RFC 854.
Milliseconds	Represents time unit value in milliseconds.
Microseconds	Represents time unit value in microseconds.
FcphVersion	Represents the version of FC-PH supported by an NxPort or FxPort.

Add the following descriptions on page 3-34 in Table 3-3:

FcFeModuleCapacity	Represents the maximum number of modules within a Fabric Element.
FcFeFxPortCapacity	Represents the maximum number of FxPorts within a module.
FcFeModuleIndex	Represents the module index within a conceptual table.
FcFeFxPortIndex	Represents the FxPort index within a conceptual table.
FcFeNxPortIndex	Represents the NxPort index within a conceptual table.

Add the following note on Page 5-13 to the end of the "SW Traps" section:

NOTE: The swGroupName, swGroupType, and swGroupMemPos variables are optional trap variables in Fabric OS v2.6.x. These variables are not supported in Fabric OS v4.x and above.

On Page 5-6 in the "swFabricWatchTrap" section the following variable is missing from the swFabricWatchTrap list of variables:

"swFwLastSeverityLevel 1.3.6.1.4.1.1588.2.1.1.1.10.3.1.12"

On page Page 8-51 in the "Unsupported Traps" section these traps are supported and the heading should read "FibreAlliance MIB Traps". Only the connUnitDeletedTrap is not supported by Brocade.

Add the following descriptions on page 8-8 in Table 8-1:

FcNameId	The Port Name for this entry in the SNS table.
FcGlobalId	An optional global-scope identifier for this connectivity unit. It MUST be a WWN for this connectivity unit or 16 octets of value zero.
FcAddressId	The Port Identifier for this entry in the SNS table.

Fabric OS System Error Message Reference Manual
(Publication number 53-0000515-10)

On Page 3-1 the BL-1000 message states, "This message occurs on only the SilkWorm 12000." However the message may also occur on a Silkworm 3250, 3850, or 3900 switch.

On Page 3-2 the BL--1003 message states, "This message occurs on only the SilkWorm 12000, 24000, and 48000." However the message may also occur on a Silkworm 3250, 3850, or 3900 switch.

The following messages were added after the document publication.

CDR-1001

Message

```
<timestamp>, [CDR-1001], <sequence-number>,, WARNING, <system-name>, Port <port
number> port fault. Please change the SFP or check cable
```

Probable Cause

Indicates a deteriorated SFP, an incompatible SFP pair, or a faulty cable between peer ports.

Recommended Action

Verify that the SFPs have not deteriorated and that the Fibre Channel cable is not faulty. Replace the SFPs or cable if necessary.

Severity

WARNING

CONF-1000

Message

```
<timestamp>, [CONF-1000], <sequence-number>,, WARNING, <system-name>,
configDownload completed successfully but Zoning and Security parts were ignored
```

Probable Cause

Indicates that the user with switchadmin role does not have permission to update Zoning and Security configurations, so Zoning and Security configurations are skipped while downloading the configuration file

Recommended Action

No action is required. Run the **configShow** command to view the configuration.

Severity

WARNING

FICU-1010

Message

```
<timestamp>, [FICU-1010], <sequence-number>,, WARNING, <system-name>, FMS Mode
enable failed due to address conflict with port <port number>.
```

Probable Cause

Indicates that the FICON Management Server mode (fmsmode) was not enabled because the specified port has an address conflict with the CUP management port.

Recommended Action

Use the **portDisable** command to disable the specified port causing the port address conflict.

Severity

WARNING

HAMK-1004

Message

```
<timestamp>, [HAMK-1004], <sequence-number>,, INFO, <system-name>, Resetting
standby CP (double reset may occur).
```

Probable Cause

Indicates that the standby CP is being reset due to a loss of heartbeat. This message is typically seen when the standby CP has been rebooted. Note that in certain circumstances a CP may experience a double reset and reboot twice in a row. A CP can recover automatically even if it has rebooted twice.

Recommended Action

No action is required.

Severity

INFO

PLAT-1001

Message

```
<timestamp>, [PLAT-1001], <sequence-number>,, INFO, <system-name>, Resetting
standby CP (double reset may occur).
```

Probable Cause

Indicates that the standby CP is being reset. This message is typically generated by a CP that is in the process of becoming the active CP. Note that in certain circumstances a CP may experience a double reset and reboot twice in a row. A CP can recover automatically even if it has rebooted twice.

Recommended Action

No action is required.

Severity

INFO

Fabric Watch User's Guide
(Publication number 53-0000524-06)

The following row replaces the existing rows "Invalid CRC Count," "Link Failure Count," and "State Changes" in Table A-6, "Port Class Threshold Defaults," on page A-6:

Area	Description	Default Threshold Settings	Default Alarm Settings	Threshold State
Link Failure Count	Monitors the number of link failures	Unit: Error(s) Time Base: minute Low: 0 High: 5 Buffer: 0	Changed: 0 Below: 0 Above: 0 In-Between: 0	Informative Informative Out_of_range In_range

Alarm behavior of CF monitor

Fabric watch alarm behavior depends on the state of Above, Below and In-between thresholds configured for a specific monitor. In case of CF monitor, the monitor state configurations are as shown below:

- Below threshold monitor state is STATE_IN_RANGE

- In-Between threshold monitor state is STATE_INFO

- Above threshold monitor state is STATE_OUT_OF_RANGE

Alarms are only generated when a monitor state changes from good to bad or from bad to good. In case of CF monitor, the alarm behavior is described below:

1. If the state changes from Below (STATE_IN_RANGE – good state) to In-between (STATE_INFO – good state), no alarm will be generated. If the CF usage is within a reasonable range, there is no need to raise an alarm.

2. An alarm is generated when CF monitor state goes from In-between (STATE_INFO – good state) to Above (STATE_OUT_OF_RANGE – bad state). This is to inform you that the CF usage is above the high boundary.

3. An alarm is generated when CF monitor state changes from Above to In-between. This informs the user that CF usage has returned to normal.

Secure Fabric OS Administrator's Guide
(Publication number 53-0000526-05)

On page 3-4, in the section "SilkWorm 48000 Director," delete the first paragraph and replace it with the following:

Figure 3-3 shows an example of the Web Tools Switch Explorer for a SilkWorm 48000 director.

On page 2-24, in the second example box, replace the following information:

"Address-file: -a
  Path/file-name of optional input file containing IP addresses or aliases of fabrics to which sessions should be established. If this argument is not provided, this data is read from the file indicated by environment variable 'FABRIC_CONFIG_FILE'."

With this information:

"Address-file: -a addr-file
  "addr-file" is the path/file-name of optional input file containing IP addresses or aliases of fabrics to which sessions should be established. If this argument is not provided, this data is read from the file indicated by environment variable 'FABRIC_CONFIG_FILE' if defined. Use Microsoft Notepad to create the addr-file."

SilkWorm 200E Hardware Reference Manual
(Publication number 53-0000633-01)

On page v, in the "How This Document Is Organized" section, a glossary is listed; however there is no glossary in this manual.

On page 2-7, "Configuring the SilkWorm 200E

On page 1-4 under the heading "Supported Fabric Configurations" the text should read:

The SilkWorm 200E is supported as an edge device in fabrics of up to 53 domains.

The order of tasks is incorrect, that is, the steps are the same but you must perform them in a slightly different order.

To configure the SilkWorm 200E, you must first:

1. Power on the switch.

2. Establish a physical serial connection to the switch.

3. Log in to the switch as the admin user using a hyperterminal application.
   Step number 1 in the document details how to set up the hyperterminal connection.

On page 2-6, Table 2-3, replace the BTU Rating and Input Electrical Power value:

Delete the following: 266 BTU

Replace with the following: "38 W/ 130 BTU"

On page 2-6, Table 2-3, above the row title "Input Voltage" add a row "Input Electrical Power":

Input Electrical Power   / 45VA

On page 2-6, Table 2-3, replace the BTU Rating and Input Electrical Power value:

Delete the following text: 266 BTU

Replace with the following text: 38 W/ 130 BTU
On page 2-8, "Set the IP Address," the note should read as follows:

Note: Any time the Ethernet or serial connection is not in use, the safety plug should be installed to protect it from dust or other foreign material.

On page 2-9, "Modify the Domain ID (Optional)" section a) delete the last sentence in the second paragraph and b) replace the fourth paragraph with the following :

The domain ID is a number assigned to the switch by the Fabric OS and is used when routing frames to the switch. If you do not set the Domain ID for the switch and it is attached to a fabric, the Principal switch of the fabric will assign it a new domain ID. If you set the domain ID for the switch, using the **configure** command, then this number must be unique to the fabric the switch is connecting to, or the switch will segment.

On page 4-1, "Management Features of the SilkWorm 200E," add the following note before the table:

Note: Some of the management tools listed below are available only with the appropriate license key installed.

SilkWorm 3250/3850 Hardware Reference Manual
(Publication number 53-0000623-02)

On page 2-3, replace the "Note" text:

"The 0° - 40° Celsius range applies to the ambient air temperature at the air intake vents on the nonport side of the switch. The temperature inside the switch can be up to 75° Celsius during switch operation.

If the internal temperature range exceeds the operating ranges of the components, the LEDs, error messages, and Fabric Watch alerts will indicate a problem. Enter the **tempShow** or Fabric Watch commands to view temperature status."

With this text:

"The 0° - 40° Celsius range applies to the ambient air temperature at the air intake vents on the nonport side of the switch. The temperature inside the switch can be up to 65° Celsius during switch operation.

If the internal temperature range exceeds the operating ranges of the components, the LEDs, error messages, and Fabric Watch alerts will indicate a problem. Enter the **tempShow** or Fabric Watch commands to view temperature status.

If the internal temperature range exceeds the safe range, the SilkWorm 3250/3850 reboots. To remove power from the SilkWorm 3250/3850, refer to "Powering the SilkWorm 3250/3850 On and Off" on page 3-1."

On page 3-1, add the following note:

NOTE: The switch and port status LEDs always indicate status based on the preset default thresholds and policies designed into the switch, regardless of Fabric Watch settings. Fabric Watch enables the user to receive alarms when the user defined thresholds are exceeded, as well as show them the current status of the switch.

SilkWorm 4020 Hardware Reference Manual
(Publication number 53-000688-01

Add the following text at the end of the first paragraph under the heading "Interpreting SilkWorm 4020 LEDs" on page 4-2:

Note: The blinking rate for RX/TX LED and the internal LED is slower than the physical switch when I/O activity is present.

SilkWorm 4100 Hardware Reference Manual
(Publication number 53-0000563-01)

On page 1-1, under the heading "Ports on Demand", replace this text:

"The SilkWorm 4100 has 32 ports. By default, ports 0-15 are enabled. To enable additional ports, you must install Ports On Demand (POD) licenses. To enable ports 16 through 23, you must install the POD1 license. To enable ports 24 through 31, you must install the POD2 license. Although you can install the POD2 license without having the POD1 license installed, you cannot use ports 16 through 23 until the POD1 license is

enabled. For detailed information on enabling additional ports using the Ports on Demand license, refer to the *Fabric OS Administrator's Guide*."

With this text:

"The SilkWorm 4100 model can be purchased with 16, 24, or 32 licensed ports. As your needs increase, you can activate unlicensed ports (up to the maximum of 32 ports) by purchasing and installing the Brocade Ports on Demand optional licensed product.

By default, ports 0 through 15 are activated on the SilkWorm 4100. Each Ports on Demand license activates the next group of eight ports, in numerical order. Before installing a license key, you must insert transceivers in the ports to be activated. Remember to insert the transceivers in the lowest group of inactive port numbers first. For example, if only 16 ports are currently active and you are installing one Ports on Demand license key, make sure to insert the transceivers in ports 16 through 23. If you later install a second license key, insert the transceivers in ports 24 through 31.

After you install a license key, you must enable the ports to complete their activation. You can do so without disrupting switch operation by using the **portEnable** command on each port. Alternatively, you can disable and reenable the switch to activate ports.

For more information on activating ports on demand, refer to the *Fabric OS Administrator's Guide*."

On page A-6, under the heading "Fibre Channel Port Specifications" (on page A-6), replace this text:

"The ports are capable of operating at 1, 2, or 4 Gbit/sec and are able to autonegotiate to the higher of 1 or 2 Gbit/sec. Operation at 4 Gbit/sec must be manually set"

With this text:

"The ports are capable of operating at 1, 2, or 4 Gbit/sec and are able to autonegotiate to the higher of 1, 2, or 4 Gbit/sec."

SilkWorm 12000 Hardware Reference Manual
(Publication number 53-0000148-05)

As long as one power supply is operating, all the card slots (1-10) have power. The statement should read:

Power to the backplane is load sharing and redundant across all power supplies. The left and right power feeds control INPUT power to power supplies 1 and 3, and 2 and 4 respectively. Because 2 power supplies are required to support a fully populated 12000, we recommend filling all 4 power supply slots to ensure that in the case of a power feed failure, the chassis will have enough power for both switches.

On page 2-2, under the heading, "Powering the SilkWorm 12000 On and Off," replace the following information:

To power the SilkWorm 12000 off:

Flip both AC power switches to "0". To remove all sources of power from the switch, disconnect both cables from the power source.

Note: Removing all power from the switch triggers a system reset. When power is restored, all devices are returned to the initial state and the switch runs POST.

With this information:

To power the SilkWorm 12000 off:

1. Shut down both logical switches (see Figure 2-1):

    a. Enter the **switchShutdown** command to ensure a graceful shutdown of Switch 1, and verify the command has completed and displayed the message "Cleaning up kernel modules.....Done".

    b. From the active CP card session, log into Switch 0 by entering the login command, logging in as admin, then entering "0" to log into Switch 0.

    c. Enter the **switchShutdown** command to ensure a graceful shutdown of Switch 0, and verify the command has completed and displayed the message "Cleaning up kernel modules.....Done".

Figure 2-1 Sample Output for the **switchShutdown** Command on Both Switches

```
SW1:admin> switchshutdown
Stopping all switch daemons...Done.
Powering off slot 7...Done.
Powering off slot 10...Done.
Checking all slots are powered off...Done.
Cleaning up kernel modules.....Done
SW1:admin>
SW1:admin> login
login: admin
Enter Switch Number to Login <0 or 1>: 0
password: xxxx
SW0:admin>
SW0:admin> switchshutdown
Stopping all switch daemons...Done.
Powering off slot 1...Done.
Powering off slot 4...Done.
Checking all slots are powered off...Done.
Cleaning up kernel modules.....Done
SW0:admin>
```

For details on the **switchShutdown** command, refer to the Fabric OS Command Reference Manual, or the online help.

2. Power off the chassis by flipping both AC power switches to "0" (LEDs inside AC power switches should turn off). See Figure 1-1 on page 1-2 for location of switches. To maintain the ground connection, leave both power cords connected to the chassis and to an electrical outlet.

Page 2-6, in Table 2-1 "Port Card LED Patterns," replace the following table information. In the line that begins "Left of each port, lower LED" next to the "No light (LED is off)" entry replace the following information:

- The port card does not have incoming power, or there is no light or signal carrier detected.
- Polling is in progress
- Connected device is configured in an offline state.

with the following:

- The port card does not have incoming power, or there is no light or signal carrier detected.
- Polling is in progress.
- Connected device is configured in an offline state.
- The port on the switch is not disabled, the state on the local port is offline, but the remote port is still emitting light.

On page 2-12, he following statement in the "Operating Information for Power Supplies" section is incorrect:

The left power connector provides power to the power supplies in power supply bays #1 and #3 (color-coded blue), which provide power to the left side of the chassis (slots 1-5). The right power connector provides power to the power supplies in power supply bays #2 and #4 (color-coded yellow), which provides power to the right side of the chassis (slots 6-10).

SilkWorm 24000 Hardware Reference Manual
(Publication number 53-0000619-01)

On page A-2, table A-1, "System Architecture," replace the following table entry:

"Switch latency <2.1 µsec any port to any port at 2 Gb/sec, cut-through routing"

With this table entry:

"Switch latency 2.05 < 2.35 µsec any port to any port at 2 Gbit/sec, cut-through routing"

Step 1 of the "Replacing a Power Supply and Filler Panel" on page 5-21 is incorrect.

Determine whether power adequate to keep the chassis operating will be available throughout the replacement. If adequate power will *not* be consistently available, shut down the SilkWorm 24000 gracefully, as follows:

a. Open a telnet session to the active CP card and log in to the switch as admin.

b. Enter the **switchshutdown** command.

c. Power off the chassis by flipping both AC power switches to the off position (the "0" on the AC switch).

Replace Step 1 with this information:

Determine whether power adequate to keep the chassis operating will be available throughout the replacement. If adequate power will *not* be consistently available, shut down the SilkWorm 24000 gracefully, as follows:

a. Open a telnet session to the active CP card and log in to the switch as root.

b. Enter the following command:
```
/usr/bin/shutdown -h now
```

c. Watch the console log for the following power down message. The director will automatically reboot, so hit the ESC key to stop at the bootprom. This will stop the standby CP from rebooting.

```
The system is going down for system halt NOW !!
INIT: Switching to runlevel: 0
INIT: Sending processes the TERM signal
2005/08/17-18:10:01, [FSSM-1003], 19,, WARNING, Silkworm12000, HA State out
of sync
Unmounting all filesystems.
The system is halted
flushing ide devices: hda
Power down.

The system is coming up, please wait...
Checking system RAM - press any key to stop test
00b00000
System RAM check terminated by keyboard
System RAM check complete
Press escape within 4 seconds to enter boot interface.

1) Start system.
2) Recover password.
3) Enter command shell.
```

d. Login to the active CP and repeat steps b and c for the active CP. Once both CPs are stopped at the boot prom, you can power off the system safely.

e. Power off the chassis by flipping both AC power switches to "0" (LEDs inside AC power switches should turn off). See Figure 1-1 on page 1-2 for location of switches. To maintain the ground connection, leave both power cords connected to the chassis and to an electrical outlet.

On Page 5-21 under the heading "Replacing the Power Supply and Filler Panel" replace Step 1 with the following:

1. Determine whether power adequate to keep the chassis operating will be available throughout the replacement. If adequate power will *not* be consistently available, shut down the SilkWorm 24000 gracefully as follows:

    a. From the active CP card, log into Switch 1 by entering the login command, logging in as admin, then entering "1".

    b. Enter the **switchShutdown** command to ensure a graceful shutdown of Switch 1, and verify the command has completed and displayed the message "Cleaning up kernel modules.....Done".

    c. From the active CP card, log into Switch 0 by entering the login command, logging in as admin, then entering "0".

    d. Enter the **switchShutdown** command to ensure a graceful shutdown of Switch 0, and verify the command has completed and displayed the message "Cleaning up kernel modules.....Done".

    e. Power off the chassis by flipping both AC power switches to the off position (the "0" on the AC switch).

On page 3-2, under the heading "Configure IP Addresses for CP Cards," remove the first sentence in the following note:

"Note: Use a block of three IP addresses that are consecutively numbered in the last octet. The IP and gateway addresses must reside on the same subnet."

Table 4-7 on page 4-15 within the "WWN Card" section in Chapter 4 needs to be revised. Replace Table 4-7 with the following:

Table 4-7 WWN Bezel LED Patterns

LED Location/Purpose	Color	Status	Recommend Action
16-Port card/CP card Power	Steady green	Power is OK.	No action required.
	Flashing green	Power to port card is OK; however, this LED flashes if the port card status LED is flashing.	Check port card status LED and determine if it is flashing slow (2 second increments) or fast (1/2 second increments) and then take appropriate action.
	No light (LED is OFF)	No port card present or power source is unavailable.	Insert port card, or check AC switch or power source.
	NOTE: Check the individual port card (see Figure 4-1 on page 4-2) or CP card power LEDs (see Figure 4-2 on page 4-6) on the port side of the chassis to confirm the LED patterns.		
16-Port card/CP card Status	Steady amber	Port card is faulty.	Check port card.
	Slow-flashing amber (on 2 seconds; then off 2 seconds)	Port card is not seated correctly or is faulty.	Pull card out and reseat it. If LED continues to flash, replace card.
	Fast-flashing amber (on 1/2 second; then off 1/2 second)	Environmental range exceeded or port card failed diagnostics (run during POST or manually).	Check for out-of-bounds environmental range and correct it. Replace card if it fails diagnostics.
	No light (LED is OFF)	Port card is either healthy or does not have power.	Verify that the port card power LED is on.

LED Location/Purpose	Color	Status	Recommend Action
	NOTE: Check the individual port card (see Figure 4-1 on page 4-2) or CP card status LEDs (see Figure 4-2 on page 4-6) on the port side of the chassis to confirm the LED patterns.		
Power supply/ Power/Status	Steady green	Power is OK.	No action required.
	Steady amber	Power supply is faulty.	Ensure that the correct AC power switch is on and the power supply is seated. If LED remains on, replace the power supply.
	Slow-flashing amber	FRU header (SEEPROM cannot be read) due to I2C problem.	Replace power supply.
	Fast-flashing amber	Power supply is about to fail due to failing fan inside the power supply.	Replace power supply.
	No light (LED is OFF)	No power supply present or is not inserted/seated properly, or power source is unavailable.	Insert power supply module, ensure it is seated properly, or check AC switch or power source.
	NOTE: Check the individual power supply LEDs on the port side of the chassis to confirm the LED patterns (see Figure 4-3 on page 4-9).		

NOTE: If a port card slot or power supply bay has a filler panel installed, the corresponding LEDs on the WWN card do not light up.

On page 5-20 , "Replacing a Power Supply and Filler Panel, add the following paragraph:

"A Silkworm 24000 that is fully populated with FC2-16 blades can function on one power supply. Redundancy of the power supply is achieved using power supply FRUs in slots 1 and 2. You can populate all 4 power supply slots in the SilkWorm 24000 for maximum redundancy. Power supply FRUs are interchangeable between Silkworm 12000 and Silkworm 24000."

SilkWorm 48000 Hardware Reference Manual
(Publication number 53-0000645-01)

On page A-8, Table A-6, replace the Heat dissipation values as follows.

Delete the following:

913 Watts or 3115 BTU (Eight
FC4-32 blades and two CP4 blades)
711 Watts or 2426 BTU (Eight
FC4-16 blades and two CP4 blades)

Replace with the following:

720 Watts or 2457 BTU (Eight
FC4-32 blades and two CP4 blades)

On page 6-7 under the heading "Disconnect from Network and Fabric" replace step 1 with the following:

1. Shut down the SilkWorm 48000.

a. Verify which CP is the active CP, and login to the active CP using a Serial Console connection.

b. On the standby CP, set the slider switch to the off position, or eject the standby CP from the chassis. This disables the standby CP.

c. Run the reboot command from the active CP. This will gracefully take down the system.

d. When you see the "Press escape within 4 seconds to enter boot interface." message, press ESC to suspend the switch.

Add the following section to the end of Appendix B Diagnostics and Troubleshooting:

Powering Off the SilkWorm 48000

To power off the SilkWorm 48000 gracefully:

1. Verify which CP is the active CP, and log in to the active CP using a Serial Console connection.

2. On the standby CP, set the slider switch to the off position, or eject the standby CP from the chassis. This disables the standby CP.

3. Run the reboot command from the active CP to take down the system gracefully.

4. When you see the "Press escape within 4 seconds to enter boot interface." message, press ESC to suspend the switch.

5. Power off the chassis by flipping both AC power switches to "0" (LEDs inside AC power switches should turn off). To maintain the ground connection, leave both power cords connected to the chassis and to an electrical outlet.

SilkWorm Director Blade Support Notes
(Publication number 53-0000761-01)

On page 11, in the section "Adding FC2-16 Blades to a SilkWorm 48000," replace the second paragraph:

If you are using **chassisConfig** mode 1 you can add **FC2-16 cards** with minimal disruption. If you are using **chassisConfig** mode 5, you MUST change to mode 1 prior to executing this procedure. Changing the **chassisConfig** mode requires a reboot and is disruptive.

Web Tools Administrator's Guide
(Publication number 53-0000522-08)

On page 3-4, in the section "SilkWorm 48000 Director," delete the first paragraph and replace it with the following:

Figure 3-3 shows an example of the Web Tools Switch Explorer for a SilkWorm 48000 director.

On page 3-7, in the section "Refresh Rates," add the following paragraph after the first paragraph:

The refresh, or polling, rates listed in this section and throughout the book indicate the time between the end of one polling and the start of the next, and *not* how often the screen is refreshed. That is, a refresh rate of 15 seconds does not mean that a refresh occurs every 15 seconds. It means that a new refresh starts 15 seconds after the previous refresh finished.

On page 3-7, in the section "Fabric Tree," delete the fourth paragraph:

The Fabric Tree is updated at time intervals depending on the number of switches in the fabric. On average, for a fabric with up to 12 switches, the Fabric Tree is updated every 30 seconds. For every additional 12 switches in the fabric, it takes an additional 30 seconds to update the Fabric Tree. The Switch Information View displays a field, "Polled At", that identifies the last time the information was updated.

On page 3-7, in the section "Fabric Tree," replace the last paragraph:

You can also manually refresh the status of a switch within the fabric by right-clicking that switch in the Fabric Tree and clicking **Refresh**.

With this paragraph:

You can manually refresh the status of a switch within the fabric by right-clicking that switch in the Fabric Tree and clicking **Refresh**.

On page 4-27, in the section "Displaying the Name Server Entries," replace the following text in the Note:

You must click **Refresh** from the Name Server window to poll Name Server entries.

You can also specify a time interval at which the Name Server entries will be automatically refreshed.

With this text:

Click **Refresh** in the Name Server window to poll Name Server entries.

You can also click the Auto Refresh checkbox and specify a time interval at which the Name Server entries will be automatically refreshed.

On page 4-27, in the section "To view a list of the switches in the Name Server," replace the following steps:

2. *Optional*: Check the **Auto Refresh** checkbox on the Name Server window.

3. *Optional*: Enter an autorefresh interval (in seconds), at a minimum of 15 seconds. The Name Server entries will refresh at the rate you set.

With this step:

2. *Optional*: Check the **Auto Refresh** checkbox on the Name Server window. Type an auto-refresh interval (in seconds); the minimum (and default) interval is 15 seconds. The Name Server entries will refresh at the rate you set.

In Chapter 12, "Administering FICON CUP Fabrics," on page 12-1, in the section "This chapter contains" add the following bullet:

*   "Enabling Port Based Routing on the SilkWorm 4100 and SilkWorm 48000," next

In Chapter 12, "Administering FICON CUP Fabrics," on page 12-1, add a new section "Enabling Port Based Routing"

*   Enabling Port Based Routing on the SilkWorm 4100 and SilkWorm 48000

Port-based path selection is a routing policy in which paths are chosen based on ingress port and destination only. This also includes user-configured paths. All SilkWorm 4100 and 48000 switches with FICON devices attached must have port-based routing policy enabled. Port-based routing is a per-switch routing policy. After port-based routing is enabled, you can continue with the rest of the FICON implementation.

To enable or disable port-based routing

1. Click a switch with FICON devices attached from the Fabric Tree.

2. Launch the Switch Admin module as described on page 4-3.

3. Click the FICON CUP tab.

The FICON CUP tab displays, with the FICON Management Server subtab in front, as shown in Figure 12-1.

4. Click the **Enable** radio button to enable the port-based routing policy.

Figure 12-1 FICON CUP Management, Port Based Routing.

On page 14-1, in the section "Monitoring Performance Using Web Tools," replace the following paragraph:

> Each graph is displayed individually in a window, so it can be minimized, maximized, resized, and closed. Graphs within the Performance Monitor module are updated every 30 seconds.

With these paragraphs:

> Each graph is displayed individually in a window, so it can be minimized, maximized, resized, and closed.

> Graphs within the Performance Monitor module are updated every 30 seconds. When you first display the graph or if you modify the graph (such as to add additional ports), you might have to wait up to 30 seconds before the new values are shown.

On page 14-3, add this paragraph after the first paragraph:

> Port throughput utilization is represented by a horizontal bar for each selected port, which gets longer or shorter depending on the percent utilization for that port at the last poll time. Thin short vertical intersecting bars give a historical perspective by representing the highest and lowest values reached for each selected port since the graph was opened. A third bar between them represents the average of all values polled. (See Figure 14-1.)

On page 14-3, replace Figure 14-1 with this one:

Figure 14-1   Accessing Performance Graphs
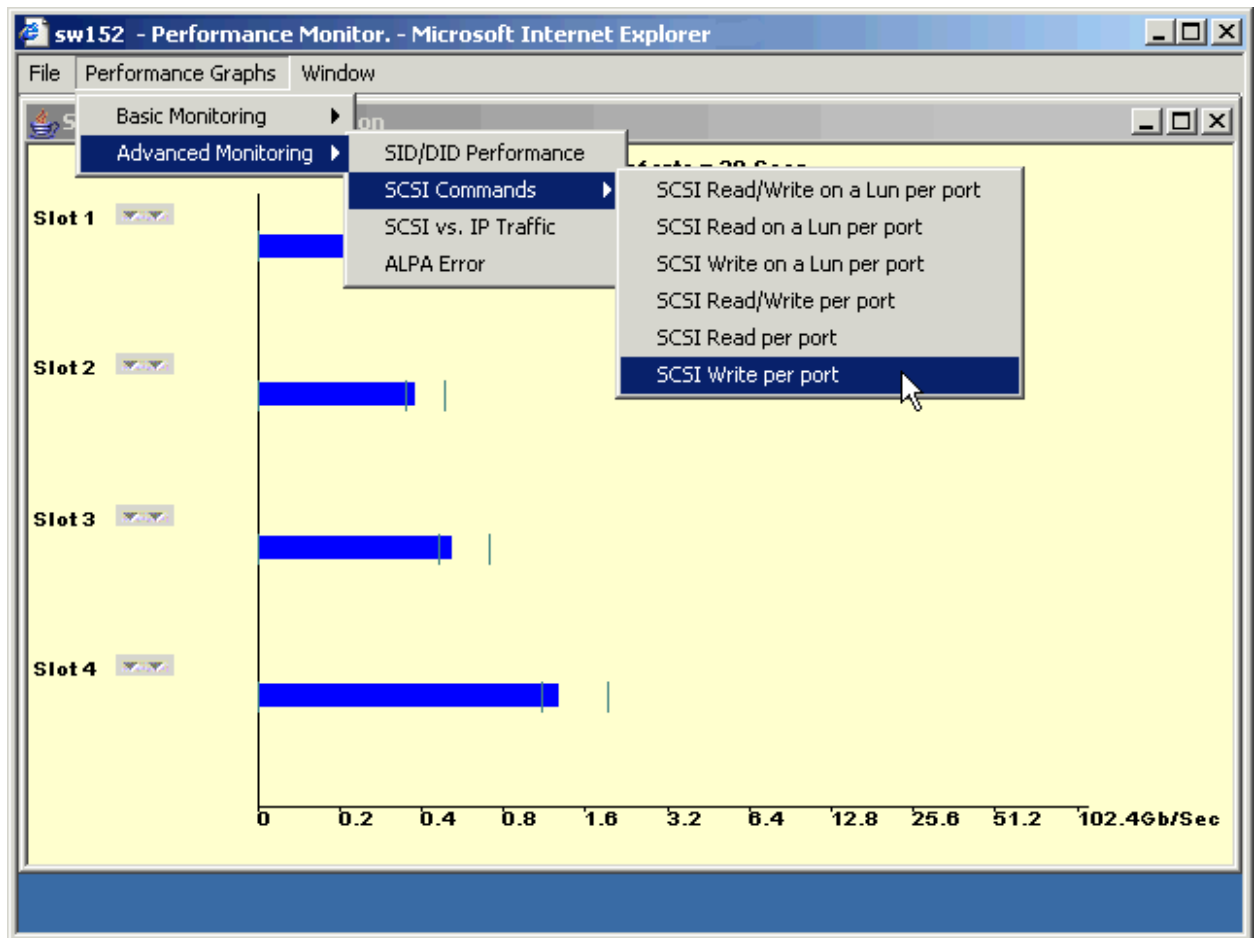
Closed Defects in Fabric OS v5.0.5b

This table lists defects newly closed in this release of Fabric OS.

Defect ID	Severity	Description
DEFECT000058918	High	Summary: HBA occasionally may not log into a 4Gb switch during auto topology discovery. Symptom: The port connecting a certain 3rd party HBA (1, 2, & 4Gb) and a 4Gb Brocade switch (SilkWorm 4100 or 48000) may not come up during initialization. Solution: Sends a CLS/LIP/OLS followed by a loop init. Workaround: Use portCfgGPort to lock the switch in G port mode. To re-establish the link on the affected port, issue portDisable and portEnable. Customer Impact: This will happen as a result of maintenance activity or a severe error condition that would cause the link to force a full reset to recover. A fabric may see a few ports not come back on line - until a portDisable / portEnable. Probability: Medium Risk of Fix: Medium Service Request# RQST00000039264 Reported in Release: V5.0.1
DEFECT000064330	High	Summary: Switch panics with Failure Detection (FDET) when a Linux host is connected, affecting the SilkWorm 3900, 24000, and 3x50 platforms. Symptom: Error message: "[BLL-1000], 2234,, CRITICAL,, ASIC driver detected Slot x port x as faulty (reason: 13)." It might also have a reason code of 9. There are other types of FDET that can report the same message. To confirm, related data should be gathered with supportsave and interpreted. Solution: Handles a corrupted hardware register and also logs a new internal RASLOG entry to identify the condition. Probability: Low Risk of Fix: Medium Service Request# RQST00000043717 Reported in Release: V5.0.1

Defect ID	Severity	Description
DEFECT000068280	High	Summary: A phantom device on an L-port disappears after HAfailover/hareboot, and a private device is no longer in the name server database, affecting the SilkWorm 3900, 12000, 24000, and 3x50 platforms. Symptom: HAFailover can cause the newly active CP to register the port without the phantom device, and without the private device in the name server. Host and device may not see each other. Solution: Handles the removed public device in the LIPA (loop initialization primitives assigned) phase and leaves the cleanup of the removed private device to the LIHA (loop initialization hardware assigned) phase. Workaround: Perform the following commands on the port: portdisable and portenable. Probability: Medium Risk of Fix: Low Service Request# RQST00000047162 Reported in Release: V5.0.1
DEFECT000069556	High	Summary: Specific HBA behavior causes a memory leak and the switch runs out of memory and reboots. Symptom: The pdshow command displays: "Out of Memory: Killed process 638 (nsd0). VM size = 18380 KB, Runtime = 9844 minutes, CPU time = 423 sec." and RASLOG entry similar to: "2006/05/06-03:01:22, [KSWD-1003], 195,, WARNING, Silkworm3900, kSWD: Detected unexpected termination of . . . " Solution: Fixed code to free IU when a specific 3rd party HBA performs a PLOGI to a target not yet registered with the switch name server and the PID in the PLOGI frame has a non-zero Arbitrated Loop Physical Access (ALPA) for the unknown target. Probability: High Risk of Fix: Low Service Request# RQST00000048426 Reported in Release: V5.0.1

Defect ID	Severity	Description
DEFECT000071137	High	Summary: Switch panics with time service daemon (tsd) core dump. Symptom: Error message on console: "2006/06/11-01:42:34, [KSWD-1003], 62/49,, WARNING,, kSWD: '[18]tsd:0'R failed to refresh SWD***." This affects all Fabric OS v4.x and v5.x based platforms. Solution: Time service daemon (tsd) is changed to update the hardware clock asynchronously. Probability: Medium Risk of Fix: Low Service Request# RQST00000049268 Reported in Release: V5.0.3
DEFECT000074725	High	Summary: Zoning change fails with message: unabortable internal zoning transaction. Symptom: Error message: "trans_begin: there is an outstanding transaction, and you are not the owner of that transaction." But cfgtransshow reports "There are no outstanding zoning transactions." The GUI interface fails with message: "Failed to commit changes to fabric; There is an unabortable internal zoning transaction (zone merge or zoning Enable/Disable issued remotely), and you are not the owner of that critical transaction. Please resubmit your zoning changes after that transaction is complete." This impacts Fabric OS 4.x/5.x platforms. Solution: Removed a condition in which zoning is blocked by the underlying transport layer. Risk of Fix: Low Service Request# RQST00000049717 Reported in Release: V5.0.3

Defect ID	Severity	Description
DEFECT000076440	High	Summary: The Web Tools GUI hangs upon invoking fabric wide polling when there is no IP connection between the Web Tools proxy switch and the other switches in the fabric. Symptom: On a large fabric (e.g., 30+ embedded switches), the GUI hangs upon launch and no access via Web Tools is possible. This impacts all platforms if the proxy has no IP connection to a large number of switches in the fabric. It may be more apparent after an upgrade from Fabric OS v4.x to v5.0.x if there are switches in the fabric previously unrecognized by Fabric OS v4.x, and then after the upgrade is newly recognized by Fabric OS 5.0.x. These switches are on a different subnet from the proxy and without an IP connection to the proxy. Solution: Provides two configurable parameters to allow the user to selectively disable fabric wide polling if there is not IP connectivity between all switches: Perform License Checking and Warning (yes, y, no, n): [yes] n Allow Fabric Event Collection (yes, y, no, n): [yes] n The first configuration parameter can be used to disable license polling to other switches in the fabric. Note: these parameters can be updated with the "configure" command while the switch is on line. The Web Tools application must be relaunched in order for the configuration to take effect. Risk of Fix: Low Service Request# RQST00000049863 Reported in Release: V5.0.3
DEFECT000076551	High	Summary: Without displaying an error message, a hardware ASIC failure causes performance degradation on a SilkWorm 48000, or a switch fault on other 4Gb switches. Symptom: Porterrshow indicates no errors on F_ports, though the ISLs on the remote switch show a large number of frame errors. Performance drops dramatically without traffic failover to the redundant data path. Solution: Added external critical RASLOG CDR-1003 message to identify the failed slot and ASIC chip. Workaround: Replace the faulted blade. Risk of Fix: Low Service Request# RQST00000052235 Reported in Release: V5.0.3

Defect ID	Severity	Description
DEFECT000057407	Medium	Summary: During stress testing, after fastboot on a fully-loaded SilkWorm 48000, a few E_ports and F_ports may stay in "Mod_Val" mode. Symptom: During stress testing, the SilkWorm 48000 is loaded with 32 ISLs and 320+ devices. A few E_ports and F_ports may stay in "Mod_Val" mode for more than 11 minutes. This only impacts the Silkworm 48000. Solution: Added code to retry i2c access to the SFP for the SilkWorm 48000. Workaround: Disable and enable the port that is stuck in the Mod-Val state. Probability: Medium Risk of Fix: Low Reported in Release: V5.0.1
DEFECT000062690	Medium	Summary: During 3rd party SFP qualification testing with certain types of SFP, a high number of cyclic redundancy check (CRC) and low level interrupt (LLI) errors were reported by porterrshow. Symptom: The errors reported were observed only on 4Gb switches. Solution: Sets the proper bit rate in the hardware register after the two end points have successfully negotiated the port speed. Risk of Fix: Low Reported in Release: V5.1.0
DEFECT000068123	Medium	Summary: On the SilkWorm 12000, logging into both logical IP addresses puts the user on the same logical switch. Symptom: Sometimes user cannot get on the intended logical switch to make changes on the SilkWorm 12000. This impacts the SilkWorm 12000 only. Solution: The fix handles a race condition that sometimes does not update a variable properly. Workaround: Use ssh (secure shell) or after logging in as root, change the logical switch instance: Enter the setswitch <instance> command where <instance> is either 0 or 1. Change 0 to 1 if needed. The PS1 line needs to be entered only once per login. Risk of Fix: Low Service Request# RQST00000047128 Reported in Release: V4.4.0

Defect ID	Severity	Description
DEFECT000069096	Medium	Summary: Port blade on a SilkWorm 24000 faulted due to Failure Detection when there is a 1Gb F_port on an even-numbered user port. Symptom: Error message: "2006/05/09-15:01:37, [BLM-5010], 212797/0,, CRITICAL, S3,P39(12): fdet(ERR_X)Slot:3, Port:12, . . . ERROR, Slot 3 set to faulty, rc=20015." Solution: Prevents hardware triggered failure detection (FDET) by having multiple ports set in a unicast route entry. This rare problem is isolated to the bloom ASIC and the specific configuration of a 1Gb F_port connection on an even numbered user port. Note that after the upgrade, the ports need to be bounced for the fix to take effect. Workaround: Connect 1Gb F_ports to odd-numbered user ports only. Probability: Low Risk of Fix: Low Service Request# RQST00000048318 Reported in Release: V5.0.4
DEFECT000076931	Medium	Summary: Inserting unsupported FC4-48 blade into a pre-Fabric OS v5.2 chassis causes FAULTY(1) error that will not clear with blade removal. HAsync is dropped, and an active CP reboot is required to recover. Symptom: HA sync will be lost and removal of the blade will not recover the HA sync. Solution: Detects and faults only the unknown blade, allowing HA to stay in sync. Workaround: Schedule window to perform reboot of the active CP. Probability: Medium Risk of Fix: Low Reported in Release: V5.0.5

Closed Defects in Fabric OS v5.0.5a

This table lists defects newly closed in this release of Fabric OS.

Defect ID	Severity	Description
DEFECT000068292	High	Summary: Host cannot see a specific device on 2Gig platforms: SilkWorm 3900, 12000, 24000. Symptom: Device sent an F_Port login (FLOGI) and immediately performed a link initialization. The FLOGI continued to be processed even though the link was initializing. This impacted the route. There is no error message. Solution: Under this condition, the incoming FLOGI is dropped. Workaround: Remove and reinsert the cable, or run portDisable/portEnable. Risk of Fix: Low Service Request# RQST00000047840 Reported in Release: v5.0.1
DEFECT000071142	High	Summary: In rare cases, the zoning daemon (zoned) asserts and causes the switch to panic during HAfailover. Symptom: If the switch panics during HAfailover, the console log displays "do_assert (expression=0x10062ae8 "ns_zt[port].ns_zonetype != 0", file_name=0x0, line_num=970) at dbg_assert.c:268." Solution: Handle the error condition properly during the zone type recovery. Risk of Fix: Low Service Request# RQST00000049244 Reported in Release: v5.0.3

Defect ID	Severity	Description
DEFECT000071707	High	Summary: An abrupt power failure on a port blade on a SilkWorm 48000 might cause the CPs to fail. Symptom: When a port blade experiences an abrupt power glitch/failure, 1 or 2 CPs might become faulted. Not seen when the port blade is powered off manually and can only be simulated via diagnostics tools. RASLOG BL-1006 and HAM-1001 together identify this problem with the following error messages: [BL-1006], 172,, INFO, , Blade 5 NOT faulted. Peer blade 1 experienced abrupt failure and [HAM-1001], 182,, CRITICAL, , Standby CP is not Healthy, device CP Blade (Rsn: 21) status BAD, severity = CRITICAL. Solution: In the event of an abrupt power failure, the fix allows a small delay for the power failure to be reflected in the power status register. The register is then read a second time to detect the correct blade to fault. Risk of Fix: Medium Service Request# RQST00000049520 Reported in Release: v5.0.3
DEFECT000069629	Medium	Summary: Unhelpful error message after portCfgGPort for embedded platform products. Symptom: Error message does not explain why the command cannot be completed. Solution: Fixes the error message to explain that the portCfgGPort command cannot be applied to ports connected to internal server blades in the embedded products. Customer Impact: Usability issue, with no impact to functionality. Probability: Medium Risk of Fix: Low Reported in Release: v5.0.5