BROCADE

Brocade Enterprise and Entry SAN Switch Modules for IBM eServer BladeCenter Design, Deployment and Management Guide DDM

Version 1.0

Publication Number: 53-0000561-01

Publication Date: May 4, 2004

 Publication Number: 53-0000561-01

Brocade Communications Systems, Incorporated

Corporate Headquarters
1745 Technology Drive
San Jose, CA 95110
T: (408) 333-8000
F: (408) 333-8101
Email: info@brocade.com

European and Latin America Headquarters
29, route de l'Aeroport
Case Postale 105
CH-1211 Geneva 15,
Switzerland
T: +41 22 799 56 40
F: +41 22 799 56 41
Email: europe-info@brocade.com

Asia-Pacific Headquarters
Shiroyama JT Trust Tower 36th Floor
4-3-1 Toranomon, Minato-ku
Tokyo, Japan 105-6036
T: +81 35402 5300
F: +81 35402 5399
Email: apac-info@brocade.com

Document History

The following table lists all versions of the *Brocade Enterprise and Entry SAN Switch Modules for IBM eServer BladeCenter Design, Deployment and Management Guide*.

Document Title	Publication Number	Summary of Changes	Publication Date
Brocade Enterprise and Entry SAN Switch Modules for IBM eServer BladeCenter Design, Deployment and Management Guide	*53-0000561-01* Version 1.0	This is the first release of this guide.	May 2004

Contents

About This Document

This document is a user guide written for SAN administrators to help you learn about many of the Brocade® licensed products and how they relate to the Brocade® Enterprise SAN Switch Module for IBM® eServer BladeCenter™ and the Brocade® Entry SAN Switch Module for IBM® eServer BladeCenter™. This guide covers both models.

"About This Document" contains the following sections:

- How This Document Is Organized
- What's New in This Guide
- Document Conventions
- Additional Information
- Getting Technical Help
- Document Feedback

How This Document Is Organized

This document is organized into two sections to help you find the particular information that you want as quickly and easily as possible.

The first section, “Section I,” describes each of the specific Brocade SAN Switch Module licensed feature products. This section provides only concepts. If you are already familiar with the licensed products described in this guide, you might want to refer to the *Brocade Fabric OS Features Guide*, the *Brocade Fabric OS Procedures Guide* or the *Brocade Web Tools Administrator's Guide* to learn how to use the products in more detail.

The second section, “Section II” is focused on covering detailed “how to” information for the Brocade SAN Switch Module from a design, deployment, and management perspective. This portion of the Design, Deployment and Management guide is intended to be used in conjunction with existing Brocade manuals, release notes, and related Brocade publications (especially the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0*).

The document contains the following components:

- The title page and "Table of Contents" provide the version number, date, and part number of this publication; copyrighted information contained in the document; Brocade Communications, Inc., locations around the world; a document history table; and the topics covered in this particular document.
- "About This Document" provides information specific to this document: how it is organized, what information has changed since its most recent publication, the typographic conventions and particular terminology that it uses, where to go for further information on the topic, how to get technical assistance with your product, and how to provide your feedback about this document.

Section I

- Chapter 1, “Product Introduction”, introduces the Brocade Enterprise and Entry SAN Switch Modules for IBM eServer BladeCenter also known as the Brocade SilkWorm 3016.
- Chapter 2, “Switch and Fabric Management”, explains the different methods used to manage a Brocade SilkWorm switch and a Brocade SAN.
- Chapter 3, “Advanced Zoning”, describes the Advanced Zoning licensed feature.
- Chapter 4, “ISL Trunking”, describes the ISL Trunking licensed feature.
- Chapter 5, “Advanced Performance Monitoring”, describes the Advanced Performance Monitoring licensed feature.
- Chapter 6, “Extended Fabrics”, describes the Extended Fabrics licensed feature.

- Chapter 7, “Remote Switch”, describes the Remote Switch licensed feature.
- Chapter 8, “Brocade Advanced Security (Secure Fabric OS)”, describes the Brocade licensed feature for security.
- Chapter 9, “Brocade Interoperability Mode”, describes heterogeneous fabric connectivity.

Section II – SAN Design

- Chapter 10, “SAN Solutions”, introduces SAN solution concepts.
- Chapter 11, “SAN Availability”, introduces SAN Availability concepts.
- Chapter 12, “SAN Scalability”, introduces SAN Scalability concepts.
- Chapter 13, “SAN Performance”, introduces SAN Performance concepts.
- Chapter 14, “ISL Trunking”, describes Trunking Design concepts for the Brocade SAN Switch Module.
- Chapter 15, “Architecting SANs with SilkWorm Switches”, describes device attachment strategies and platform specific features.
- Chapter 16, “Zoning Design”, guidelines for Advanced Zoning.
- Chapter 17, “Security Design”, guidelines for Advanced Security (Secure Fabric OS).

Section II – SAN Deployment

- Chapter 18, “SAN Deployment Overview”, covers overview of deployment strategies.
- Chapter 19, “Planning”, introduces deployment planning concepts.
- Chapter 20, “Staging”, introduces deployment checklists for a Brocade SAN.
- Chapter 21, “Validation”, introduces SAN validation concepts.
- Chapter 22, “Maintenance”, introduces SAN maintenance concepts.

Section II – SAN Management

- Chapter 23, “SAN Management Overview”, introduces management concepts for a Brocade SAN.
- Chapter 24, “Brocade SAN Switch Module Management Tools”, introduces an overview of the management tools that can be employed to manage a SAN.
- Chapter 25, “SNMP”, introduces SNMP management concepts and examples.
- Chapter 26, “Fabric Watch”, introduces Fabric Watch management concepts and examples.
- Chapter 27, “Advanced Performance Monitoring”, introduces Advanced Performance Monitoring management concepts and examples.

What's New in This Guide

This is the first release of this guide.

The *Brocade Enterprise and Entry SAN Switch Modules for IBM eServer BladeCenter Design, Deployment and Management Guide* describes the following Brocade optionally licensed products:

- Brocade Advanced Web Tools
- Brocade Advanced Zoning
- Brocade Fabric Watch
- Brocade ISL Trunking
- Brocade Advanced Performance Monitoring
- Brocade Extended Fabrics
- Brocade Remote Switch
- Brocade Secure Fabric OS® (Advanced Security)
- Brocade Fabric Manager

The information for *Advanced Zoning*, *ISL Trunking*, *Advanced Performance Monitoring*, *Extended Fabrics* and *Remote Switch* has been consolidated into the *Brocade Fabric OS Features Guide*.

The other optional products have their own documentation and can be found in the following additional user's manuals: *Advanced Web Tools Administrator's Guide* (for procedures using Advanced Web Tools), *Fabric Watch User's Guide*, *Secure Fabric OS User's Guide, Secure Fabric OS QuickStart Guide* and the *Fabric Manager User's Guide* (for procedures using Fabric Manager). These products have separate books, because the books are too large to be easily incorporated into the *Brocade Fabric OS Features Guide*. Also, because Brocade Advanced Web Tools and Brocade Fabric Manager are alternate ways of managing a fabric, they are documented separately, as the counterpart to the *Brocade Fabric OS Procedures Guide*.

There are several other manuals that provide more detailed information on the CLI commands, SNMP, and the Diagnostics available in this Fabric Operating System. The *Brocade Fabric OS Reference Manual* covers command line descriptions, the *MIB Reference Manual* covers the SNMP agent supplied by the switch, and the *Diagnostic and System Error Messages Reference Manual* covers the switch diagnostic commands and messages.

The *Brocade Fabric OS Procedures Guide* completes the set of manuals provided and offers more in-depth descriptions of the commands that can be used and how to use them to manage Brocade Fabrics to help storage area network (SAN) administrators like you configure and manage your Brocade SilkWorm SAN.

In addition there are two hardware reference manuals that explain how to plug the Brocade SAN Switch Module into the IBM eServer BladeCenter chassis and get the unit up and running: The *SilkWorm 3016 Hardware Reference Manual* and the *SilkWorm 3016 QuickStart Guide*.

Document Conventions

This section describes text formatting conventions, important notices formats, and terms as they are used in this document.

Text Formatting

The following table describes the narrative-text formatting conventions that are used in this document.

Convention	Purpose
bold text	• Identifies command names • Identifies GUI elements • Identifies keywords/operands • Identifies text to enter at the GUI or CLI
italic text	• Provides emphasis • Identifies variables • Identifies paths and internet addresses • Identifies document titles
`code` text	• Identifies CLI output • Identifies syntax examples

Notes, Cautions, and Warnings

The following notices appear in this document.

Note	A note provides a tip, emphasizes important information, or provides a reference to related information.
Guideline	Guidelines are recommendations for consideration. The adoption of these guidelines is a function of the user's ability to interpret and correlate relevant SAN information and make decisions based upon their organization and SAN requirements.
Caution	A caution alerts you to potential damage to hardware, firmware, software, or data. Cautions indicate that a particular action or type of connection is not recommended as it may cause failure of the switch or fabric.
Warning	A warning alerts you to potential danger to personnel.

Additional Information

This section lists additional Brocade, IBM and industry-specific documentation that you might find helpful.

Brocade Resources

The following related documentation is provided on the Brocade SAN Switch Module for IBM eServer BladeCenter Documentation CD-ROM, the IBM eServer BladeCenter Web Site, the IBM TotalStorage SAN Switch Web Site or on the Brocade Web Site, through Brocade Connect:

Fabric OS

- Brocade Fabric OS Procedures Guide
- Brocade Fabric OS Reference Manual
- Brocade Diagnostic and System Error Messages Reference Manual
- Brocade MIB Reference Manual
- Brocade Fabric OS v4.2.1 Release Notes

Fabric OS Optional Features

- Brocade Fabric OS Features Guide
- Brocade Advanced Web Tools Administrator's Guide
- Brocade Fabric Watch User's Guide
- Brocade Secure Fabric OS User's Guide
- Brocade Secure Fabric OS QuickStart Guide

Fabric Management Software

- Brocade Fabric Manager User's Guide

Brocade SAN Switch Module for IBM eServer BladeCenter (SilkWorm 3016)

- SilkWorm 3016 Hardware Reference Manual
- SilkWorm 3016 QuickStart Guide

Additional Resource Information

For information about how to use many of the features in this document in a SAN solution, refer to the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version3.0* (53-0000366-03). The latest DDM can be found through the Brocade portal link on the IBM TotalStorage SAN Switch Support Website:
Go to:
http://www.storage.ibm.com/ibmsan/products/2109/library.html#support

Click on any of the links at the bottom of the page for Software Product Manuals. At the redirected site navigate to:
- Software Product Manuals
- Technical "How To" Guides
Home > Technical Resource Center > Documentation Library > Technical "How To" Guides >

There are several "How To" Guides there. The latest DDM Guide is the document entitled "Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0"

The Guides can also be obtained through the Brocade Connect Website:

http://www.brocadeconnect.com

The following related Whitepapers are provided on the Brocade web site and are an excellent resource for additional information.

http://www.brocade.com/san/white_papers.jsp

- *Building Enterprise SANs Through Intelligent Networking*
- *Zoning Implementation Strategies For Brocade San Fabrics*

The following related publications are provided on the Brocade Partner Web Site or the IBM TotalStorage SAN Switch Support Web Site and are an excellent resource for additional information.

- *SAN Migration Guide*
- *LAN Guidelines For Brocade SilkWorm Switches*
- *SAN Security: A Best Practices Guide*

For practical discussions about SAN design, implementation, and maintenance, you can obtain Building SANs with Brocade Fabric Switches through:

http://www.amazon.com

For additional Brocade documentation, visit the Brocade SAN Info Center and click the Resource Library location:

http://www.brocade.com

Release notes are bundled with the Fabric OS.

IBM Resources

The following related documentation is provided on the Brocade SAN Switch Module for IBM eServer BladeCenter Documentation CD-ROM, on the IBM eServer BladeCenter Web Site or on the IBM TotalStorage SAN Switch Web Site.

IBM eServer BladeCenter Documentation

- Brocade Enterprise SAN Switch Module for IBM eServer BladeCenter and Brocade Entry SAN Switch Module for IBM eServer BladeCenter Installation Guide
- IBM eServer BladeCenter Management Module Installation Guide
- IBM eServer BladeCenter Management Module User's Guide
- IBM eServer BladeCenter Management Module Command Line Interface Reference Guide
- IBM eServer BladeCenter Fibre Channel Expansion Card Installation and User's Guide
- IBM eServer BladeCenter Type 8677 Installation and User's Guide
- IBM eServer BladeCenter Type 8677 Hardware Maintenance Manual and Troubleshooting Guide
- IBM eServer BladeCenter Type 8677 Rack Installation Instructions
- IBM Configuration Options Guide

IBM eServer xSeries BladeCenter Marketing Web Site

For additional marketing resource information, visit the IBM eServer BladeCenter Marketing Web site:

http://www.ibm.com/servers/eserver/bladecenter

IBM eServer BladeCenter Support Websites

For additional support resource information, visit the IBM eServer BladeCenter Web site. This Web site provides additional information on the IBM BladeCenter:

http://www.ibm.com/pc/support

http://www-306.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-54597

IBM RedBooks Websites

For additional IBM eServer BladeCenter integration resource information, visit the IBM RedBooks Website and search for "BladeCenter":

http://www.redbooks.ibm.com

IBM TotalStorage SAN Switch Support Websites

For additional resource information, on the complete range of Brocade switch products from IBM visit the IBM TotalStorage Web site. This Web site provides additional information on the IBM TotalStorage SAN Switch product line that is fully compatible with the Brocade SAN Switch Module for IBM eServer BladeCenter:

TotalStorage SAN Switch Marketing Website:

http://www.storage.ibm.com/ibmsan/products/2109/san_switch_solu.html

TotalStorage SAN Switch Support Site:

http://www.storage.ibm.com/ibmsan/products/2109/library.html#support

IBM Support Web Site

For additional general support information, visit the IBM Support Web site.

http://www.ibm.com/pc/support

If you have any questions or problems go to the following Help Center World Telephone Numbers URL:

http://www.ibm.com/planetwide

Other Industry Resources

For additional resource information, visit the Technical Committee T11 Web site. This Web site provides interface standards for high-performance and mass storage applications for fibre channel, storage management, as well as other applications:

http://www.t11.org

For information about the Fibre Channel industry, visit the Fibre Channel Industry Association Web site:

http://www.fibrechannel.org

Getting Technical Help

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

1. General Information
   - Technical Support contract number, if applicable
   - Switch model
   - Switch operating system version
   - Error messages received
   - **`supportshow`** 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, on the side of the unit. The serial number is 12 characters long and looks something similar to this:

   SN: ZXXXXX3WA069

   The serial number label is located as follows:

   - *SilkWorm 3016 switch:* Side of switch module

   It can also be obtained:

   - From a telnet session to the switch, use the "**`chassisshow`**" command it is located in the "Serial Num:" field.
   - From a Web Tools session, Click on the "Info" tab it is located in the "Supplier Serial #" field.
   - From Fabric Manager, Click on the "Switches" tab it is located in a column titled "Supplier Serial Number".
   - From the IBM eServer BladeCenter Management Module, Click on the "Hardware VPD" tab it is located in a column titled "FRU Serial No.".
3. License ID Information. (Required to obtain optional licenses)
   - SilkWorm 3016: Provide the license ID.

   It can be obtained:

   - From a telnet session to the switch, use the "**`licenseidshow`**" command to display the license ID.
   - From a Web Tools session, Click on the "Info" tab it is located in the "LicenseID" field.

Document Feedback

Because quality is our first concern at Brocade, we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you. Forward your feedback to documentation@brocade.com. Provide the title and version number and as much detail as possible about your issue, including the topic heading and page number and your suggestions for improvement.

Chapter 1

Section I

Product Introduction

The *Brocade Enterprise and Entry SAN Switch Modules for IBM eServer BladeCenter Design, Deployment and Management Guide* provides detailed information to help you better utilize the Brocade® SAN Switch Module, licensable features and software. Brocade SilkWorm® switches are available in models ranging from 8-port switches with a few available features to high port count Director class switches providing many licensable features to help you better utilize your fabric and investment.

The following topics are covered:

- IBM eServer BladeCenter Overview
- Brocade SAN Switch Module (SilkWorm 3016)
- Brocade Fabric OS v4.2.1

IBM eServer BladeCenter Overview

IBM eServer BladeCenter is an innovative and manageable modular computing platform that provides outstanding performance density and affordable availability features. This highly integrated infrastructure building block is designed to reduce datacenter complexity and simplify deployment

IBM eServer BladeCenter is a modular server design that gathers computing resources into a cost-effective, high-density enclosure that supports hot-swappable, high-performance 2-way and 4-way Intel processor-based and new 2-way POWER processor-based server blades. Figure 1-1.

A general overview of the features of the IBM eServer BladeCenter Chassis include:

- Space-saving form factor 7U
- High-availability midplane
- Server blade bays - up to 14, 2-way, and up to 7, 4-way
- Standard media USB, CD-ROM and Floppy diskette drive accessible from each server blade
- Switch modules – up to 4 switch module bays
- (2) Power supply modules are standard and there is an option for two more
- (2) Cooling modules are standard
- (1) KVM/management module is standard and there is an option for a redundant KVM/management module
- Innovative design reduces cables by up to 83% saving installation time and money
- Delivers pay as you grow scalability
- Provides fault tolerant connection from the server blade to ALL modular components
- Inventive design allows you to you to upgrade to new technologies and preserve your original investment

Fibre Channel SAN Switch Module options for the IBM eServer BladeCenter:

- Up to 2 Fibre Channel switch modules are supported in switch module bays 3 or 4
- Each switch module has at least one connection to each server blade
- 2 switch modules provide the ability for redundant fabric connections from the chassis
- Switch modules work in active-active redundant fabric configuration
- Each module supports two external port connections

- Each module supports 14 internally facing ports for server blade connectivity

The Fibre Channel Expansion Card option for the IBM eServer BladeCenter Server Blades include:

- Two Fibre Channel ports
- One port connects to a switch module located in bay three
- The second port connects to a switch module located in bay four

Some of the other Server Blade features for the IBM eServer BladeCenter include:

- 1-4 Processors (depending on Server Blade)
- Up to 16GB Memory (depending on Server Blade)
- Optional Hot-swap SCSI or Fixed IDE Drives
- Up to 4 Integrated Gigabit Ethernet controllers (depending on Server Blade)
- Up to 2 Expansion Card connections (depending on Server Blade)

Figure 1-1. Front view of the IBM eServer BladeCenter showing the Standard chassis with 14 Server Blades.

Brocade SAN Switch Module (SilkWorm 3016)

Throughout this document, the term "Brocade SAN Switch Module", "Brocade Switch Module", "SilkWorm 3016" or "switch module" refers to the Brocade® Enterprise SAN Switch Module for IBM ® eServer BladeCenter™ or the Brocade® Entry SAN Switch Module for IBM ® eServer BladeCenter™.

The Brocade SAN Switch Module is a 16-port embedded switch. It supports link speeds up to 2 Gbit/sec. The Brocade SAN Switch Module is based on the Brocade Fabric Operating System™ (Fabric OS) version 4.x, and is compatible with the entire Brocade SilkWorm product family. The main features are indicated below. The Brocade SAN Switch Module is pictured in Figures 1-2 and 1-3.

- 1 or 2 Brocade SAN Switch Modules per IBM eServer BladeCenter Chassis to be placed in Bays 3 or 4
- 14 Internal Ports
    - Ports 1 to 14
    - Connect to IBM eServer BladeCenter Server Blades
    - Fixed at 2Gbit/sec and Server Blades log into the switch as F-ports (requires optional Fibre Channel Expansion Card installed on Server Blade)
- 2 External Ports
    - Ports 0 and 15
    - Connect to existing Fibre Channel SAN switches or directly to Fibre Channel Storage devices
    - Auto-negotiate link speed (1Gbit/sec or 2Gbit/sec)
    - U-port initialization (E-port, F-port or FL-port)
    - Can form a single 4Gbit/sec ISL Trunk (requires optional license)
- Two internal full-duplex 100Mbps Ethernet interfaces, terminated at a single MAC
- Hot Code Activation
- Frame-filtering technology that enables Advanced Zoning and Advanced Performance Monitoring capabilities
- Redundant Power and Cooling provided by IBM eServer BladeCenter chassis

The Brocade SAN Switch Module can only be used inside the IBM eServer BladeCenter chassis. The Brocade SAN Switch Module can be directly connected to external Fibre Channel storage devices or to an existing or new Fibre Channel SAN fabric. To extend the SAN outside of the IBM eServer BladeCenter chassis refer to the IBM TotalStorage SAN Switch family of products. For reference the entire IBM TotalStorage SAN Switch family names and their associated Brocade names are listed in Appendix A.

Brocade Fabric OS 4.2.1

Brocade Fabric OS provides the core infrastructure growing businesses need to deploy scalable and robust Storage Area Networks (SANs). Fabric OS runs on the SilkWorm family of Fibre Channel switches. It supports scalable SAN fabrics that interconnect thousands of devices while ensuring high-performance data transfer among connected resources and servers. Fabric OS easily manages both large switch fabrics and Fibre Channel Arbitrated Loop (FC-AL) configurations. Moreover, Fabric OS is highly flexible, making it easy for network administrators to add functionality and scale their SANs at the speed of business.

Figure 1-2. Brocade SAN Switch Module for IBM eServer BladeCenter.

Figure 1-3. Rear View of the IBM eServer BladeCenter chassis showing two Brocade SAN Switch Modules on the far left in switch module Bays 3 and 4.

Chapter 2

Switch and Fabric Management

This chapter explains the different methods used to manage the Brocade SAN Switch Module and a Brocade SAN.

The following information is discussed:

- Overview of Switch Management
- Fabric Management Recommendations
- High Availability
- User Access Levels
- Switch Licensing
- IBM eServer BladeCenter Management Module
- Fabric OS Command Line Interface
- Advanced Web Tools
- Fabric Watch
- Fabric Manager
- SNMP
- Fabric OS Access Layer (API)
- Management Server

Overview of Switch Management

The Brocade SAN Switch Module can be managed using many local and remote access methods including access from the IBM eServer BladeCenter Management Module Web interface. To manage a switch, you must have access to one of the following available management methods:

- IBM eServer BladeCenter's Management Module Web Interface
- Fabric OS Command Line Interface (CLI) through Telnet or Secure Shell (ssh)
- Advanced Web Tools through a supported Web Browser
- Brocade Fabric Manager a stand alone JAVA application
- SNMP through a Third-Party SNMP management application
- Brocade Fabric Access Application Programming Interface (API) through a Third-Party management product or a Scripting ToolKit Interface
- Management Server through a Third-Party management product

Telnet, Advanced Web Tools, Fabric Manager, SNMP, and the Fabric Access API require that the switch be accessible using a network connection. The IBM eServer BladeCenter provides this network connection through the Management Module's Ethernet port (out of band). The switch module must be configured with an IP address to allow for the network connection. Refer to the Brocade *SilkWorm 3016 Hardware Reference Manual* or to the *Brocade Enterprise SAN Switch Module for IBM eServer BladeCenter and Brocade Entry SAN Switch Module for IBM eServer BladeCenter Installation Guide* for specific information on physically connecting to the switch.

Fabric Management Recommendations

Listed next are some recommendations for managing multiple Brocade switches and fabrics:

- Mixed fabrics

  In a mixed fabric containing Fabric OS v4.x, v3.x, and v2.x switches, manage the fabric via the switch with the latest Fabric OS version as the first criteria, then the most advanced as the second criteria.

- Multiple connections

  Switches can be accessed simultaneously from different connections (for example, Advanced Web Tools, CLI, API and the IBM eServer BladeCenter Management Module). If this happens, changes from one connection might not be updated to the other, and some modifications might be lost. When connecting with simultaneous multiple connections, make sure that you do not overwrite the work of another connection.

- Fabric-level tasks

  A number of management tasks (whether executed from the CLI, Advanced Web Tools, or other management interfaces) are designed to make fabric-level changes: for example, the zoning commands. When executing fabric-level configuration tasks, make sure to allow time for the changes to propagate across the fabric before executing any sequential tasks. For a large fabric, this might be up to a few minutes.

- Command letter casing

  In this guide, all commands are shown as lowercase. Fabric OS v4.x, unlike previous versions of the Fabric OS, is case-sensitive. For backward compatibility, you can enter the commands using the legacy mixed-case notation (for example, **`portCfgShow`**). Lowercase is recommended and always works on both v3.x and v4.x.

High Availability

This section provides information on the Brocade SAN Switch Module High Availability (HA) features.

Brocade SAN Switch Module HA Features

The Brocade SAN Switch Module switches deliver a number of HA features, including:

- Non-disruptive firmware download (external server to flash)
- Non-disruptive code activation

The same mechanisms used to provide non-disruptive features on the SilkWorm 3850/3250 are used to perform non-disruptive reboots and code activation on the Brocade SAN Switch Module.

When a reboot or code activation process is initiated, a reboot management utility is launched. This program creates a standby image of the current switch state. The reboot management program then initiates synchronization, causing all components of the active image to replicate to the standby image. As the state updates are synchronized, messages are saved in Flash.

When all components have finished the synchronization, the active image is disengaged and the actual reboot or code activation operation occurs. Control of the switch is passed seamlessly to the standby image. As a result, there is no effect whatsoever on the flow of data between logged in hosts and storage devices. Read and write tasks proceed with no delay or interruption whatsoever. No frames are dropped. No devices that are currently logged in have to re-login.

User Access Levels

There are four levels of user access to the Brocade SAN Switch Module:

- root
- factory
- USERID
- user

Not all commands are available to all levels. Commands are assigned a minimum login level to execute. Root level has access to all commands; user level has limited access to commands. USERID replaces the "admin" account that exists on all other Brocade SilkWorm switches, and has the same level of administration.

Note Because the USERID account provides access to all the commands needed to manage and configure a switch or fabric, it is the recommended login level.

Switch modules can be accessed simultaneously from different connections (for example, through the CLI and Advanced Web Tools). If this happens, changes from one connection might not be updated to the other, and some changes might be lost. When connecting with simultaneous multiple connections, make sure that you do not overwrite the work of another connection.

In Fabric OS v4.x, each user access level can have the number of simultaneous logins shown:

User Name	Maximum Number of Simultaneous Sessions
root	4
factory	4
USERID	2
user	4

Switch Licensing

The Brocade SAN Switch Module will be shipped in two versions. The first version will be called the Brocade® Entry SAN Switch Module for IBM ® eServer BladeCenter™. This switch will be shipped with the following licenses:

- Advanced Web Tools (Web)
- Advanced Zoning (Zoning)
- Fabric Watch (Fabric Watch)
- Two Domain Fabric License (2 Domain Fabric)

The Two Domain license limits the size of the fabric that this Brocade Entry SAN Switch Module can be a member of. In this case the number is "2". The license can be upgraded using the Brocade Entry Switch Full SAN Upgrade license. This license may be purchased from your switch supplier.

The second version will be called the Brocade® Enterprise SAN Switch Module for IBM ® eServer BladeCenter™. This switch will be shipped with the following licenses:

- Advanced Web Tools (Web)
- Advanced Zoning (Zoning)
- Fabric Watch (Fabric Watch)
- Full Fabric License (Fabric)

With the Full Fabric license the theoretical size limit of the fabric that this Brocade SAN Switch Module can be a member of is 239 switches. Refer to your switch supplier for maximum supported fabric size limits.

The following software features are optional:

- Brocade ISL Trunking for IBM eServer BladeCenter
- Brocade Advanced Performance Monitoring for IBM eServer BladeCenter
- Brocade Performance Bundle for IBM eServer BladeCenter (includes ISL Trunking and Advanced Performance Monitoring)
- Brocade Extended Fabrics for IBM eServer BladeCenter
- Brocade Remote Switch for IBM eServer BladeCenter
- Brocade Advanced Security (Secure OS) for IBM eServer BladeCenter
- Brocade Entry Switch Full SAN Upgrade for IBM eServer BladeCenter (Upgrades from Two Domain to Full Fabric - Entry version only)
- Brocade Fabric Manager v4.x for IBM eServer BladeCenter (stand alone application, must be purchased separately)

To activate and use these optional software features, you must purchase the corresponding license keys. Visit the IBM eServer BladeCenter Web Site for more details.

IBM eServer BladeCenter Management Module

The IBM eServer BladeCenter Management Module provides a Web browser interface that enables you to monitor and manage several key aspects of the individual switch modules from a standard workstation. Only the switch modules enclosed in one IBM eServer BladeCenter chassis are displayed in the Web interface. Detail for the Brocade SAN Switch Modules can be found in the IO Modules sections of the individual panes.

Following are some of the features that make the management module an important part of the switch management and administration process:

- Provides the Hardware and Firmware Vital Product Data (VPD) of the switch module for asset management and support.
- Can power-on, power-off or restart the switch module and determine if extended diagnostics are run on reboot.
- Allows the user to configure or manage the IP Address configuration.
- Displays Power-on-Self-Test (POST) results.
- Can program whether Fast POST is enabled or disabled
- Can program whether the External ports (0 and 15) are persistently enabled or disabled after a reboot or power cycle. If the ports are disabled the Port diagnostics LEDs will flash slowly in amber.
- Can program whether External Management is enabled on the external ports
- Can program whether the IP Address configuration changes are preserved on the switch module
- Restore the switch configuration to factory defaults.
- Send Ping Requests to ensure Ethernet connectivity through the management module.
- Provide a launch location for Telnet and Web Tools sessions to the switch module.

The status of the Brocade SAN Switch Module can be found in the IO Modules section of System Status. See Figure 2-1.

Figure 2-1. IBM eServer BladeCenter Management Module System Status of IO Modules.

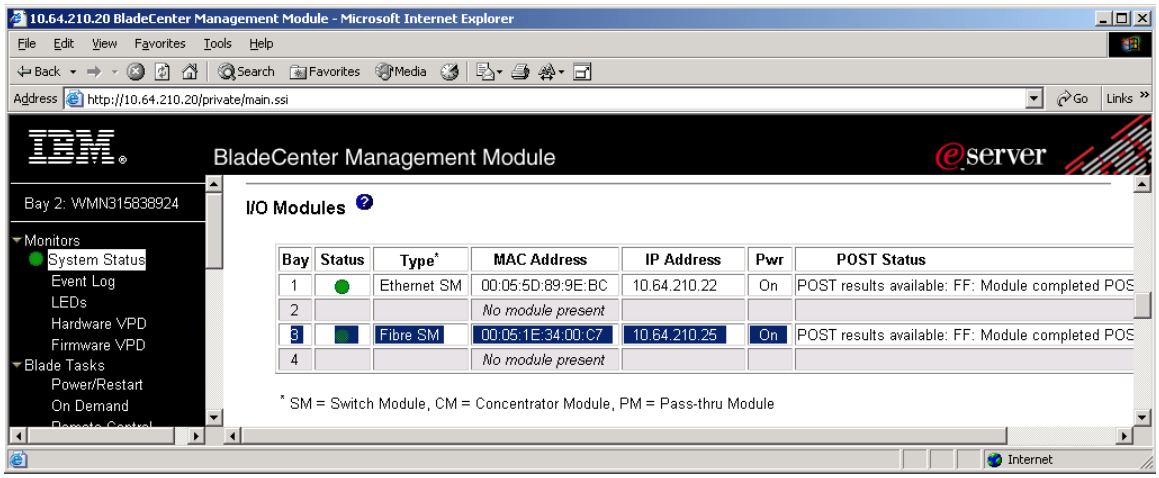

The Hardware VPD tab in the IO Modules section for the Brocade SAN Switch Module contains the switches Serial Number information. See Figure 2-2.

Figure 2-2. Brocade SAN Switch Module Hardware VPD.

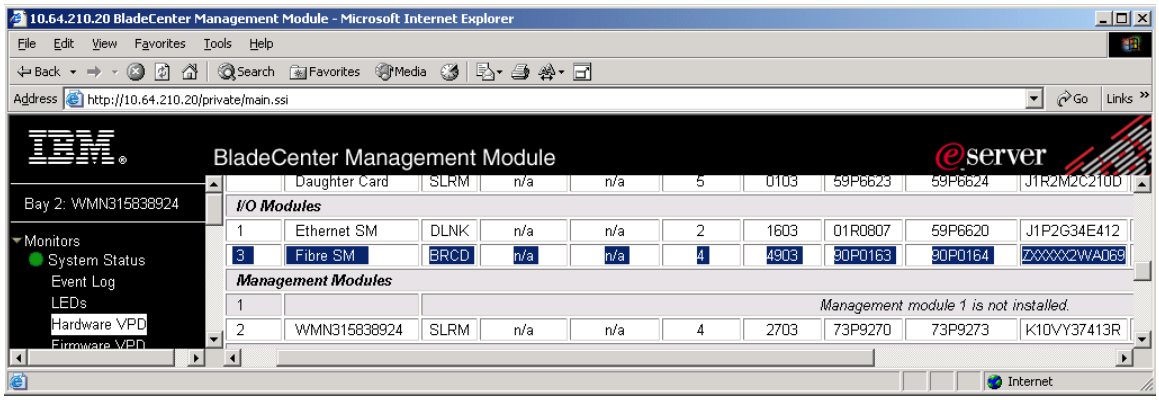

The Firmware VPD tab in the IO Modules section for the Brocade SAN Switch Module contains the firmware revision. See Figure 2-3.

Figure 2-3. Brocade SAN Switch Module Firmware VPD.

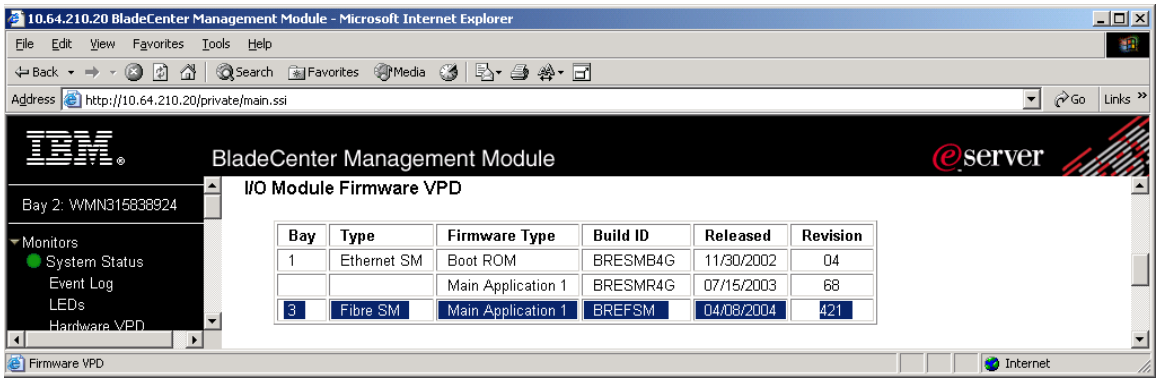

In the IO Modules Task tab click on Management to manage the Brocade SAN Switch Module from the IBM eServer BladeCenter Management Module. See Figure 2-4.

Figure 2-4. IBM eServer BladeCenter Management IO Module Task.

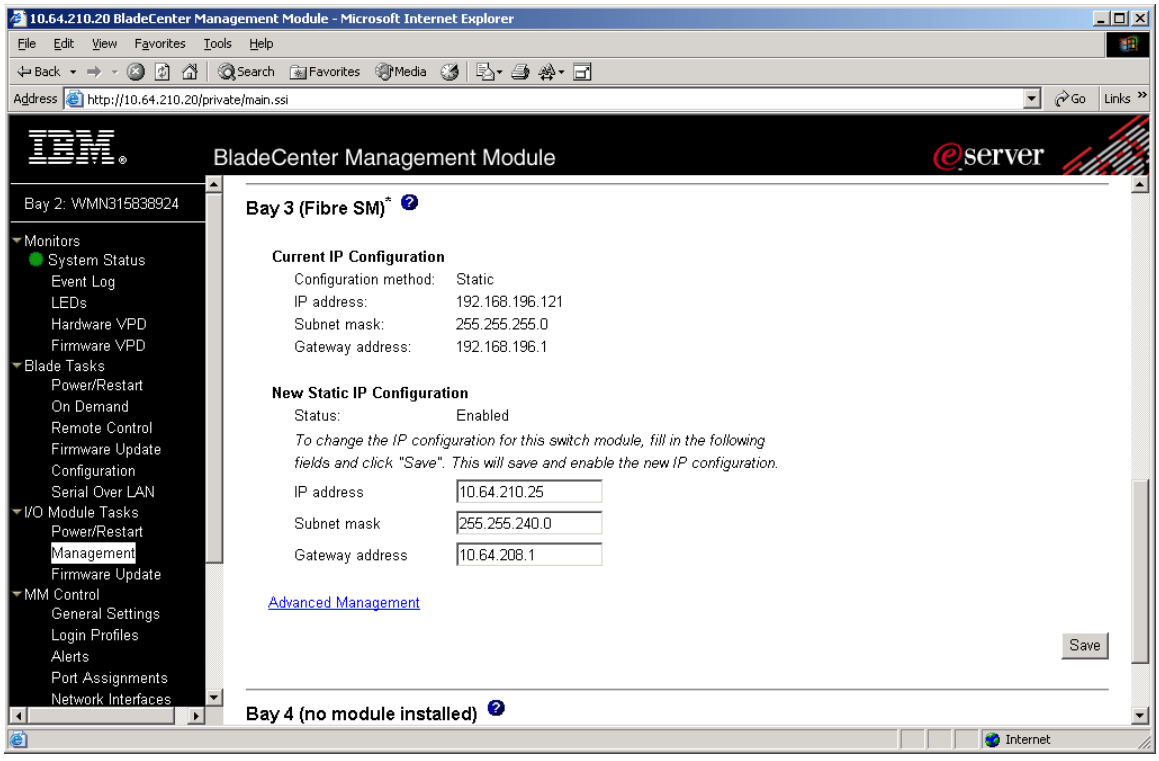

From here the Advanced Management functions of the IBM eServer BladeCenter Management Module for the Brocade SAN Switch Module can be performed. See Figure 2-5.

Figure 2-5. Advanced Management features for the Brocade SAN Switch Module.

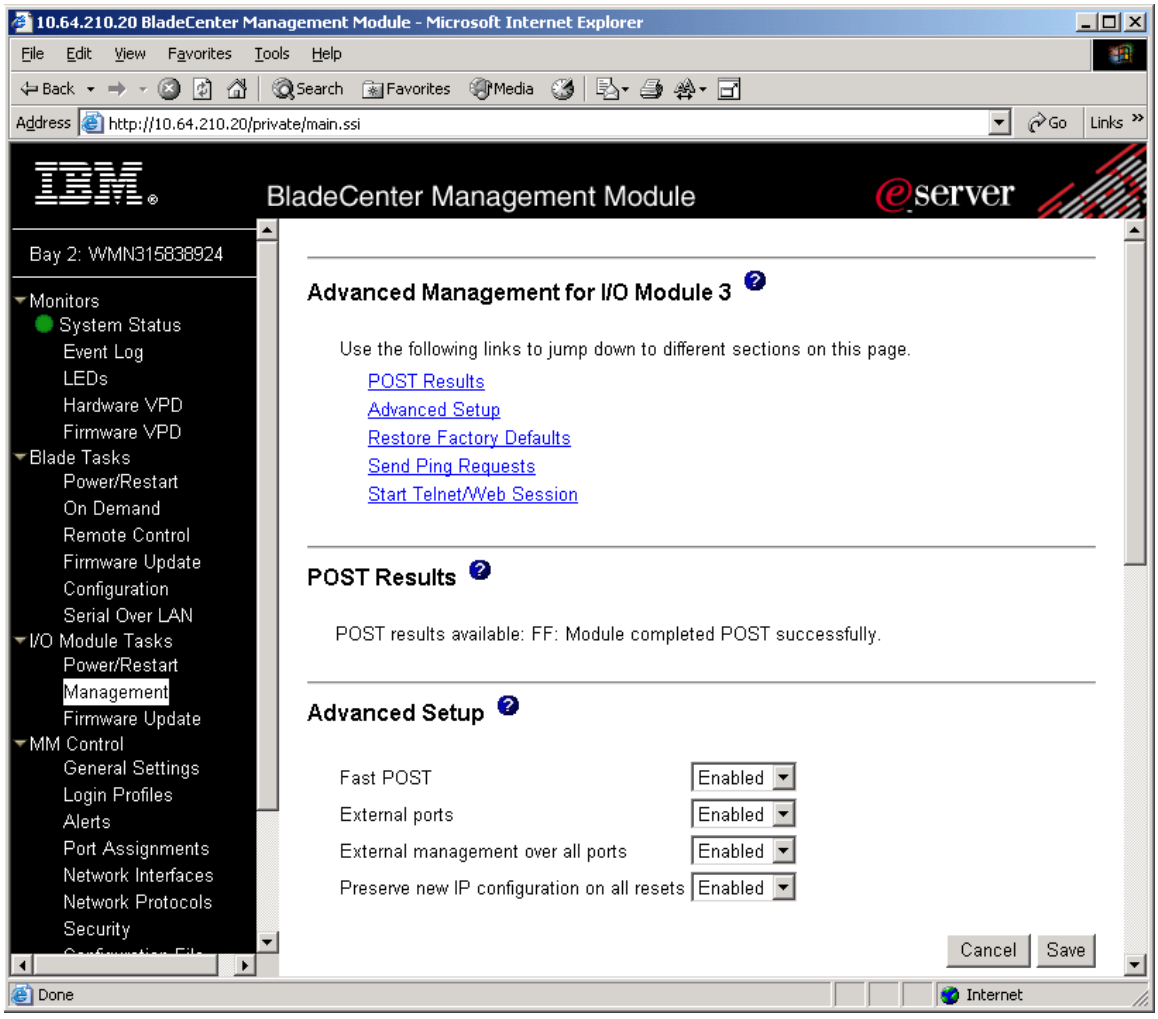

Advanced Management functions of the IBM eServer BladeCenter Management Module for the Brocade SAN Switch Module continued. See Figure 2-6.

Figure 2-6. Advanced Management features for the Brocade SAN Switch Module.

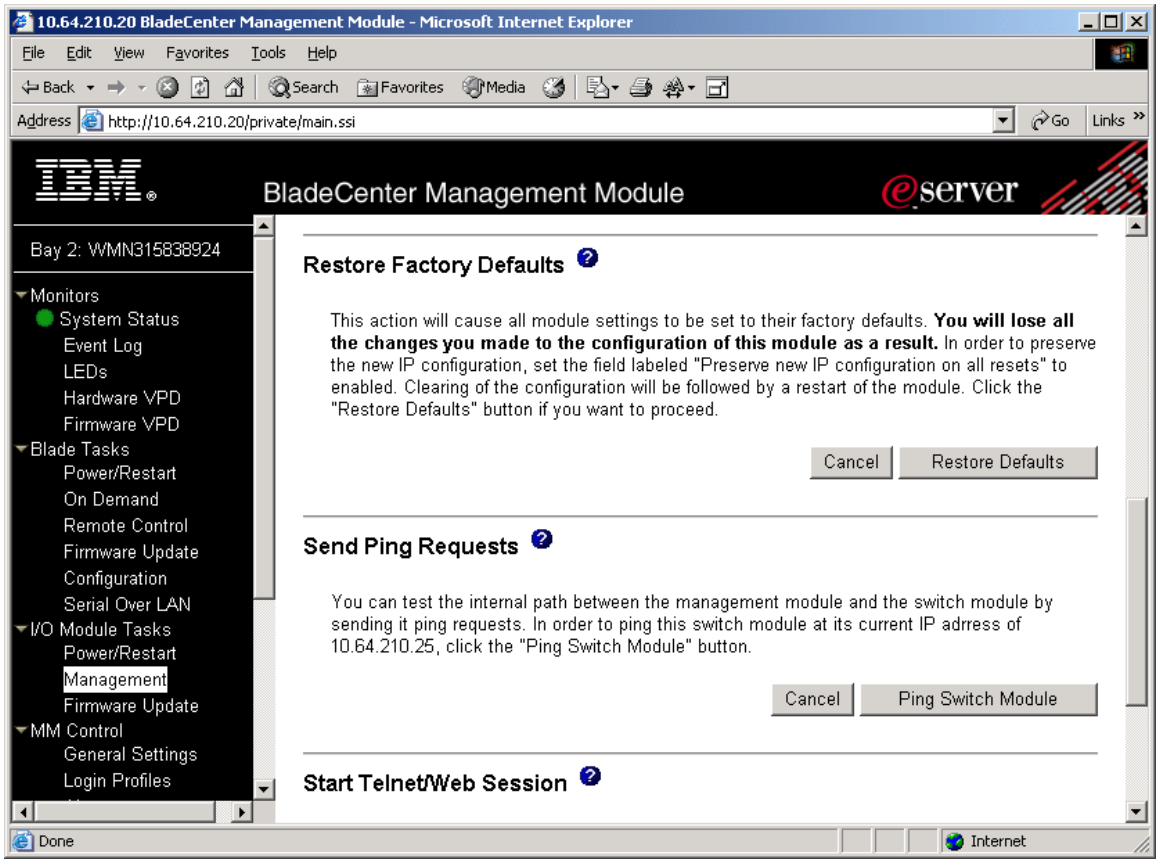

Note — The IBM eServer BladeCenter Management Module should have the latest firmware version. This firmware file is contained in the Management Module FW update package and can be found on the IBM Support website.

For more information, refer to the *Brocade Enterprise SAN Switch Module for IBM eServer BladeCenter and Brocade Entry SAN Switch Module for IBM eServer BladeCenter Installation Guide*.

Fabric OS Command Line Interface

The Fabric OS command line interface (CLI) accessed through telnet (there is no direct Serial Console port on the Brocade SAN Switch Module) provides the user with a full range of management capabilities. The Fabric OS CLI enables an administrator to monitor and manage individual switches, and ports from a standard workstation. Many commands can monitor and manage certain aspects of a fabric as well.

Access is controlled by a switch-level password for each user level (root, factory, USERID, user). The commands available through the CLI are based on the user's login level and the license keys used to unlock certain features.

All configuration and management tasks are available using the USERID or user level; the root and factory levels should be used with caution. The *Brocade Fabric OS Reference* lists all the commands available to the user- and USERID-level users.

Fabric OS CLI is a complete switch management tool for Brocade SAN Switch Modules, providing the following advantages:

- Access to the full range of Fabric OS features, based on which license keys you purchase
- A full set of tools for configuring, monitoring, dynamic provisioning, and daily management of every aspect of storage area networks.
- The ability to configure and manage a Brocade switch on multiple efficient levels
- Fine grain management of every aspect of the switch and its features
- Extensive diagnostic capabilities

The command line interface can be accessed from the Advanced Management link of the IO Module Management Task through a Telnet connection. See Figure 2-7.

Figure 2-7. Advanced Management link of IO Module Management Task.

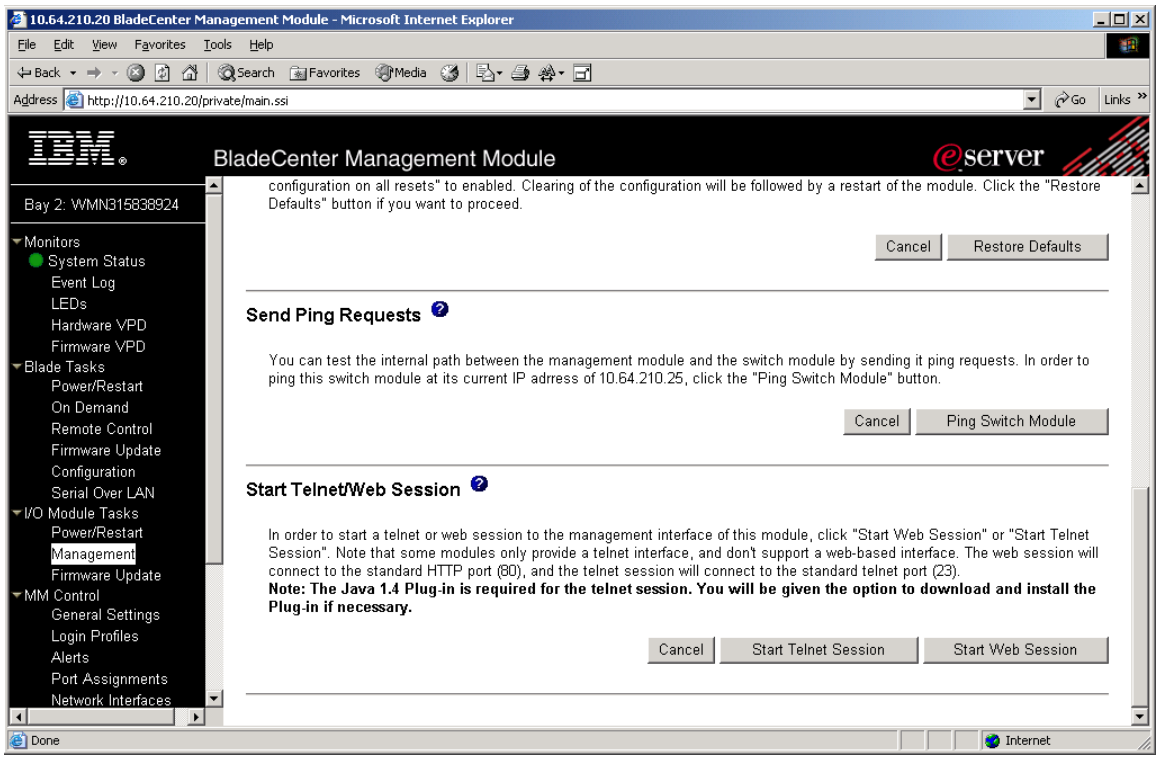

The Telnet Session will look similar to that shown in Figure 2-8.

Figure 2-8. Brocade SAN Switch Module Telnet Session.

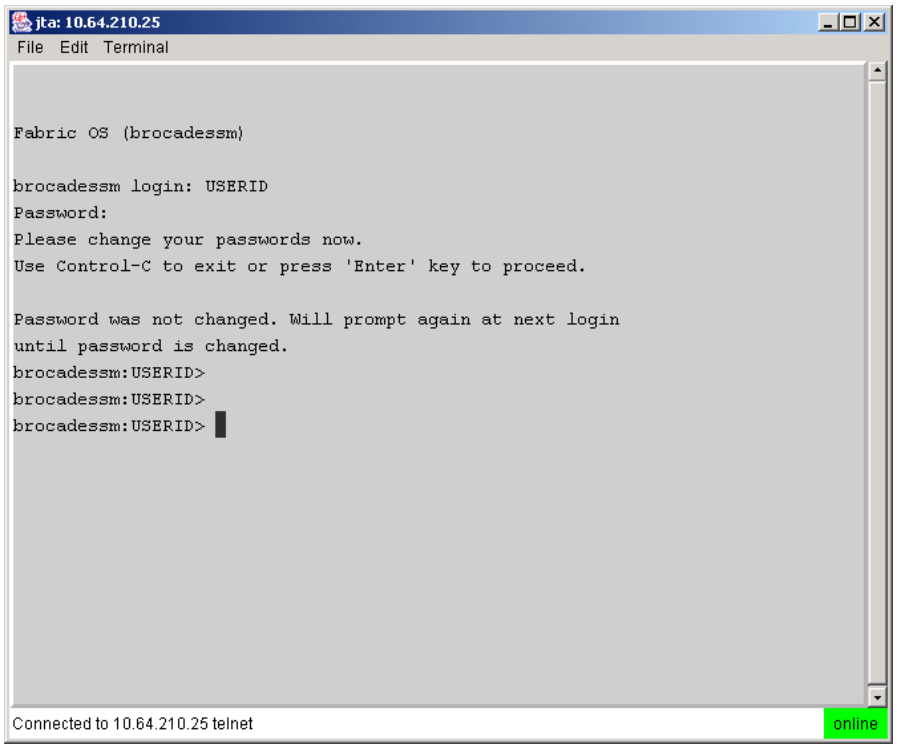

Note — On the first telnet login to the Brocade SAN Switch Module, the user will be challenged to change passwords for three of the four accounts on the system (root, factory and user). The four accounts are: root, factory, USERID and user. This step can be skipped, by using the CONTROL-C key combination, however the challenge will continue at every new login until all of the accounts passwords have been changed. Store the passwords in a safe location.

For more information about using the command line interface, refer to the *Brocade Fabric OS Procedures Guide* and the *Brocade Fabric OS Reference Manual*.

Advanced Web Tools

Brocade Advanced Web Tools is an excellent partner to the traditional Fabric OS CLI commands; in many ways, it provides faster and more effective results than can be achieved strictly through the CLI.

Advanced Web Tools provides a graphical interface that enables you to monitor and manage individual switches and ports from a standard workstation. It is a licensed product that runs on Fabric OS. All switches in the fabric are displayed in the main window of Advanced Web Tools, including switches that do not have an Advanced Web Tools license; however, only those switches that have an Advanced Web Tools license installed can be managed through Advanced Web Tools (other switches must be managed through telnet).

Following are some of the features that make Advanced Web Tools an important part of the switch management and administration process:

- Advanced Web Tools can be used simultaneously with Fabric OS CLI commands. Simply open a second window and you can take advantage of the benefits of both interfaces at the same time.
- Advanced Web Tools can be used from a standard workstation, allowing you to be virtually "in front of" any fabric, switch, or port.
- Advanced Web Tools allows you to zone your configuration by dragging addresses and port numbers rather than laboriously typing them out.
- Advanced Web Tools provides a Performance Monitor feature, enabling you to view the status and traffic of a switch or port in seconds via a variety of effective graphs. Refer to Chapter 5 Advanced Performance Monitoring for more information on this feature.
- Advanced Web Tools is easy and intuitive to use.

All Brocade SAN Switch Modules for IBM eServer BladeCenter have the Web Tools license installed from the factory.

An example of the Advanced Web Tools Graphical User Interface is shown in Figure 2-9.

Figure 2-9. Advanced Web Tools GUI.

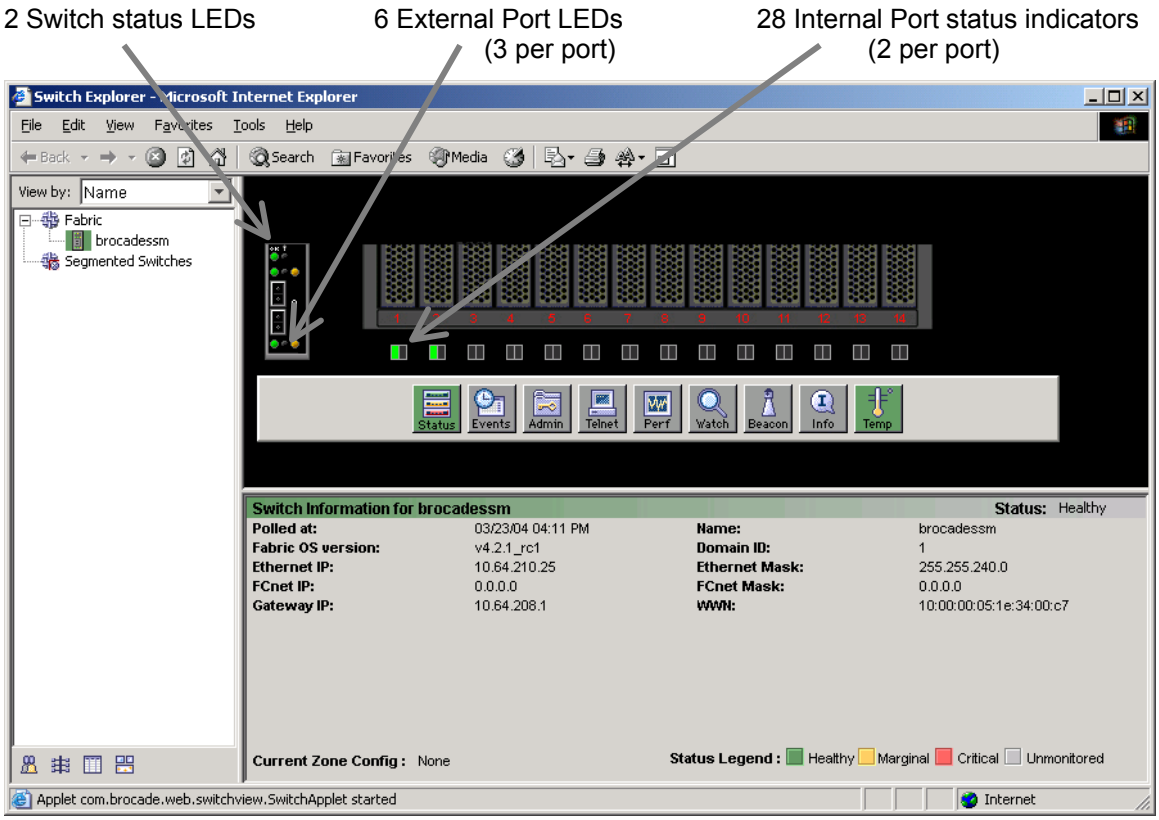

The switch icon for the Brocade SAN Switch Module switch consists of the following:

- External ports and status LEDs
- Internal ports and status indicators
- Switch status LEDs

For more information on the purpose of these LEDs and indicators refer to the *SilkWorm 3016 Hardware Reference Manual* and the *Brocade Enterprise SAN Switch Module for IBM eServer BladeCenter and Brocade Entry SAN Switch Module for IBM eServer BladeCenter Installation Guide*.

For more information on Web Tools itself refer to the *Brocade Advanced Web Tools Administrator's Guide*.

Fabric Watch

Brocade Fabric Watch is a licensed product that monitors the health and performance of Fibre Channel networks and SilkWorm switches, alerting you when problems arise before they become costly failures. SAN managers can configure Fabric Watch software to monitor any of the following:

- Environmental conditions (such as temperature on the Brocade SAN Switch Module)
- SFP behavior (such as temperature, current, voltage and transmit and receive power [serial diagnostic SFPs only])
- Port behavior (such as state changes, errors, and performance)
- Fabric events (such as fabric reconfigurations, segmentations, and zone changes)
- Advanced Performance Monitoring (provided a license key has been purchased and installed)
- Security (such as management violations, security policy violations, and login violations [provided a license key has been purchased and installed])
- Switch Availability (such as down time, up time)

With Fabric Watch software, you can place limits, or thresholds, on the behavior of different switch and fabric elements. Fabric Watch then monitors these behavior variables, or counters, and issues an alarm to address problems when a counter exceeds a threshold. An alarm can send a message to the Event Log (Error Log), SNMP trap, Fabric Access API, E-mail or it can even lock the port log, depending on how you configure the alarm.

All Brocade SAN Switch Modules for IBM eServer BladeCenter have the Fabric Watch license installed from the factory.

An example of the Fabric Watch management tab in the Advanced Web Tools Interface is shown in Figure 2-10.

Figure 2-10. Fabric Watch.

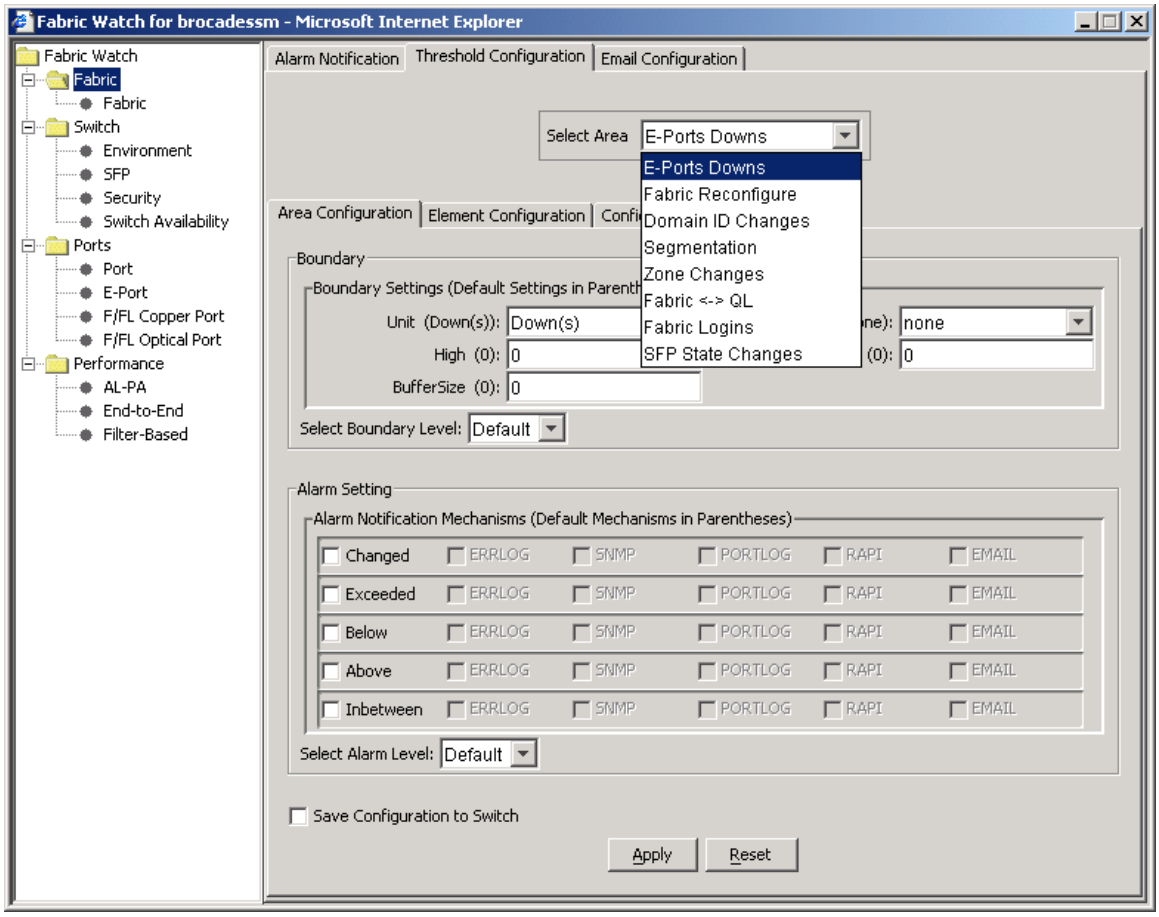

For more information, refer to the *Brocade Fabric Watch User's Guide*.

Fabric Manager

Brocade Fabric Manager is a java-based management application that provides a central point of control to manage multiple fabrics. Fabric Manager is tightly coupled with Brocade Web Tools, Fabric Watch, and Advanced Performance Monitor. Fabric Manager provides a graphical interface for monitoring and managing multiple IBM eServer BladeCenters in multiple fabrics comprised of Brocade switches from a standard workstation. The GUI simplifies task administration at the fabric, switch, and port levels in a medium-to-large size Brocade SAN environment.

Fabric Manager can be used to manage multiple switch fabrics in addition to individual Brocade switches. It provides consolidated high-level information about all the switches in the fabric, launching the Advanced Web Tools application when more detailed information is required for a particular switch. The launching of Advanced Web Tools is transparent, providing a seamless user interface and experience.

Fabric Manager is installed on a workstation. All switches in the fabric are represented in the main window of Fabric Manager, but only those with an Advanced Web Tools license can be managed through Fabric Manager.

Fabric Manager is the complete SAN management tool for Brocade SANs, providing the following advantages to administrators:

- A highly scalable, Java-based application that manages multiple switches and multiple fabrics in real-time
- Time savings by enabling global integration and execution of processes across multiple fabrics, through a single-point SAN management platform
- More effective management by providing rapid access to critical SAN information

- The ability to identify, isolate, and manage SAN events across multiple switches and fabrics
- Drill-down capability to individual SAN components, through tightly coupled Advanced Web Tools, Fabric Watch and Advanced Performance Monitoring integration
- Discovery of all SAN components and ability to view the real-time state of all Brocade fabrics
- Multi-fabric administration of Advanced Security (Secure Fabric OS) through a single encrypted console
- Call Home facility integrated into the IBM TotalStorage support system.

Brocade Fabric Manager 4.1.1 key features are briefly described below.

- SAN Discovery and Topology Display
- SAN Logical Grouping (Switch Groups and Port groups)
- Switch/Port Administration
- At-A-Glance views of switch and device information in the fabric
- Fabric Change Management
- ISL Checking
- Fabric Merge Checking (checks for Zoning and Switch Configuration parameter conflicts)
- Web Tools Device Management launcher
- Telnet (Sec-Telnet) launcher
- Consolidated Event Monitoring
- Switch Firmware Download
- Switch Sequence Reboot
- Switch Configuration Download/Upload (Base-lining)
- Complete Fabric Backup and Compare
- Zoning Management
- Advanced Security (Secure Fabric OS) Policy Management
- Set Time on all switches in a Fabric
- Call Home support
- Switch License Management
- Firmware Download to FDMI capable Host Bus Adapter

Fabric Manager is a stand-alone Software Program that must be purchased from your switch supplier.

Note — The Brocade SAN Switch Module is supported by Fabric Manager version 4.1.1 or higher

---

An example of the Fabric Manager application is shown in Figure 2-11, Figure 2-12 and Figure 2-13.

Figure 2-11. Fabric Manager At-A-Glance Switch view.

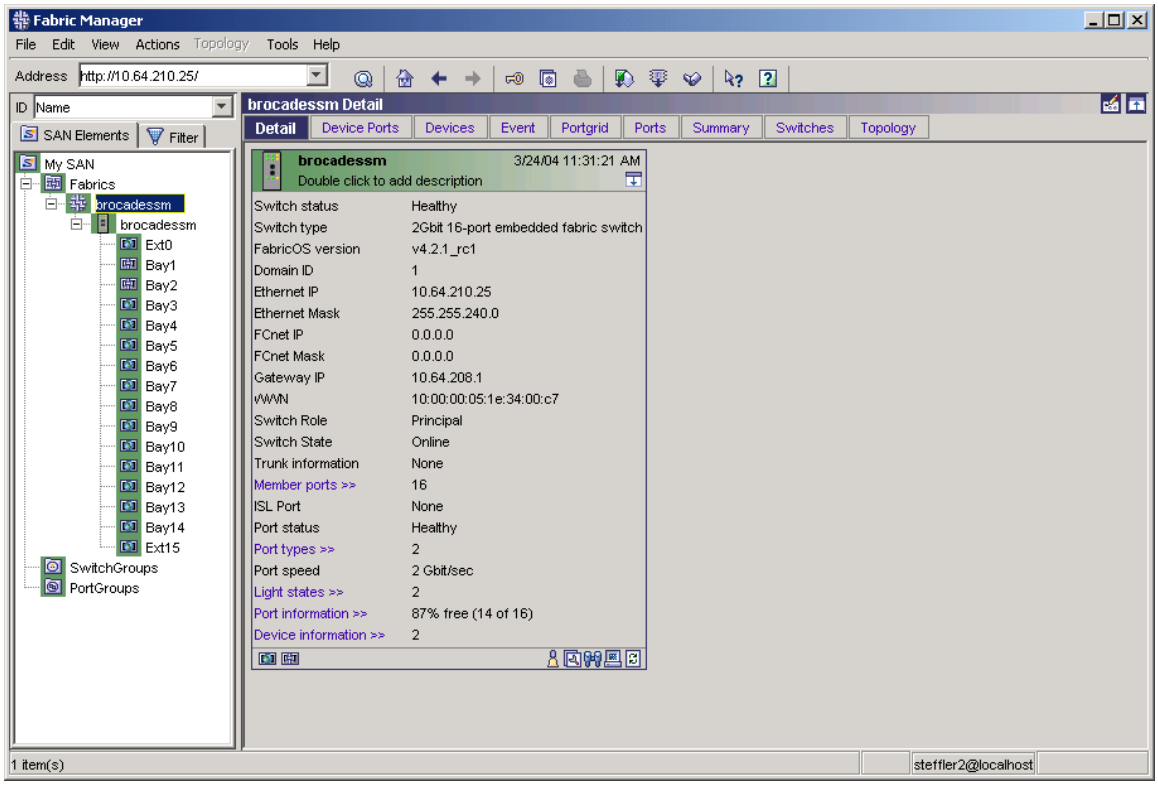

Figure 2-12. Fabric Manager At-A-Glance Fabric view.

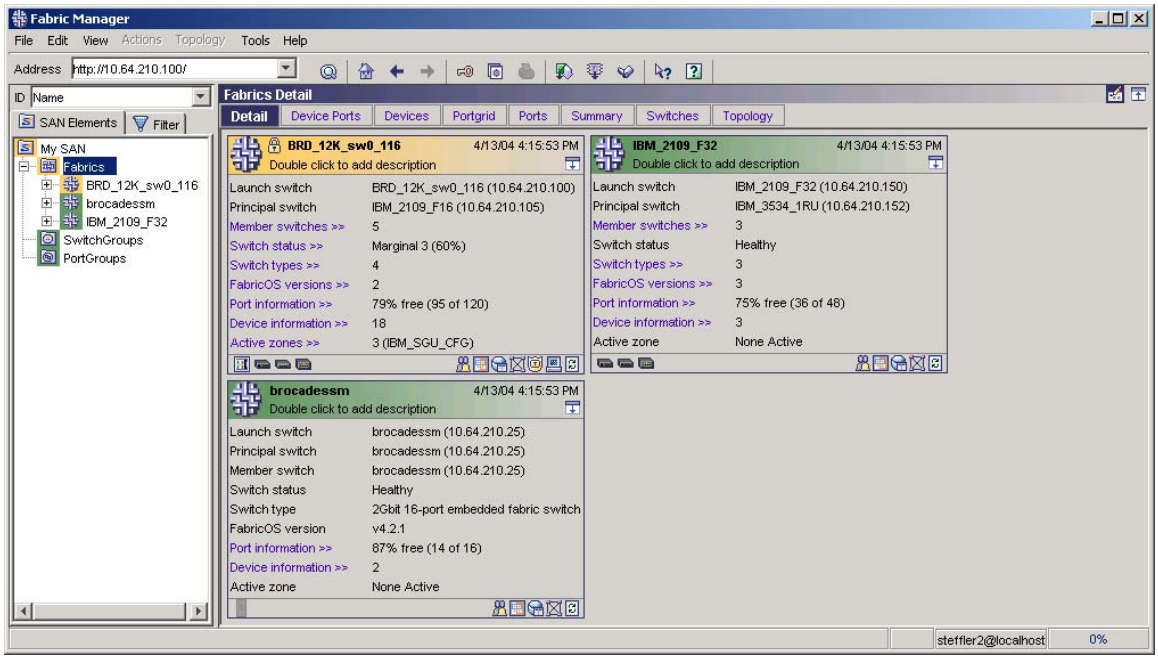

Figure 2-13. Fabric Manager Topology View (2 distinct fabrics shown). The Brocade SAN Switch Module is the top right switch in the fabric on the left.

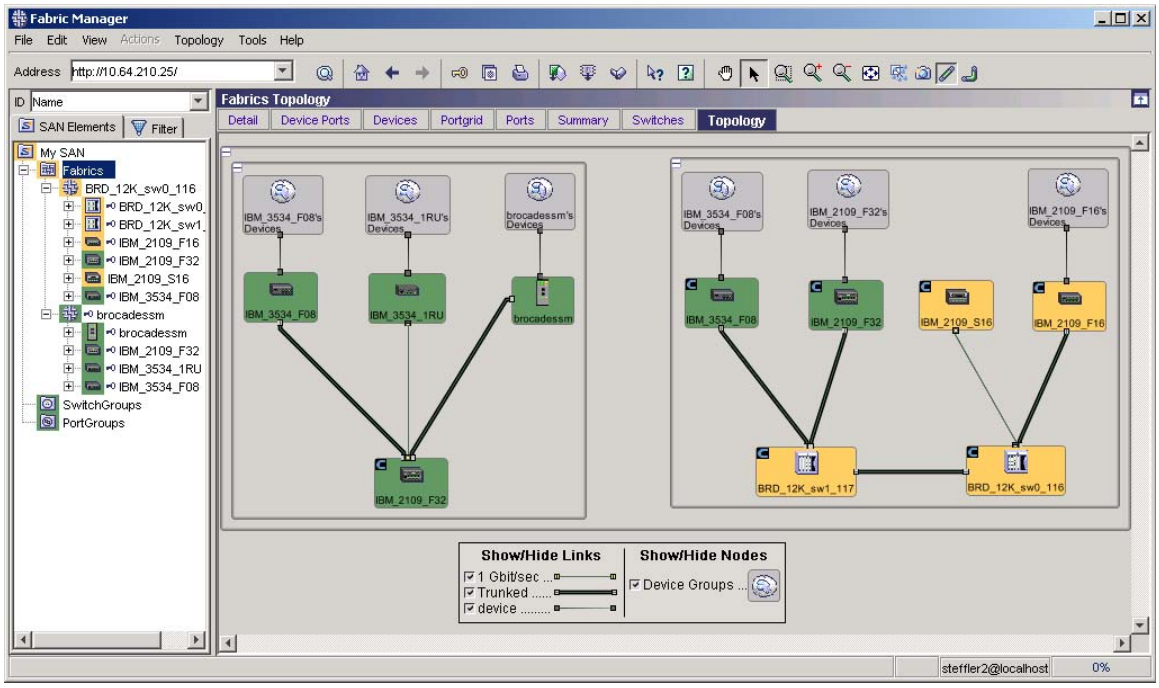

For more information, refer to the *Brocade Fabric Manager User's Guide* and Fabric Manager online help.

SNMP

A common method for monitoring and managing a network device is using the Simple Network Management Protocol (SNMP). Understanding the components of SNMP make it possible to use any tool to view, browse, and manipulate Brocade switch variables as well as set up an enterprise-level management process. Every Brocade switch supports SNMP and carries an *agent* and management information base (MIB). The agent accesses information about a device and makes it available to a network manager station.

When active, the SNMP management station inspects (**get**) or alters (**set**) variables when it queries an agent. The **get**, **getnext**, and **set** commands are sent from the SNMP management station, and the agent replies once the value is obtained or altered. Agents use variables to report such data as the number of bytes and packets in and out of the device, or the number of broadcast messages sent and received. These variables are also known as *managed objects*. All managed objects are contained in the MIB.

When passive, the SNMP management station receives an unsolicited message (trap) from the switch agent if an unusual event occurs.

The Brocade SNMP agent can receive queries from one or more SNMP management stations and can send traps to up to *six* SNMP management stations.

The Brocade SAN Switch Module has two customized SNMP parameter settings. The first parameter that has been customized is the “sysDescr” OID of the MIB-II (RFC1213-MIB). The typical value for all Brocade switch products is “Fibre Channel Switch”. It has been customized to “Brocade SAN Switch Module for IBM eServer BladeCenter”. This can be identified in the CLI “agtcfgset” command or by viewing the “SNMP” pane of the Switch Admin window of Web Tools. See Figure 2-14.

The second parameter that has been customized is the “sysObjectID” OID of the MIB-II (RFC1213-MIB). Its customized value is “1588.2.1.1.22” and is required for IBM Director integration.

Figure 2-14. SNMP configuration settings.

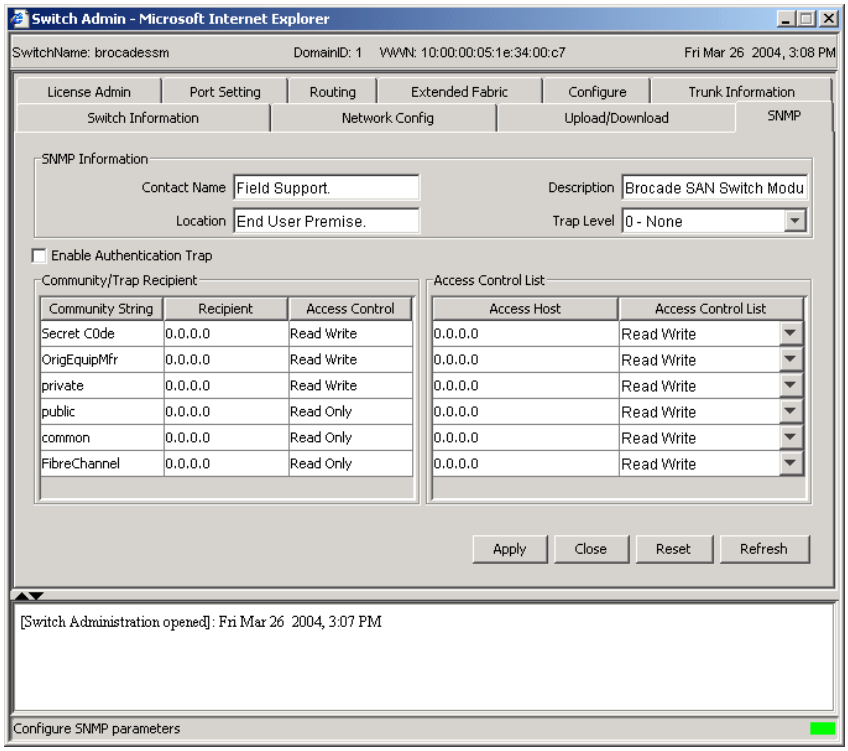

For more information about using Brocade's SNMP agent, refer to the *Brocade MIB Reference Manual*.

Brocade Fabric Access API

The Brocade Fabric Access API is an application programming interface that enables any application to access critical information about a Brocade SAN. With Fabric Access, an application can query or control individual switches or the entire fabric.

The Fabric Access API has the following advantages:

- It can create SAN management applications specific to your needs, using the tools available in the Fabric OS.
- It uses third-party software to manage a Brocade fabric (i.e. IBM Tivoli Storage Area Network Manager)

Note The Brocade SAN Switch Module is supported by Fabric Access API v3.0.2 or higher.

An example of a product that uses the Fabric Access API is IBM Tivoli Storage Area Network Manager.

For more information on the Brocade Fabric Access API, refer to the following web site link:

http://www.brocade.com/products/fabric_access_api.jsp

Management Server

The Fabric OS includes a Distributed Management Server. The Management Server (MS) allows a SAN management application to retrieve information and administer the fabric and interconnected elements, such as switches, servers, and storage devices. The Management Server is located at the Fibre Channel well-known address *FFFFFAh*.

The implementation of the Management Server provides the following management services:

- Fabric Configuration Service (Provides basic configuration management for topology information [referred to as Topology Discovery].)
- Unzoned Name Server access (Provides a management view of the Name Server information for all devices in a fabric, regardless of the active zone set.)
- Fabric Zone Service
- FDMI

The services provided by the Management Server assist in the auto-discovery of switch-based fabrics and their associated topology. A client of the Management Server can determine basic information regarding the switches in the fabric and use this information to construct topology relationships.

In addition, the basic configuration services provided by the management server allow certain attributes associated with switches to be obtained and, in some cases, modified. For example, logical names identifying switches can be registered with the Management Server.

For more information, refer to the *Fabric OS Procedures Guide*.

Note — The Management Server is disabled if the switch is in secure mode. Refer to the Secure Fabric OS User's Guide for more information.

Chapter 3

Advanced Zoning

This chapter contains the following information:

- Overview of Advanced Zoning
- Other Aspects of Advanced Zoning
- Using Zoning to Administer Security
- Zoning Architecture
- Managing Zoning

Overview of Advanced Zoning

Advanced Zoning is a licensed Brocade product that allows partitioning of a storage area network (SAN) into logical groupings of devices that access each other. These logical groupings are called *zones*.

You can use Advanced Zoning to customize environments and optimize resources:

- ***Customize Environments***: You can use zones to create logical subsets of the fabric to accommodate environments such as closed user groups or functional areas within the fabric. For example, you can identify selected devices within a zone for the exclusive use of zone members, or you can define a zone to create separate test or maintenance areas within the fabric.
- ***Optimize Resources***: You can use zones to logically consolidate equipment for efficiency or to facilitate time-sensitive functions: for example, to create a temporary zone to back up non-member devices.

A zone is a specified group of fabric-connected devices, also called *zone objects*, which have access to one another. Zone objects are grouped into zones, and zones are grouped into a zone configuration. Any zone object connected to the fabric can be included in one or more zones. Objects in a zone can communicate only with other objects in the same zone. Zones can overlap; that is, a zone object can belong to more than one zone and a fabric can have multiple zones. A switch can have any number of resident zone configurations; however, only one active configuration can be enabled at a time.

Note — After zoning is enabled, if a device is not explicitly defined in a zone, that device is isolated and inaccessible by other devices in the fabric.

Other Aspects of Advanced Zoning

In addition to the above, zones can be configured dynamically and can vary in size, depending on the number of fabric-connected devices.

Zoning can be disabled at any time. When zoning is disabled, the fabric is in a “non-zoning” mode, and devices can freely access other devices in the fabric.

All devices connected to a fabric can be configured into one or more zones. Every zone must have at least one zone object. Empty zones are not allowed.

Zone configurations are consistent across reboots and power cycles. If two switches are connected in a fabric, they can become isolated (for example, due to an ISL failure); however, when rejoined, they maintain the same fabric configuration unless one of the switches has had a configuration change. In a non-Secure Fabric OS environment the Zoning database is distributed to every switch in the fabric by the Principle switch.

Using Zoning to Administer Security

Zones can provide controlled access to fabric segments and establish barriers between operating environments. They isolate systems with different uses, protecting individual systems in a heterogeneous environment; for example, when zoning is in secure mode, no merge operations occur.

In an Advanced Security (Secure Fabric OS) environment Brocade Zoning is managed from the primary Fabric Configuration Server (FCS). The primary FCS switch makes zoning changes and other security-related changes. The primary FCS switch also distributes zoning to all other switches in the secure fabric. All existing interfaces can be used to administer zoning (depending on the policies--refer to the *Secure Fabric OS User's Guide* for information about security policies).

You must perform zone management operations from the primary FCS switch using a zone management interface, such as telnet or Advanced Web Tools. You can alter a zoning database, provided you are connected to the primary FCS switch.

When two secure fabrics join, the traditional zoning merge does not occur, refer to the *Secure Fabric OS User's Guide* for information for more information.

Note — Advanced Security (Secure Fabric OS) requires the activation of a Brocade Security license and an Advanced Zoning license.

Zoning Architecture

Zoning commands are executed under the transaction model. A working copy of the defined configuration is created at the start of a transaction. Each zoning command is executed from the working copy.

When a transaction is opened, all new zoning information is placed in a transactional buffer. The new changes are not applied to the fabric until the transaction is closed. A transaction is aborted when another switch closes its transaction or by issuing the command ***cfgTransAbort*** . When a transaction is closed, all new and existing zoning information is applied to the fabric and saved to flash memory.

Zone configuration is managed on a fabric basis. Zoning can be implemented and administered from any switch in the fabric that has a Zoning license enabled. When a change in the configuration is saved, enabled, or disabled per the transactional model, it is automatically (by closing the transaction) distributed to all switches in the fabric, preventing a single point of failure for zone information.

Managing Zoning

You can manage Zoning using the following methods:

- Telnet command line interface
- Advanced Web Tools
- Fabric Manager
- Secure Shell
- Fabric Access (API) via third-party software

All Brocade SAN Switch Modules for IBM eServer BladeCenter have the Advanced Zoning license installed from the factory.

An example of the Zoning graphical user interface launched from Advanced Web Tools is shown in Figure 3-1.

Figure 3-1. Advanced Zoning Interface launched from Web Tools.

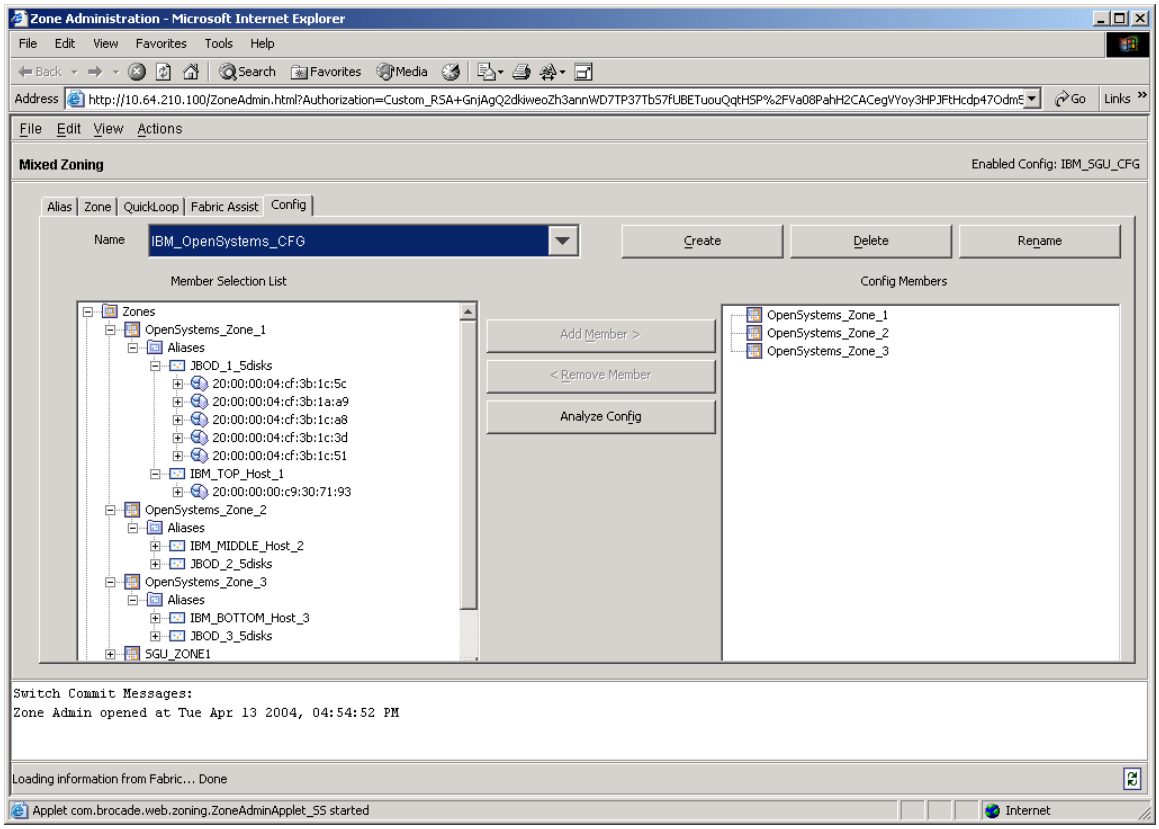

For more detailed information on Advanced Zoning, refer to the following publications and material:

- *Brocade Fabric OS Features Guide,* Chapter 2
- *Brocade Advanced Web Tools Administrator's Guide,* Chapter 6 entitled “Zone Administration”
- *Brocade Fabric OS Procedures Guide,* Chapter 10 entitled “Zoning Procedures”
- *Brocade Fabric Manager User's Guide*, Chapter 6 entitled “Zoning”
- *Fabric OS Reference Manual,* (CLI commands used for managing zoning)
- *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version3.0,* Section 3.3 entitled “Zoning Design Considerations & Guidelines” under SAN Design and Section 5.8 entitled “Zoning Plan” under SAN Deployment
- Zoning Implementation Strategies For Brocade San Fabrics (Whitepaper)

Note, that in terms of Zoning behavior, the Brocade SAN Switch Module behaves exactly like the SilkWorm 3850/3250 as may be described in the above publications.

Chapter 4

ISL Trunking

This chapter provides the following information:

- Overview of ISL Trunking
- ISL Trunking Architecture
- Trunking Ports
- Designing the Fabric for Optimal Use of ISL Trunking
- Managing ISL Trunking

Overview of ISL Trunking

ISL Trunking is an optionally licensed product available for the Brocade SAN Switch Module and is licensed on a per-switch basis. The ISL Trunking feature is provided with the Fabric OS and can be activated by entering a license key available from the switch supplier. It optimizes network performance by forming trunking groups that can distribute traffic across the shared bandwidth of all the ISLs (inter-switch links) in the trunking group. It is compatible with both short wavelength (SWL) and long wavelength (LWL) fiber optic cables and transceivers.

The Brocade SAN Switch Module, ISL Trunking feature allows up to two ISL connections between itself and any other Brocade switch that has an ISL Trunking license installed. This allows the two ISLs to merge logically into a single 4 Gbit/sec link between the two switches (Figure 4-1). This enables traffic to be routed through any available ISL in the group rather than being restricted to a specific, potentially congested ISL. This feature has significant advantages for this architecture. ISL Trunking can result in more optimal throughput by avoiding congestion. ISL Trunking distributes traffic dynamically across the merged ISLs at the fibre channel frame level while preserving in-order delivery of the frames.

ISL Trunking Architecture

The ISL Trunking software identifies and constructs trunking groups as soon as the ISL Trunking license is activated. A license must be activated on each switch that will participate in trunking. The ISLs and ports that participate in trunking groups are referred to as *trunking ISLs* and *trunking ports*.

ISL Trunking makes it possible to accomplish the same fabric performance with fewer ISLs, resulting in simplified fabric design and management, lowered cost of ownership, increased fabric performance, and increased data availability.

Figure 4-1. One Logical ISL made up of 2 physical ISL links.

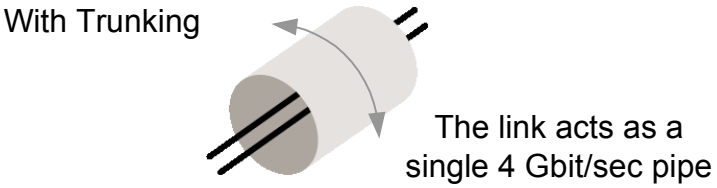

Trunking Ports

The first ISL link that is connected to another Brocade switch that supports ISL Trunking will be the master ISL (*Trunk master*). The master ISL port works to direct traffic over the trunked ISLs. If the second ISL is connected to the same switch and the same Trunk Group (4 port grouping also known as a "quad") on that switch then the second ISL link becomes the secondary E-port of that Trunk Group. If it is connected to different Trunk Group (quad) or a different switch then it will become its own master link and no Trunk Group will form on the Brocade SAN Switch Module.

For more information about Trunk Groups on other Brocade switch models, refer to the *Brocade Fabric OS Features Guide*.

On the Brocade SAN Switch Module all trunking ports must meet the following criteria:

- There can only be one Trunk Group on the Brocade SAN Switch Module consisting of two ports.
- The port speed must be set to auto-negotiate or to 2 Gbit/sec (the default speed is auto-negotiate).
- If Extended Fabrics is in use, the ports must be in L0 mode. Trunking is supported for normal E_Ports (referred to as L0 in the ***portcfglongdistance*** command) with LWL media up to 5km at the full speed permitted by the link. With LWL media, the throughput begins to fall off beyond 5km, due to normal latency effects. ISL Trunking does not support the LE, L0.5, L1, or L2 portcfglongdistance modes at this time. For information about these modes and Extended Fabrics in general, refer to *Fabric OS Features Guide* and the *Fabric OS Procedure's Guide*.

Note — On the Brocade SAN Switch Module external ports 0 and 15 are in the same four-port grouping or "quad" as internal ports 9 and 10. The two internal ports (9 and 10) do not have the Trunking parameter enabled and so will not be affected if Trunking is configured on the external ports (0 and 15).

The diagram in Figure 4-2 shows the possible Trunk Groups on the Brocade 3016.

Figure 4-2. Trunk Port Groupings on a Brocade SAN Switch Module.

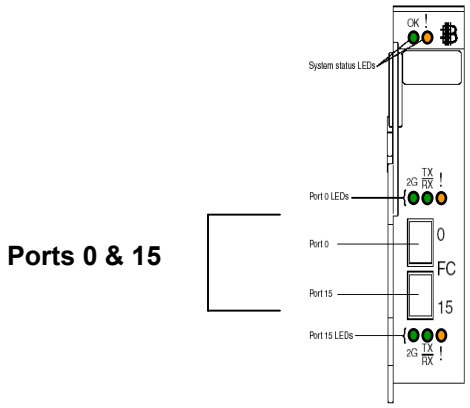

Note — ISL Trunking mode is enabled by default for ports 0 and 15 on the Brocade SAN Switch Module. For instructions on disabling and enabling trunking capability for individual ports or all the ports on the switch, refer to the *Fabric OS Procedures Guide*.
(The Trunking mode is enabled by default, but a Trunking license is still required.)

To facilitate ease of deployment and configuration the Brocade SAN Switch Module has customized port settings. The External Ports 0 and 15 have been configured to auto-negotiate their Link Speed and have their Trunking parameter enabled. A Trunking license must be installed on the Brocade SAN Switch Module as well as a remote Brocade switch. Provided the above guidelines have been followed the switches will automatically create the appropriate Trunking Groups when connected. These settings are the same on all Brocade switch products.

For the internal 14 server connected ports the Trunking parameter has been disabled by default because these ports will never be connected to another switch, they will only exist as server connected ports. In addition, the Link Speed has been configured for a fixed 2Gbit speed for all of the internal ports, as that is the only supported Fibre Channel Expansion Card configuration on the server blades. This can be identified in the CLI "**`portcfgshow`**" command or by viewing the "Port Setting" pane of the Switch Admin window of Web Tools.

As you can see in Figure 4-3, Port Number 0 has a check mark on the "Enable Trunking" box and a value of "Negotiate" in the "Change Speed" box. Port Number 1 does not have a check mark on the "Enable Trunking" box and has a value of "2G" in the "Change Speed" box.

Figure 4-3. Default Port Settings on the Brocade SAN Switch Module.

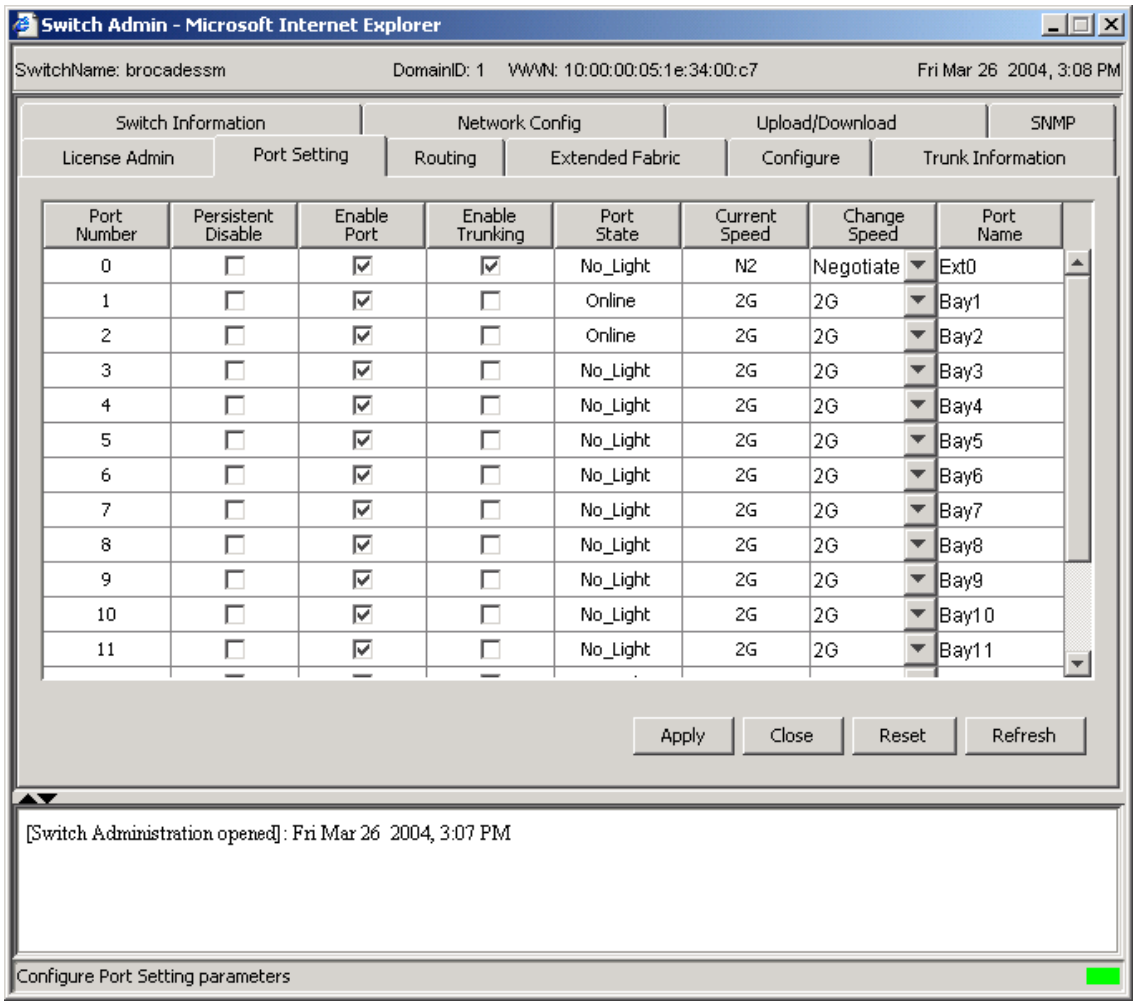

Note Do not change the default port settings on the Brocade SAN Switch Module for any of the Internal ports (1 to 14).

Designing the Fabric for Optimize Use of ISL Trunking

ISL Trunking can be used to simplify SAN (storage area network) design and improve performance. When designing the SAN, consider the following recommendations (in addition to the standard guidelines for SAN design):

- Evaluate the traffic patterns within the fabric. (Discussed in more detail in the *Fabric OS Features Guide* and the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version3.0*). This allows your fabric to implement trunking groups that will help to optimize its performance.
- Activate an ISL Trunking license on each switch that is expected to participate in a trunking group.
- Verify that the lengths of the ISLs in the group do not differ by more than the recommended values. (The lengths of the ISLs in the group should differ by less than 30 meters [recommended], and *must* differ by less than 400 meters. Large differences in latency decrease the efficiency of load-sharing. If cable lengths differ by 400 meters or more, Trunking groups form only for the ISLs that have lengths that differ by less than 400 meters.)
- When connecting two switches with two or more ISLs, ensure that all trunking requirements are met to allow a trunking group to form.
- Trunking groups can be used to resolve ISL over-subscription if the total capability of the trunking group is not exceeded.

Note SWL and LWL fiber optic cables and transceivers can be used in the same trunking group.

Managing ISL Trunking

You can manage Trunking using the following methods:

- Telnet command line interface
- Advanced Web Tools
- Fabric Manager
- SNMP (to view but not set)

The Brocade ISL Trunking license must be purchased from your switch supplier.

An example of the Trunking interface launched from Advanced Web Tools is shown in Figure 4-4. Visibility to existing Trunk groups and connections can be seen in the Topology view of Fabric Manager, Figure 4-5.

Figure 4-4. Trunking Interface launched from Web Tools.

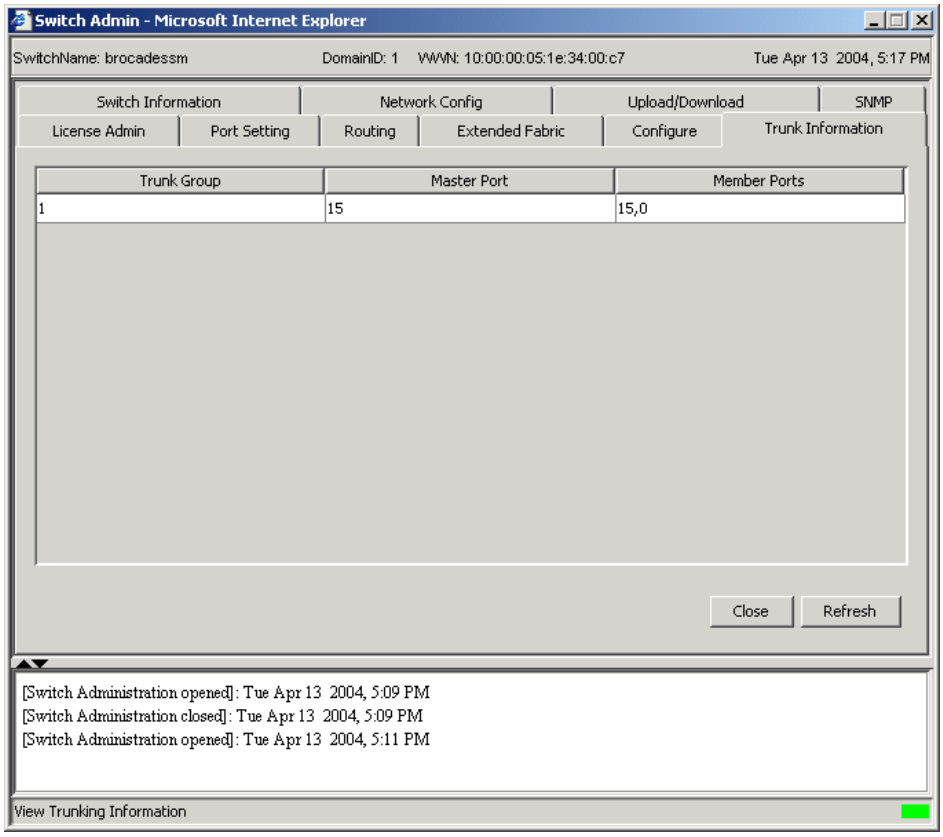

Figure 4-5. Topology view showing Trunked ports in Fabric Manager.

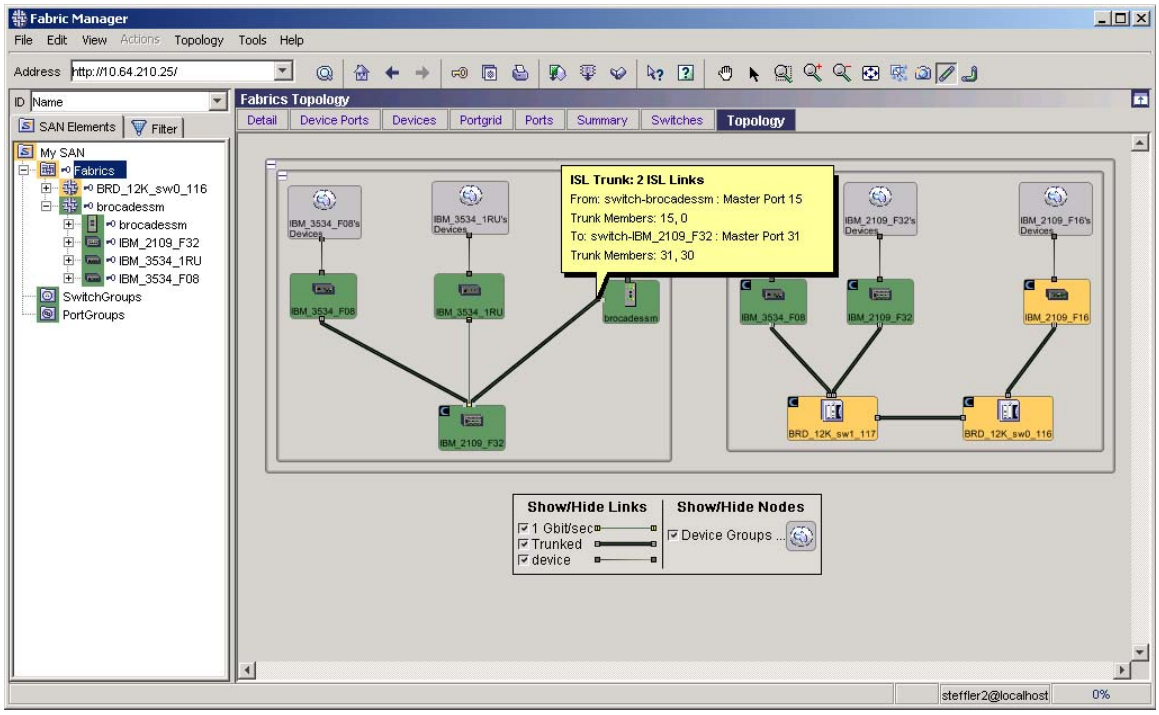

For more detailed information on ISL Trunking, refer to the following publications and material:

- *Brocade Fabric OS Features Guide,* Chapter 5
- *Brocade Advanced Web Tools Administrator's Guide,* Chapter 3 entitled “Managing your Fabrics, Switches and Ports”
- *Brocade Fabric OS Procedures Guide,* Chapter 9 entitled “ISL Trunking Procedures”
- *Brocade Fabric Manager User's Guide*, Chapter 11 entitled “ISL Checking”
- *Fabric OS Reference Manual,* (CLI commands used for managing trunking)
- *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version3.0,* Section 2.3 entitled “Trunking”, Section 2.4 entitled “ISL Over-subscription Ratios” and Section 3.1.1 entitled “Trunk and ISL Connections” under SAN Design

Note, that in terms of Trunking behavior, the Brocade SAN Switch Module behaves similarly to the SilkWorm 3850/3250 as may be described in the above publications (except where noted in this section).

Chapter 5

Advanced Performance Monitoring

This chapter contains the following sections:

- Overview of Advanced Performance Monitoring
- Advanced Performance Monitoring Architecture
- Managing Advanced Performance Monitoring

Overview of Advanced Performance Monitoring

Advanced Performance Monitoring is an optionally licensed product used for monitoring the performance of networked storage resources. This tool helps reduce over-provisioning while enabling SAN performance tuning and increasing administrator productivity. The Advanced Performance Monitoring feature is provided with the Fabric OS and can be activated by entering a license key available from the switch supplier.

Advanced Performance Monitoring provides SAN performance monitoring through an end-to-end monitoring system with the following benefits:

- Increased end-to-end visibility into the fabric.
- More accurate reporting for service-level agreements and charged access applications.
- Shortened troubleshooting time.
- Better capacity planning.
- Increased productivity via preformatted and customized screens and reports.
- The Advanced Performance Monitoring product provides the following functionality:
- Measures the bandwidth consumed by individual routes (host-target pairs).
- Provides device performance measurements by port, AL_PA, and LUN.
- Reports CRC error measurement statistics.
- Compares IP versus SCSI traffic on each port.
- Includes a wide range of predefined reports.
- Allow you to create customized user-defined reports.

Advanced Performance Monitoring Architecture

Using Advanced Performance Monitoring, you can track the following:

- Number of CRC errors for AL_PA devices
- Number of words received and transmitted in Fibre Channel frames with a defined SID/DID pair
- Number of frames with CRC errors received at the port with a defined SID/DID pair
- Number of times a particular filter pattern in a frame is transmitted by a port

These functions are broken down into three configuration groups:

- AL_PA monitoring
- End-to-end monitoring
- Filter-based monitoring

AL_PA Monitoring

AL_PA performance monitoring tracks and displays the number of CRC errors that have occurred on frames sent from each AL_PA on a specific port. AL_PA-based performance monitoring does not require explicit configuration. The switch hardware and firmware automatically monitors CRC errors for all valid AL_PAs.

End-to-End Monitoring

End-to-end monitoring provides information regarding performance between a source (SID) and a destination (DID) on a fabric or a loop. Up to eight SID-DID pairs can be specified per port. For each of the SID-DID pairs, the following information is available:

- Number of fibre channel words received by the port for the SID-DID pair (RX_COUNT)
- Number of fibre channel words transmitted from the port for the SID-DID pair (TX_COUNT)
- Number of frames with CRC errors received at or transmitted from the port for the SID-DID pair (CRC_COUNT)

To enable end-to-end performance monitoring, you must configure an end-to-end monitor on a port, specifying the SID-DID pair. The monitor counts only those frames with matching SID and DID. Each SID or DID has three fields, listed in the following order:

- Domain ID (DD)
- Area ID (AA)
- AL_PA (PP)

For example, the SID 0x118a0f has domain ID 0x11, area ID 0x8a, and AL_PA 0x0f. (The prefix "0x" denotes a hexadecimal number.)

The monitor counts the number of words received, number of words transmitted, and number of CRC errors detected in frames qualified using either of the following two conditions:

- For frames received at the port (with end-to-end monitor installed), the frame SID is the same as "SourceID" and the frame DID is the same as "DestID." Both RX_COUNT and CRC_COUNT are updated accordingly.
- For frames transmitted from the port (with end-to-end monitor installed), the frame DID is the same as "SourceID" and the frame SID is the same as "DestID." TX_COUNT and CRC_COUNT are updated accordingly.

Where to Add End-to-End Monitors

Depending on the application, any port along the routing path can be selected for such monitoring.

Figure 5-1. Setting End-to-End Monitors on a Port shows two devices:

- Host A, connected to domain 0x05, switch area port 12 (area ID 0x0c), AL_PA 0x00 on Switch X
- Device B, connected to domain 0x11, switch area port 30 (area ID 0x1e), AL_PA 0xef on Switch Y

Figure 5-1. Setting End-to-End Monitors on a Port

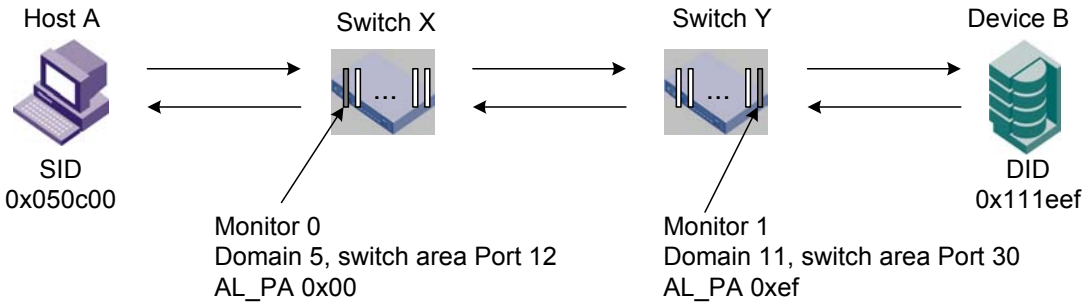

To monitor the traffic from Host A to Device B, add a monitor to port 12, specifying 0x050c00 as the SID and 0x111eef as the DID. To monitor the traffic from Device B to Host A, add a monitor to port 30, specifying 0x111eef as the SID and 0x050c00 as the DID.

Monitor 0 counts the frames that have an SID of 0x050c00 and a DID of 0x111eef. For monitor 0:

RX_COUNT = the number of words from Host A to Dev B
TX_COUNT = the number of words from Dev B to Host A
CRC_COUNT = the number of frames in both directions with CRC errors

Monitor 1 counts the frames that have an SID of 0x111eef and a DID of 0x050c00. For monitor 1:

RX_COUNT = the number of words from Dev B to Host A
TX_COUNT = the number of words from Host A to Dev B
CRC_COUNT = the number of frames in both directions with CRC errors

---

Notes

End-to-end performance monitoring, monitors traffic on the receiving port respective to the SID only. In Setting End-to-End Monitors on a Port, if you add a monitor to Port 30 specifying Device B as the SID and Host A as the DID, no counters, except CRC, will be incremented.

Filter-Based Monitoring

Filter-based monitoring counts the number of times a frame with a particular pattern is transmitted by a port. Filter-based monitoring is achieved by configuring a filter for a particular purpose. The filter can be a standard filter (for example, a read command filter that counts the number of read commands transmitted by the port) or a user-defined filter that you customize for your particular use. The maximum number of filters is eight per port; however, depending on your combination of standard and user-defined filters, that number might be reduced.

Standard Filter-Based Monitors

Using the standard filters, you can collect the following fibre channel frame statistics:

- Number of SCSI read, write, or read/write commands
- Number of SCSI traffic frames
- Number of IP traffic frames

User-Defined Filter-Based Monitors

In addition to the standard filters (read, write, read/write, frame count, and IP count), you can create custom filters to qualify frames to gather statistics to fit your needs.

To define a custom filter, use the **`perfadduermonitor`** telnet command.

Managing Advanced Performance Monitoring

You can manage Advanced Performance Monitoring using the following methods:

- Telnet command line interface
- Advanced Web Tools
- Fabric Manager
- Fabric Watch
- SNMP (to view but not set)
- Use Advanced Web Tools to view, customize, or monitor performance by :
- Viewing predefined reports for AL_PA, end-to-end, and filter-based performance monitoring.
- Creating user-definable reports.
- Viewing performance canvas for application-level or fabric-level views.
- Accessing configuration editor (save, copy, edit, and remove multiple configurations).
- Saving persistent graphs across reboots (saves parameter data across reboots).

Predefined Performance Graphs

Advanced Web Tools provides graphs to simplify performance monitoring. A wide range of end-to-end fabric, LUN, device, and port metrics graphs are included. The table “Predefined Performance Graphs” lists the performance graphs available. You can access the basic monitoring graphs on switches that do not have an Advanced Performance Monitoring license activated. The advanced monitoring graphs give more detailed performance information to help you manage your fabric.

Predefined Performance Graphs

Basic Monitoring Graphs	Description
Port Throughput	Displays the performance of a port in bytes per second for frames received and transmitted.
Switch Aggregate Throughput	Displays the aggregate performance of all ports on a switch.
Blade Aggregate Throughput	Displays the aggregate performance of all ports on a blade.
Switch Throughput Utilization	Displays the port throughput at the time the sample is taken.
Port Error	Displays a line of CRC errors for a given port.
Switch Percent Utilization	Displays the percentage of usage of a chosen switch at the time the sample is taken.

Port Snapshot Error	Displays the CRC error count between sampling periods for all the ports on a switch.
Advanced Monitoring Graphs	**Description**
SID/DID Performance	Displays the traffic between SID and DID pair on the switch being managed.
SCSI vs. IP Traffic	Displays percentage of SCSI versus IP frame traffic on each individual port.
AL_PA Errors	Displays CRC errors for a given port and a given AL_PA.
SCSI Commands by port and LUN (R, W, R/W)	Displays the total number of read/write commands on a given port and read/write commands to a specific LUN.

User-Defined Graphs

You can modify the predefined graphs based on parameter fields such as SID/DID, LUN, AL_PA, and port to create your own customized graphs. These user-defined graphs can be added and saved to canvas configurations.

Using Advanced Web Tools, Advanced Performance Monitoring allows you to set up a *canvas* of performance graphs. The canvas can hold up to eight graphs per window. Multiple canvases can be set up for different users or different scenarios.

Up to 20 individual canvases, each with up to eight graphs, can be saved. Each canvas is saved with a name and an optional brief description.

In addition to the graphs, the Performance Monitoring Resource Usage Display shows which filter slots have been used for each port, and which are available.

Advanced Performance Monitoring also has an extensive set of CLI commands for additional information, refer to the *Fabric OS Reference Manual* and the *Fabric OS Procedures Guide* .

The Brocade Advanced Performance Monitoring license must be purchased from your switch supplier.

An example of the Basic and Advanced Performance Monitoring graphical user interfaces choices launched from Advanced Web Tools is shown in Figures 5-2 and 5-3.

Figure 5-2. Basic Monitoring Graphs provided by Advanced Web Tools.

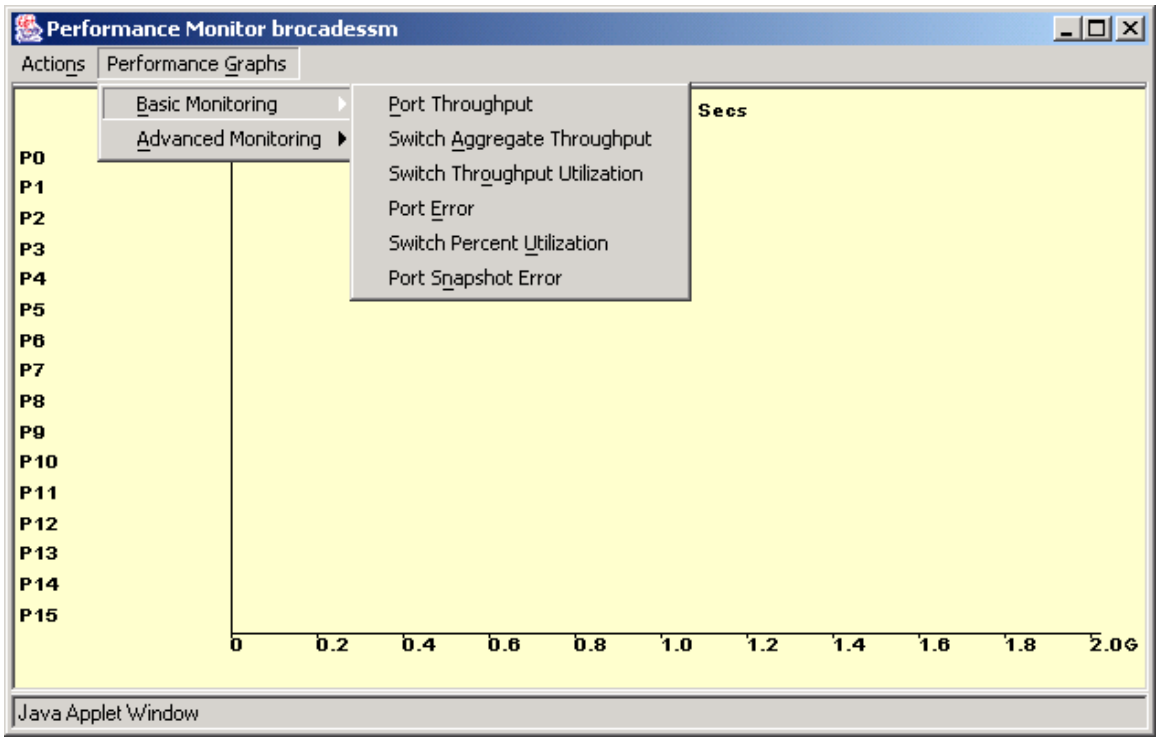

Figure 5-3. Advanced Performance Monitoring Graphs provided by Advanced Web Tools.

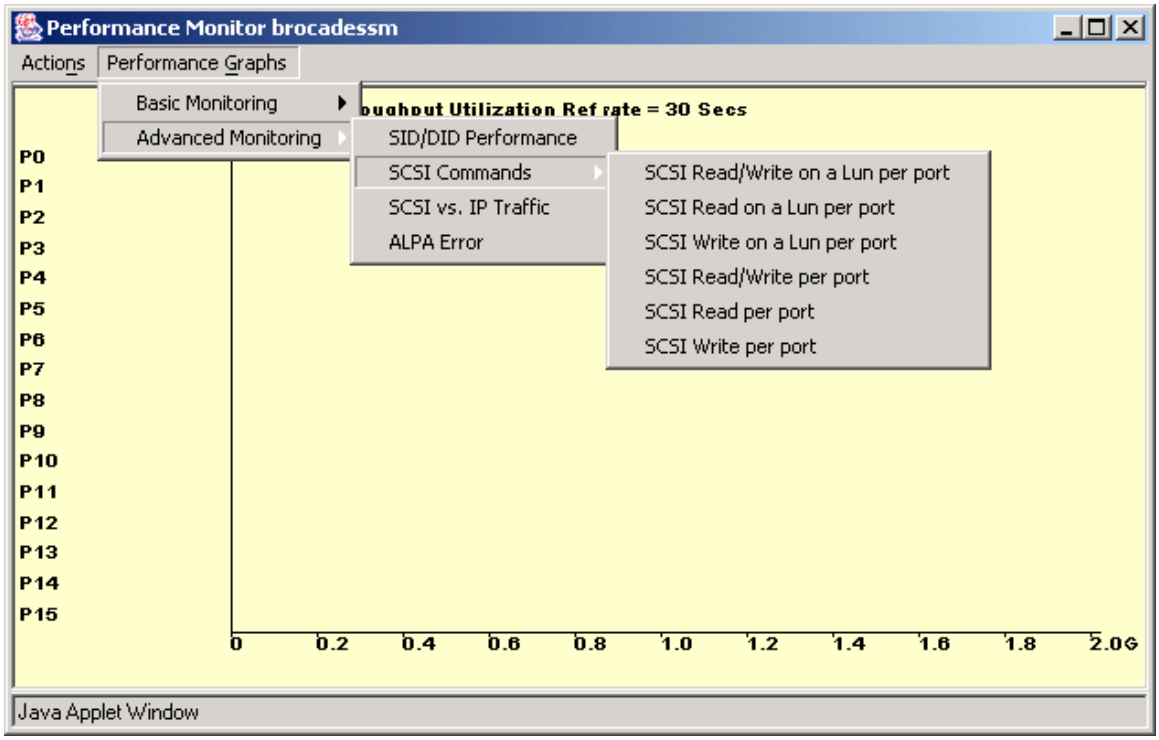

A real life example of using “Port Throughput” from the Basic menu option and “SID/DID Performance” from the Advanced Performance Monitoring menu is shown in Figure 5-4. In this example Ports 0 and 15 are Trunked and each sending a shared load of approximately 60 MB/s (each Trunk sending 30MB) seen in the top graph. There are two Internal Server Blades each talking to their own Disk in a Fibre Channel JBOD that exists on another switch in the SAN.

The SID of the Server in Bay 1 (Port 1) is 0x010100 and the Disk it is talking to on a remote switch has a DID of 0x0307ba (This is shown in the middle graph). This server is communicating with the Disk at about 35 MB/s.

The SID of the Server in Bay 2 (Port 2) is 0x010200 and the Disk it is talking to on a remote switch has a DID of 0x0307bc (This is shown in the middle graph). This server is communicating with the Disk at about 25 MB/s.

Figure 5-4. Advanced Performance Monitoring Graphs in Web Tools.

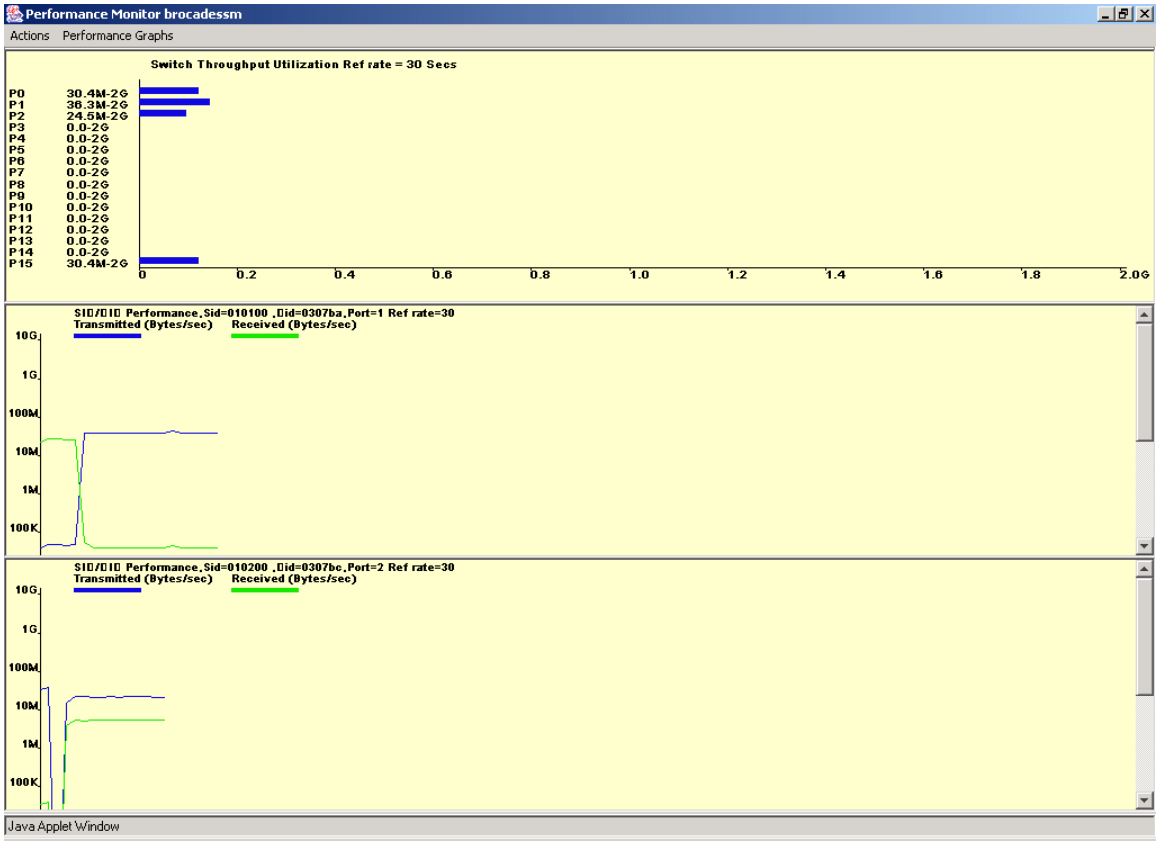

For more detailed information on Advanced Performance Monitoring, refer to the following publications and material :

- *Brocade Fabric OS Features Guide,* Chapter 3
- *Brocade Advanced Web Tools Administrator's Guide,* Chapter 5 entitled “Performance Monitoring Administration”
- *Brocade Fabric OS Procedures Guide,* Chapter 8 entitled “Performance Monitoring Procedures”
- *Brocade Fabric Watch User's Guide*, information on “Performance Monitor Class”
- *Fabric OS Reference Manual,* (CLI commands used for managing performance monitoring)
- *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version3.0,* Chapter 14 entitled “Advanced Performance Monitor Overview” under SAN Management

Note, that in terms of Advanced Performance Monitoring behavior, the Brocade SAN Switch Module behaves exactly like the SilkWorm 3850/3250 as may be described in the above publications.

Chapter 6

Extended Fabrics

Brocade Extended Fabrics is an optionally licensed product that enables inter-switch links (ISLs) to extend up to 100 km. This is achieved by optimizing the internal buffering algorithm used by Brocade switches. It provides maximum buffering between E_Ports connected over an extended distance. The buffer reconfiguration results in line speed performance of 2 Gbit/sec for switches interconnected at up to 60 km (at the moment), and of 1 Gbit/sec for switches interconnected at up to 100km (at the moment). The Extended Fabrics feature is provided with the Fabric OS and can be activated by entering a license key available from the switch supplier.

The Extended Fabric feature achieves long-distance connections by allocating more frame buffers for fibre channel traffic. Long-distance connections require more frame buffers than regular ISL connections. The greater the distance, the more frame buffers required. If the long-distance port is part of a quad, this limits the buffer space left over for the remaining ports in the quad, which must, therefore, be configured appropriately.

Note	On the Brocade SAN Switch Module external ports 0 and 15 share a buffer pool with internal ports 9 and 10 (part of the same “quad”).

The user must be careful when using the Extended Fabric licensed feature with the Brocade SAN Switch Module. The Extended Fabric feature allows the user to configure external ports (0 and 15) for long distance performance. However, certain long distance configurations can disable the other external (and possibly some internal) ports. This could cause a disruption in traffic. When considering configuring external ports for long distance, both the port speed (1 or 2Gbit/sec) and the distance setting (LE, L0.5, L1, L2, and LD) must be considered. The two internal ports 9 and 10 may be disabled due to long distance configuration of the external ports.

Caution	The two external ports of the Brocade SAN Switch Module can be configured as long distance ports, but the user must be aware that these ports (0 and 15) share buffers with the internal ports 9 and 10. Certain long-distance configurations, depending on the length and speed of the links involved, might affect the performance of servers in bays 9 and 10 of the IBM eServer BladeCenter, and in the most extreme cases, can prevent these ports from coming up.

For external ports operating at 2 Gbit/sec the following restrictions should be observed for the Brocade SAN Switch Module.

External Port 0 or 15	External Port 0 or 15	Notes
L2 (60km)	L2 (60km)	Not Allowed
L2 (60km)	L0.5 (25km)	Two internal ports disabled
L2 (60km)	E-port	One internal port disabled
L1 (50km)	L1 (50km)	Two internal ports disabled
L1 (50km)	L0.5 (25km)	One internal port disabled
L1 (50km)	E-Port	One internal port disabled
L0.5 (25km)	L0.5 (25km)	One internal port disabled

Note — When using the LD setting of external ports. LD mode auto-senses the actual cable length and depending on its distances, internal ports 9 and 10 may be disabled.

For external ports operating at 2 Gbit/sec the following are valid Extended Fabrics configurations for the Brocade SAN Switch Module. The internal ports are not affected in the following configurations :

External Port 0 or 15	External Port 0 or 15
E-Port	E-Port
LE (10km)	LE (10km)
L1 (50km)	LE (10km)
L0.5 (25km)	LE (10km)
L0.5 (25km)	E-Port

The long distance Extended Fabric configuration can be established among SilkWorm 3016, 3200, 3250, 3800, 3850, 3900, 12000, and 24000 switches. Long-distance ports consume more buffers than regular ISL ports, which means that configuration of a long-distance port could disable the port itself or other ports in the same quad due to lack of buffers.

Note — Long distance among SW3016, SW3200, SW3250, SW3800, SW3850, SW3900, SW12000, and SW24000 ports is *not* supported when the long distance fabric-wide parameter *fabric.ops.mode.longDistance* is set.

Note — ISL Trunking is *not* supported on an Extended distance ISL at this time.

Guideline	Due to the Extended Fabrics and ISL limitations (only 2 external ISLs) of the Brocade SAN Switch Module, consider using separate switches to connect remote sites and then connect the Brocade SAN Switch Modules to these switches.

The Brocade Extended Fabrics license must be purchased from your switch supplier.

An example of the Extended Fabrics tab launched from Advanced Web Tools is shown in Figure 6-1.

Figure 6-1. Extended Fabrics tab in Advanced Web Tools.

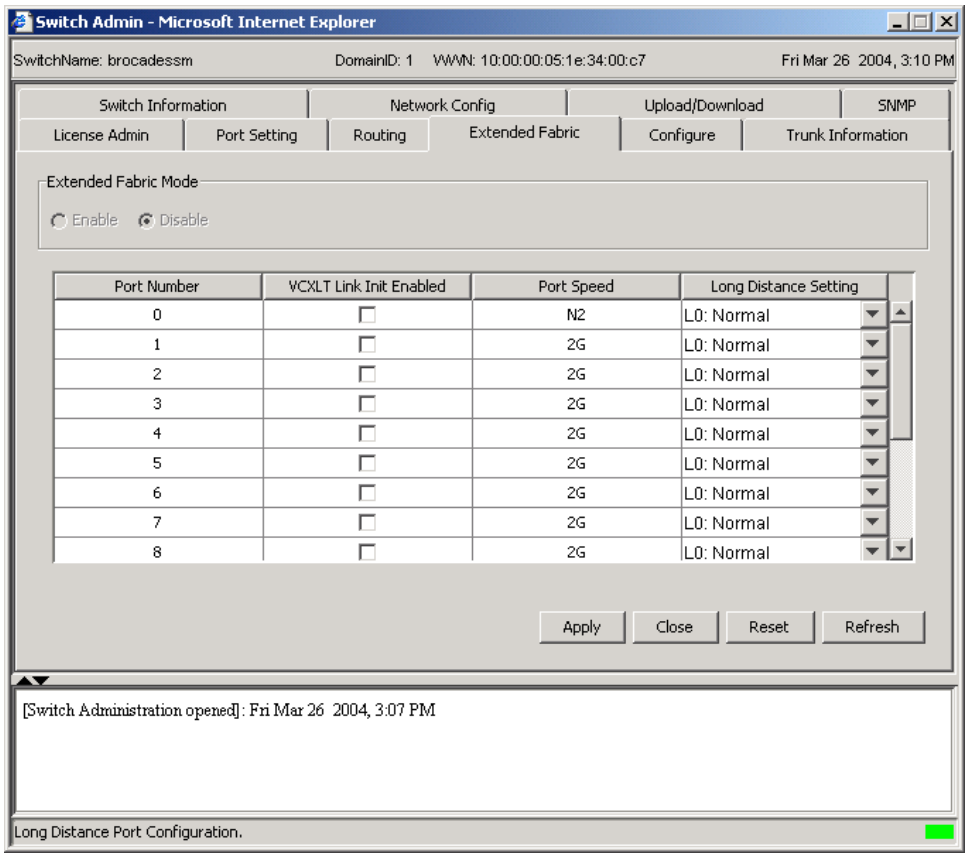

For more detailed information on Extended Fabrics, refer to the following publications and material :

- *Brocade Fabric OS Features Guide,* Chapter 6
- *Brocade Fabric OS Procedures Guide,* Chapter 6 entitled “Distributed Fabrics Procedures”
- *Fabric OS Reference Manual,* (CLI commands used for managing extended fabrics)
- *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version3.0,* Section 5.16 entitled “Brocade Extended Fabrics Planning” under SAN Deployment and especially Appendix B Long-Distance Technologies for Storage Area Networks.

Note, that in terms of Extended Fabrics behavior, the Brocade SAN Switch Module behaves similarly to the SilkWorm 3850/3250 as may be described in the above publications (except where noted in this section).

Chapter 7

Remote Switch

Remote Switch is an optionally licensed product that enables you to connect two remote Brocade fabrics over an IP network, enabling communication of IP or ATM protocols as well as fibre channel traffic. The Remote Switch feature is provided with the Fabric OS and can be activated by entering a license key available from the switch supplier.

The Brocade Remote Switch feature functions with the aid of a "bridging device" or fibre channel gateway. The gateway supports a fibre channel physical interface and a secondary, non-fibre channel physical interface, such as IP, SONET, or ATM. Remote Switch functions over E_Port connections. With Remote Switch on both fabrics, the gateway accepts fibre channel frames from one fabric, tunnels them across the network, and passes them to the other fabric. From the viewpoint of the connected hosts and storage devices, fabrics using Remote Switch interact the same as locally connected switches.

Remote Switch provides many of the same capabilities of normal ISL links including:

- Coordinated fabric services (The Remote Switch fabric configuration fully supports all fabric services, including Distributed Name Service, Registered State Change Notification, and Alias Service.)
- Distributed management (Management tools such as Advanced Web Tools, Fabric OS, and SNMP are available from both the local switch and the remote switch. Switch management is routed through the Fibre Channel connection.)
- Support for inter-switch links (ISLs) (Sites requiring redundant configurations can connect multiple E_Ports to remote sites by using multiple gateways. Standard Fabric OS routing facilities automatically maximize throughput and provide automatic failover during interruption on the WAN connection.)

The Remote Switch feature operates in conjunction with a gateway. The gateway provides an E_Port interface that links to the SilkWorm E_Port. After the link between the two E_Ports has been negotiated, the gateway E_Port moves to pass through mode and passes fibre channel traffic from the SilkWorm E_Port to the WAN.

The gateway accepts fibre channel frames from one side of a Remote Switch fabric, transfers them across a WAN, and passes them to the other side of the Remote Switch fabric.

Remote Switch is automatically activated when you enable the license key. The only required action is to connect the fabrics through the gateway device and make sure that the ***configure*** parameters are compatible with the gateway device.

The Brocade Remote Switch license must be purchased from your switch supplier.

For more detailed information on Remote Switch, refer to the following publications and material :

- *Brocade Fabric OS Features Guide,* Chapter 7
- *Brocade Fabric OS Procedures Guide,* Chapter 6 entitled “Distributed Fabrics Procedures”
- *Fabric OS Reference Manual,* (CLI commands used for managing remote switch)
- *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version3.0,* Section 5.16 entitled “Brocade Extended Fabrics Planning” under SAN Deployment and Appendix B Long-Distance Technologies for Storage Area Networks.

Note, that in terms of Remote Switch behavior, the Brocade SAN Switch Module behaves exactly like the SilkWorm 3850/3250 as may be described in the above publications.

Chapter 8

Brocade Advanced Security (Secure Fabric OS)

This chapter includes the following sections:

- Overview of Advanced Security (Secure Fabric OS)
- Advanced Security (Secure Fabric OS) Architecture
- Managing Advanced Security (Secure Fabric OS)

Overview of Advanced Security (Secure Fabric OS)

Advanced Security (Secure Fabric OS) is an optionally licensed product that provides customizable security restrictions through local and remote management channels on a Brocade fabric. It is available for the Brocade SAN Switch Module and is licensed on a per-switch basis. The Advanced Security (Secure Fabric OS) feature is provided with the Fabric OS and can be activated by entering a license key available from the switch supplier.

- Advanced Security (Secure Fabric OS) provides the ability to do the following:
- Create policies to customize fabric management access
- Specify which switches and devices can join the fabric
- View statistics related to attempted policy violations
- Manage the fabric-wide Advanced Security (Secure Fabric OS) parameters through a single switch
- Create temporary passwords specific to a login account and switch
- Enable and disable Advanced Security (Secure Fabric OS) as desired

Advanced Security (Secure Fabric OS) uses digital certificates based on PKI (Public Key Infrastructure) to provide switch-to-switch authentication.

Advanced Security (Secure Fabric OS) can be used to increase the security of the local and remote management channels, including Fabric Manager, Web Tools, standard SNMP applications, Management Server, and a supported command line interface (CLI) client such as sec-telnet. The access through a channel can be restricted by customizing the Advanced Security (Secure Fabric OS) policy for that channel. Advanced Security (Secure Fabric OS) policies are available for telnet (includes sec-telnet and Secure Shell), SNMP, Management Server, HTTP, and API. Fabric Manager, Web Tools, and API all use both HTTP and API to access the switch. To use any of these management tools to access a fabric that has Secure Mode enabled, ensure that the workstation computers can access the fabric by both API and HTTP. If an API or HTTP policy has been created, it must include the IP addresses of all the workstation computers.

After a digital certificate has been installed on the switch, Fabric OS v2.6.2, v3.1.2, v4.2.0 and v4.2.1 or higher encrypt sec-telnet passwords automatically, regardless of whether Advanced Security (Secure Fabric OS) is enabled.

Advanced Security (Secure Fabric OS) uses digital certificates based on public key infrastructure (PKI) and switch WWNs to identify the authorized switches and prevent the addition of unauthorized switches to the fabric. A PKI certificate installation utility (PKICERT) is provided for generating certificate signing requests (CSRs) and installing digital certificates on switches.

Advanced Security (Secure Fabric OS) Architecture

Fabric Configuration Server Switches

Fabric Configuration Server (FCS) switches are one or more switches that are specified as "trusted" switches (switches that are in a physically secure area) for use in managing Advanced Security (Secure Fabric OS). These switches should be both electronically and physically secure. At least one FCS switch must be specified to act as the primary FCS switch, and one or more backup FCS switches are recommended to provide failover ability in case the primary FCS switch fails.

FCS switches are specified by listing their WWNs in a specific policy called the FCS policy. The first switch that is listed in this policy and participating in the fabric acts as the primary FCS switch; it distributes the following information to the other switches in the fabric:

- Zoning configuration
- Advanced Security (Secure Fabric OS) policies
- Fabric password database
- SNMP community strings
- System date and time

Note — The role of the FCS switch is separate from the role of the principal switch, which assigns Domain IDs. The role of the principle switch is not affected by whether Secure Mode is enabled.

When Secure Mode is enabled, only the primary FCS switch can propagate management changes to the fabric. When a new switch joins the fabric, the primary FCS switch verifies the digital certificate; then it provides the current configuration, overwriting the existing configuration of the new switch. Since the primary FCS switch distributes the zoning configuration, zoning databases do not merge when new switches join the fabric. Instead, the zoning information on the new switches is overwritten when the primary FCS switch downloads zoning to these switches, if Secure Mode is enabled on all the switches. For more information about zoning, refer to the *Fabric OS Features Guide*. For more information about merging fabrics, see the Chapter 4 of the *Secure Fabric OS User's Guide*.

The remaining switches listed in the FCS policy act as backup FCS switches. If the primary FCS switch becomes unavailable for any reason, the next switch in the list becomes the primary FCS switch. A minimum of one backup FCS switch is strongly recommended to reduce the possibility of having no primary FCS switch available. It is possible to designate as many backup FCS switches as desired; however, all FCS switches should be physically secure. Any switches not listed in the FCS policy are defined as non-FCS switches. The root and factory accounts are disabled on non-FCS switches. For information about customizing the FCS policy, see Chapter 4 of the *Secure Fabric OS User's Guide*.

Fabric Management Policy Set

Advanced Security (Secure Fabric OS) supports the creation of several types of policies that can be used to customize various aspects of the fabric. By default, only the FCS policy exists when Secure Mode is first enabled. Advanced Security (Secure Fabric OS) policies can be created and managed by the CLI or Fabric Manager. Advanced Security (Secure Fabric OS) policies can be created, displayed, modified, and deleted. They can also be created and saved without being activated immediately, to allow implementation at a future time. Saved policies are persistent, meaning that they are saved in flash memory and remain available after switch reboot or power cycle.

The group of existing policies is referred to as the "fabric management policy set" referred to as FMPS, which contains an active policy set and a defined policy set. The active policy set contains the policies that are activated and currently in effect. The defined policy set contains all the policies that have been defined, whether activated or not. Both policy

sets are distributed to all switches in the fabric by the primary FCS switch. Advanced Security (Secure Fabric OS) recognizes each type of policy by a predetermined name.

Available Advanced Security (Secure Fabric OS) Policies

Advanced Security (Secure Fabric OS) supports the following policies:

- **FCS** policy. Use to specify the primary FCS and backup FCS switches. This is the only required policy.
- **Management Access Control** (MAC) policies: Use to restrict management access to switches. The following specific MAC policies are provided:
  - Telnet policy. Use to restrict which workstations can use sec-telnet or Secure Shell to connect to the fabric (telnet is not available when Advanced Security (Secure Fabric OS) is enabled).
  - HTTP policy. Use to restrict which workstations can use HTTP to access the fabric.
  - Read and Write SNMP policies. Used to restrict which SNMP hosts are allowed read and write access to the fabric.
  - API policy. Use to restrict which workstations can use API to access the fabric.
  - SES policy. Use to restrict which devices can be managed by SES.
  - Management Server policy. Use to restrict which devices the management server can access.
  - Serial Port policy. Use to restrict which switches can be accessed by serial port.
  - Front Panel policy. Use to restrict which switches can be accessed by front panel.
- **Options** policy. Use to restrict the types of WWNs that can be used for zoning.
- **Device Connection Control** (DCC) policies. Use to restrict which fibre channel device ports can connect to which fibre channel switch ports.
- **Switch Connection Control** (SCC) policy. Use to restrict which switches can join the fabric.

The Brocade SAN Switch Module switch has a different default username than “admin,” which exists on all other SilkWorm switch products. Due to this change, a new command, **`userrename`,** must be used to rename the default “USERID” user account to “admin” before connecting the Brocade SAN Switch Module to a secure fabric made up of other Brocade SilkWorm switches.

When using Advanced Security (Secure Fabric OS), rename the admin-level ID to the Brocade-specific default of "admin" and the user-level ID to the Brocade-specific default of “user” before enabling security. Otherwise, the switch will not be allowed in the secure fabric.

The following error message will display “`Error from domain <domain ID>: Switch does not have all default account names.`”

Note

To rename the admin-level ID from “USERID” to “admin”:

```
switch:admin> userRename USERID admin
```

Managing Advanced Security (Secure Fabric OS)

You can manage Advanced Security (Secure Fabric OS) using the following methods:

- Telnet command line interface (Secure Telnet only)
- Secure Shell (required in a Advanced Security (Secure Fabric OS) environment)
- Fabric Manager

- Fabric Access (API) via third-party software

The Brocade Advanced Security (Secure Fabric OS) license must be purchased from your switch supplier.

An example of the Advanced Security (Secure Fabric OS) graphical user interface launched from Fabric Manager is shown in Figure 8-1.

Figure 8-1. Advanced Security (Secure Fabric OS) Interface launched from Fabric Manager.

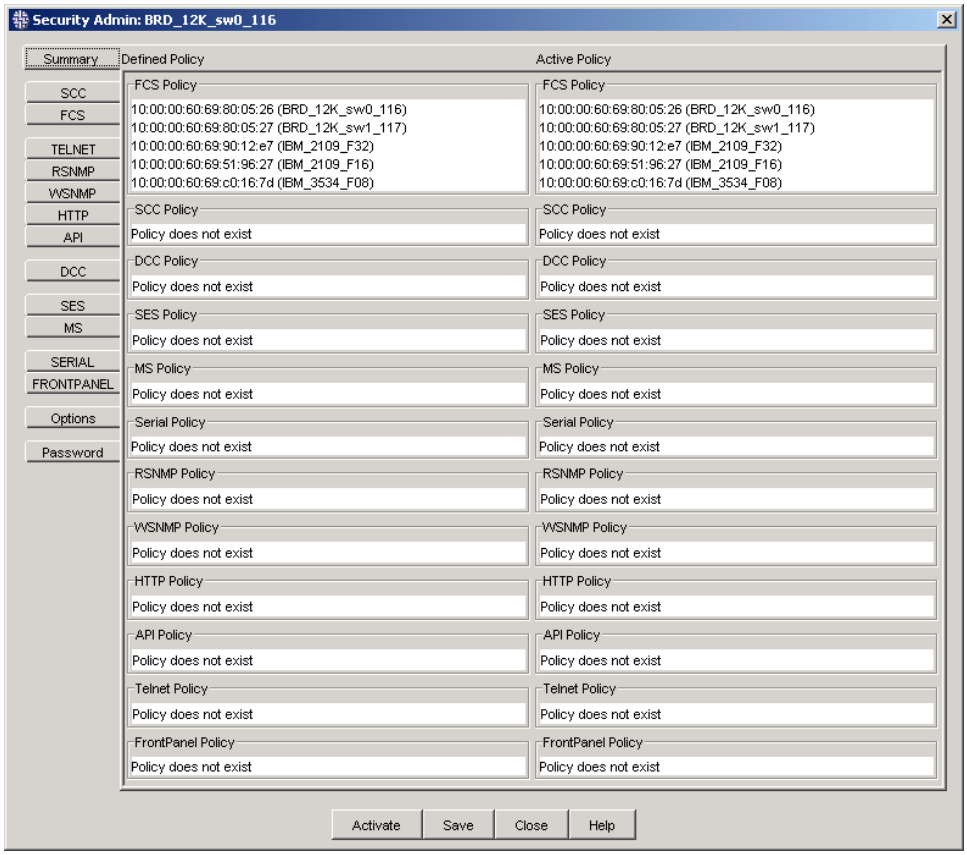

For more detailed information on Advanced Security (Secure Fabric OS), refer to the following publications and material :

- *Brocade Secure Fabric OS User's Guide*
- *Brocade Secure Fabric OS QuickStart Guide*
- *Brocade Fabric OS Procedures Guide,* Chapter 3 entitled "Securing Fabric OS"
- *Brocade Fabric Manager User's Guide*, Chapter 9 entitled "Security Management"
- *Fabric OS Reference Manual,* (CLI commands used for managing security)
- *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version3.0,* Section 3.4 entitled "Designing SANs With Secure Fabric OS" under SAN Design and Section 5.10 entitled "Secure Fabric OS Planning" and Section 6.5 entitled "Staging SAN Fabrics with Secure Fabric OS" under SAN Deployment and Section 9.6.16 entitled "Secure Fabric OS Policy Management" under SAN Management.
- *SAN Security: A Best Practices Guide*

Note, that in terms of Advanced Security (Secure Fabric OS) behavior, the Brocade SAN Switch Module behaves similarly to the SilkWorm 3850/3250 as may be described in the above publications (except where noted in this section).

Chapter 9

Brocade Interoperability Mode

This chapter provides information on setting up a heterogeneous fabric, that is, a fabric that includes both Brocade switches and other manufacturer's switches. For more detailed information on Interoperability, refer to the *Brocade Fabric OS Procedures Guide,* Chapter 12 entitled "Using Interoperability Mode".

- Interoperability
- Brocade Switch Requirements
- McDATA Firmware Requirements
- Supported Brocade Features
- Unsupported Brocade Features
- Configuration Recommendations
- Configuration Restrictions
- Pre-Configuration Planning
- Enabling Interoperability Mode
- Disabling Interoperability Mode

Interoperability

Interoperability mode enables Brocade switches and other manufacturer switch fabrics to exchange interoperability parameters in such a way that both fabrics merge and form one single fabric with one principal switch and all unique domain IDs.

In a heterogeneous fabric, several features are not available in order to provide maximum compatibility between switches.

Use the **`interopmode`** command to enable or disable interoperability mode for individual Brocade switches. This feature enables other manufacturers' switches to be used in a Brocade fabric. This command must be executed on all Brocade switches in the fabric. The switch must be rebooted after changing interoperability mode. Other manufacturers' switches may require the execution of one or more commands that select interoperability mode for their switches.

Brocade SAN Switch Module Requirements

- The Brocade SAN Switch Module must be running 4.2.1 (or later) firmware.
- A Zoning license and a Fabric license must be installed on each Brocade SAN Switch Module.

McDATA Firmware Requirements

- McDATA ES-3016 or equivalent OEM versions that are plug-compatible

---

Note — Contact your switch supplier for their currently supported vendors switches and firmware versions.

Supported Brocade Features

The following features are supported on Brocade switches only:

- Brocade Advanced Web Tools
- Brocade Fabric Watch
- Brocade Fabric Access API functions can be accessed from Brocade switches only, but other manufacturers' switch information may be reported. The object information and zoning actions are configurable from the API.
- Brocade's translative mode, which registers private storage target devices into the fabric, can be used in a heterogeneous fabric as long as the devices are directly connected to Brocade switches. The devices will be accessible from any port on the fabric.

Unsupported Brocade Features

In a heterogeneous fabric, the following Brocade optional licensed features are not supported and cannot be installed on any switch in the Fabric:

- QuickLoop, QuickLoop Fabric Assist or QuickLoop Zoning
- ISL Trunking
- Extended Fabrics
- Remote Switch
- Advanced Security (Secure Fabric OS)

In a heterogeneous fabric, the following Brocade Fabric OS services are not supported on any switch in the Fabric:

- Open E-Port
- Extended Edge PID format
- Broadcast Zoning
- Management Server Service
- Alias Server
- Platform Service
- Virtual Channels
- IP over FC

Configuration Recommendations

The following is recommended when configuring an interoperable fabric:

- Avoid Domain ID conflicts before fabric reconfiguration. There should not be duplicate domain IDs for switches joining the fabric.
- Add switches to the fabric slowly. You should wait for a fabric reconfiguration after adding each switch, when adding multiple switches to a fabric.
- Remove switches from the fabric slowly. You should wait for a fabric reconfiguration after removing each switch, when removing multiple switches from the fabric.

Configuration Restrictions

In interoperable fabrics, the following restrictions apply:

- There is an architecture maximum of 31 switches. However, the actual configuration tested is much less.
- Domain IDs must be in the 97 to 127 value range for successful connection to McDATA switches. The firmware automatically assigns a valid domain ID, if necessary, when the **`interopmode`** command is enabled on the switch.
- **`fabricShow`** only shows the WWN and Domain ID for McDATA. It will not indicate the IP address or switch name. Brocade switches will show all of the above.
- When in Interoperability mode, if managing zoning from the Brocade switches then all Brocade switches must have at least one direct connection to another Brocade switch. For example, you cannot have a McDATA switch in between two Brocade switches if you are managing zoning from the Brocade switches.
- JDSU SFPs will come up as In_Sync and an ISL cannot be established if they are connected to a McDATA ISL. Note : A work around is documented in the *Brocade Fabric OS v4.2.1 Release Notes*.
- When a Brocade switch gets a new domain ID assigned through a fabric reconfiguration, it will write the new domain ID to flash and the old domain ID value will be overwritten. When a McDATA switch gets a new domain ID assigned through a fabric reconfiguration, it will keep the original domain ID in flash. So then, when the domain ID of a McDATA switch and a Brocade switch is changed via fabric reconfiguration, on the next and subsequent fabric reconfiguration, the Brocade switch will try to use the new domain ID (from the flash) while McDATA will try to use it's old domain ID (from the flash). This situation may cause a domain ID overlap to occur during multiple fabric reconfigurations. Domain ID overlap is not supported for Brocade/ McDATA interoperability.
- Between Brocade switches, you can connect more than one ISL when in Interoperability mode.

Note — To determine whether or not a mixed vendor SAN is supported, it will be necessary to work with your switch provider to determine if your SAN design is valid. Important variables that determine the supportability of a particular mixed vendor SAN include the number of switches, version of Fabric OS, the topology, number of ISLs, number of connected devices, and hop count.

Zoning Restrictions

Zoning has the following restrictions in interoperable fabrics:

- Zoning must be consistently and entirely managed from a single selected switch type for the lifetime of the "Effective Configuration". Attempting to manage from both Brocade and McDATA switches may result in the loss of zoning information. The McDATA switch forwards only the Active Zone Set. The Full Zone Set information is not forwarded to the neighboring switches.
- Only zoning by port WWN is allowed.
- Zone members specified by node WWN will be ignored.
- Zone configurations that use either domain, port number or port IDs are not supported in interopmode.
- When there is no zoning configuration in effect, the default effective configuration is all ports are isolated and traffic is not permitted. This is in contrast to the Brocade standard behavior – when interoperability mode is off - where all data traffic is enabled.
- Web Tools can be used for zone configuration as long as Brocade switches are directly connected to each other. If Web Tools is used to setup zoning, then Web Tools must be used as the only zone management method.
- Brocade switches behind a McDATA switch will only receive the effective configuration when a zone merge occurs. This is because McDATA only has an effective configuration and will discard the defined configuration when it sends merge info to the Brocade switch. However, a zone update will send both defined and effective configurations to all switches.

- When a Brocade switch is reconfiguring, do not perform any zoning commands that are supposed to propagate until the fabric routes are fully set up. Use the **fabricShow** command to verify that all of the fabric routes are set up and all of the switches IP addresses and names are present. This does not apply to McDATA as it will only show the WWN and domain ID.
- The maximum number of items that can be stored in the zoning configuration database depends on the switches in the fabric, whether or not interopmode is enabled, and the number of bytes required for each item. The number of bytes required for an item depends on the specifics of the fabric, but cannot exceed 64 bytes per item. At 64 bytes per item you can have 498 entries for a fabric consisting solely of 4.x switches and interopmode enabled.

You can use the **cfgSize** command to check both the maximum available size and the currently saved size. If you believe you are approaching the maximum, you can save a partially completed zoning configuration and use the cfgSize command to determine the remaining space

Zone Name Restrictions

The name field must contain the ASCII characters that actually specify the name, not including any required fill bytes. Names must adhere to the following rules:

- A name must be between 1 and 64 characters in length;
- All characters must be 7 bit ASCII characters;
- The first character of a given name must be a letter. A letter is defined as either an upper case (A-Z) character or a lower case (a-z) character;
- Any character other than the first character must be a lower case character (a-z), an upper case character (A-Z), a number (0-9), or one of the following symbols ($-^_).

Note — When interopmode is in effect, the space available for the zoning database is only half the normal size.

Pre-Configuration Planning

Before enabling interoperability mode, the individual fabrics should be inspected for compatibility.

- Zones should be inspected to ensure that they meet the zone criteria and restrictions. Refer to Zoning Restrictions above.
- Remove or disable any unsupported optional features.
- Disable the Platform Management functions using the **msplmgmtdeactivate** command.

Enabling Interoperability Mode

To enable interoperability mode:

1. Verify that you have implemented all the Brocade prerequisites necessary to enable interoperability mode on the fabric. Refer to Configuration Restrictions and Pre-Configuration Planning.
2. Connect to the switch as the administrator.
3. Disable the first switch, using the **switchdisable** command.
4. At the command line enter the **interopmode 1** command to enable interoperability. This command resets a number of parameters and enables interactive mode.
5. Reboot the switch after changing the interoperability mode.
6. Repeat this procedure on all Brocade switches in the fabric.
7. Other manufacturers switches may require the execution of a similar command to enable interoperability.
8. Once you have enabled inter operability mode on the Brocade switches and other manufacturer's switches, you can cable the other manufacturers switches into the Brocade fabric, one at a time.

Example

```
switch:admin> switchdisable
switch:admin> interopmode 1
The switch effective configuration will be lost when the operating mode is changed; do
you want to continue? (yes, y, no, n): [no] y
done.
Interopmode is enabled
Note: It is recommended that you reboot this switch for the new change to take effect.
switch:admin>
```

Disabling Interoperability Mode

To disable interoperability mode:

1. Connect to the switch as the administrator.
2. Enter the **switchdisable** command to disable the switch.
3. At the command line enter the **interopmode 0** command to disable interoperability. This command resets a number of parameters and disables interactive mode.
4. Reboot the switch after changing the interoperability mode.
5. Wait for a fabric reconfiguration after adding each switch.
6. Repeat this procedure on all Brocade switches in the fabric.

Example

```
switch:admin> switchdisable
switch:admin> interopmode 0
The switch effective configuration will be lost when the operating mode is changed; do
you want to continue? (yes, y, no, n): [no] y
done.
Interopmode is disabled
Note: It is recommended that you reboot this switch for the new change to take effect.
switch:admin>
```

Section II

The second section of the *Brocade SAN Switch Module for IBM eServer BladeCenter Design, Deployment and Management Guide* is focused on covering detailed "how to" information for the Brocade SAN Switch Module from a design, deployment, and management perspective. This portion of the Design, Deployment and Management guide is intended to be used in conjunction with existing Brocade manuals, release notes, and related Brocade publications (especially the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0*).

The flow and organization of this section follows the process of first designing a Brocade SAN with focus on the Brocade SAN Switch Module for IBM eServer BladeCenter, followed by the deployment and operation of that SAN. It is important to understand how new features such as Advanced Security (Secure Fabric OS), non-disruptive code activation, and scalability impact a SAN and how these topics relate across the SilkWorm and Fabric OS family. Many Brocade features span the disciplines of SAN design, deployment, and management. For example, ISL Trunking influences a SAN design, has specific deployment tips, and can be managed via various interfaces such as the CLI, Web Tools, and Fabric Manager.

Discussed in the *SAN Design* section, are topics such as device attachment strategies, switch placement in a fabric, design related Security topics, and Zoning guidelines. Once the SAN design and other planning has taken place, the switches and devices require deployment. The *SAN Deployment* section, covers subjects such as preparation, planning, usage of new features, validating a fabric, and fabric troubleshooting. A rich set of management interfaces exists for the SilkWorm family of switches. Effectively integrating a particular management interface, such as Fabric Manager or Fabric Watch, into the enterprise management system, capacity planning, and SAN management with SNMP are just a few examples of topics addressed in the section on *SAN Management*.

This guide is targeted for use by storage administrators, SAN administrators, system administrators, SAN architects, systems engineers, and SAN operators that are involved with the design, deployment, and management of SANs. Background information and supporting information for a particular topic are kept to a minimum and as appropriate, the reader is referred to supporting documentation. The reader is expected to have working experience with Brocade products. General computer system level troubleshooting skills are always important when configuring sophisticated enterprise solutions. System administration or storage administration experience is also helpful in comprehending this document.

Guidelines are provided throughout the document. Guidelines are recommendations for consideration. The adoption of these guidelines is a function of the user's ability to interpret and correlate relevant SAN information and make decisions based upon their operational policies and SAN requirements.

SAN Design

This section discusses fundamental SAN Design & Architecture concepts and associated guidelines for developing an effective SAN design. A background on SAN design solutions, topologies, availability, scalability, performance, and Trunking is provided as a foundation for the more advanced topics discussed in Chapter 14, *Architecting SANs With SilkWorm Switches*.

This section contains the following chapters:

- SAN Solutions
- SAN Availability
- SAN Scalability
- SAN Performance
- ISL Trunking
- Architecting SANs With SilkWorm Switches
- Zoning Design
- Security Design

Chapter 10

SAN Solutions

While many SAN users begin their SAN experience with one particular SAN solution, the SAN quickly becomes the basis for many other applications. For example, a company may start out with SAN-based backup and quickly integrate storage consolidation and clustering into the existing SAN foundation. In that respect, a SAN decision is a strategic one, and should receive an appropriate level of attention. The adoption of SANs is being driven by a variety of objectives. Some examples are:

- The need for more efficient usage of enterprise storage arrays
- Decreasing size of backup/restore windows
- Increasing size of data set to be backed up
- The need for improved high availability and disaster tolerance
- The need to enhance storage resource management
- Decreased total cost of ownership for storage

Four popular SAN solution categories are Storage Consolidation, LAN-Free Backup, High Availability, and Extended Distance Solutions. Each of these SAN solutions is generically described in the following sections and key attributes of each are discussed in terms of their affect on SAN design.

This chapter contains the following sections:

- Storage Consolidation
- Backup
- High Availability/Clustering
- Extended Distance Solutions

Storage Consolidation

Storage consolidation is a way of optimizing storage resource utilization. It is often the result of migrating directly attached storage (DAS) and hosts to a SAN environment. In a SAN, it is no longer necessary to have a one-to-one correspondence between a host port and a storage port. Instead, many hosts can share a single storage port, and a single host port can access many storage devices. This immediately reduces cost on hosts because fewer HBAs are needed, and on storage because fewer controllers are needed. In addition, savings can be accrued by reducing storage management, power, cooling, and, floor space costs. However, the greatest savings comes from improved utilization of free space on enterprise storage subsystems. With the lowering cost of Fibre Channel HBAs and switch infrastructure, the storage consolidation value proposition has never been better.

Assume that 20 hosts each connect to 100 GB of storage in a direct attach environment, requiring a total 2000 GB of storage. Some space on each system is free. This is known as white space, or headroom. The average utilization of this directly attached storage (DAS) is 50%, leaving 50% white space. The total storage utilized is 1000 GB, which leaves 1000 GB of white space.

With the use of a SAN, it is possible to achieve much higher utilization since every host has access to all storage in the SAN. In this example, a modest 10-20% improvement in storage utilization could result in a savings of several hundred GB of storage. In addition, a reduction in associated ownership costs of that surplus storage would occur. In the storage consolidation model, if a host is not using all of its storage, it is possible to rapidly reallocate this extra storage to a different host. It is also possible to add additional storage for all servers to access, rather than having to purchase storage for specific hosts. In a direct attach environment, it is more difficult to do so, forcing the need to have very high white space overhead to allow growth.

Since many hosts depend upon continuous access to their storage in a storage consolidation solution, designing a highly available SAN to ensure this continuous access is critical. Resilient and redundant fabric designs are highly recommended, especially in large storage consolidation solutions. In a storage consolidation solution, many devices contend for a shared storage port. The performance-limiting factor is often the over-subscription or fan-out ratio of that port, and not the network. Because of this, it is possible to design SANs with a certain amount of over-subscription without adversely affecting application performance. Because the benefits of storage consolidation grow proportionally with the number of hosts and storage, the capability for a SAN to scale is important. It is possible to choose a SAN architecture that can grow from tens of ports to hundreds, and in some cases, thousands of ports, while minimizing or eliminating downtime. Topologies such as the core/edge are optimal for enabling this type of scaling and availability.

Backup

A SAN-based backup is, in some respects, a form of storage consolidation in that an I/O device (the tape drive) is available to be shared by many hosts. The difference is that the shared device is tape, rather than disk. This distinction can affect SAN design in several ways:

- Currently, tape libraries tend to be single-attach, so the multi-pathing approaches used in storage consolidation will usually not work.
- Backup devices tend to be more sensitive to I/O disruption than disk arrays. Arrays can recover from small glitches; tape solutions sometimes do not recover as easily. This is a known issue in the industry and something being addressed with the emergence and adoption of the FC-TAPE standard.
- The availability of tape drives is usually not as critical as that of disk arrays.
- The performance requirements of tape drives typically involve the streaming of large blocks of I/O with bandwidth requirements of 30-60 Mbyte/sec.

Non-SAN based backups take the form of direct attach tape drives, or backup over IP networks. IP backups contend with the normal day-to-day traffic already on the Local Area Network (LAN). Using direct attach tape on each host is costly because of the number of tape devices, tape management, and increased infrastructure cost for floor space, power, cooling, etc.

High-speed SAN-enabled backups reduce backup and restore windows and can enable disaster tolerance by locating libraries at remote sites. SAN based backup improves on traditional backup by enabling the sharing of fewer, larger tape libraries and by minimizing or eliminating the performance issues associated with traditional backup architectures.

A disruption in a backup SAN is usually not as critical as a disruption in a storage SAN. Mission critical applications require continuous access to storage, while a tape backup normally can be restarted without end users seeing the effect. Therefore, a SAN architecture solely used for backups may not require the highest availability enabled by a dual fabric architecture, a single resilient fabric may provide sufficient availability.

High Availability / Clustering

High-availability (HA) clusters are used to support critical business applications. They provide a redundant, fail-safe installation that can tolerate equipment, software, and/or network failures, and continue running with as little impact upon business as possible. HA clusters have been in use for some time now. However, until the advent of Fibre Channel, they were very limited in size and reliability. This is because clusters require shared storage, and sharing SCSI storage subsystems is difficult and unreliable. In fact, sharing a SCSI device between more than two initiators can be difficult due to SCSI cabling limitations, and the SCSI support level for multiple initiators.

Clustering technology has therefore been greatly enhanced by the network architecture of SANs. SANs provide ease of connectivity, and the ability to interconnect an arbitrarily large number of devices. SANs can support as few as two hosts in a failover configuration, and can be expanded to support “many-to-one” configurations. The primary advantages that a SAN affords a cluster are connectivity, scalability, and reliability.

Extended Distance Solutions

SANs enable the ability to replicate data over distances of hundreds of meters to thousands of kilometers. The data is replicated via mirroring at the host level or at the storage level. Data can also be backed up to a remote site. In a Business Continuance example disk mirroring can occur by day and tape vaulting, can occur over IP, by night. Should

the primary site become disabled, the data located at the remote site is accessible immediately so that critical applications can continue operation. Some business continuance implementations deploy standby hosts that remain idle until needed. Other sites utilize the idle equipment for testing and development purposes until that equipment is needed for production purposes. A key element to this strategy is having the hosts boot from the SAN and the existence of a boot image of the primary site hosts located at the remote site.

Underlying connections to a remote site of up to 100 KM can be enabled by use of Fibre Channel technology such as long wave length lasers (LWL), extended long wave length lasers (ELWL), or wave division multiplexing (WDM). Use of wide area network technology such as IP (internet protocol) / Fibre Channel bridges or SONET / Fibre Channel bridges enables the remote site to be separated from the primary site by thousands of kilometers.

Performance over long distance links can vary for multiple reasons. The number of buffer-to-buffer credits determines the number of Fibre Channel frames that a switch can transmit on a link at one time before requiring an acknowledgement back from the receiver. If there are not enough frames to fill the pipe, then performance degradation may result. Another performance factor on long distance links is the response time for SCSI transactions.

Chapter 11

SAN Availability

A computer system or application is only as available as the weakest link. To build a highly available computer system it is not sufficient to only have a highly available SAN. It is necessary to account for availability throughout the entire computer system: dual HBAs, multi-pathing software, highly available and multi-ported storage subsystems, and clustering software are some of the components that may make up such a system.

This chapter contains the following sections:

- Fabric Resiliency
- SAN Availability Classifications
- Redundant Fabric SANs
- SAN Topologies

Fabric Resiliency

Many fabric topologies are available that provide at least two internal fabric routes between all switches that comprise that fabric. These topologies are considered resilient because each topology can withstand a switch or ISL failure while the remaining switches and overall fabric remain operational. This self-healing capability is enabled by the Brocade-authored Fabric Shortest Path First (FSPF) protocol. While originally a Brocade-only protocol, the standards bodies have accepted FSPF as the standard protocol for Fibre Channel fabric routing.

Figure 11-1 depicts the failure of a switch in a Cascade topology. Switches A and B are unable to communicate with the remaining switches when the switch marked with the "X" fails, resulting in the fabric segmenting into three separate fabrics.

However, a switch failure in a Core/Edge, or other resilient topology fabric does not cause a loss of communication with the remaining switches, as shown in Figure 11-2. If switch B fails, switch A can still communicate with switch C through the alternate path indicated by the arrows. The fail over to alternate paths is effectively transparent to the attached devices. This fail over is performed by FSPF, which automatically reroutes the data around the failure.

Figure 11-1. In a Cascaded Topology there is no resilience to Fabric failures.

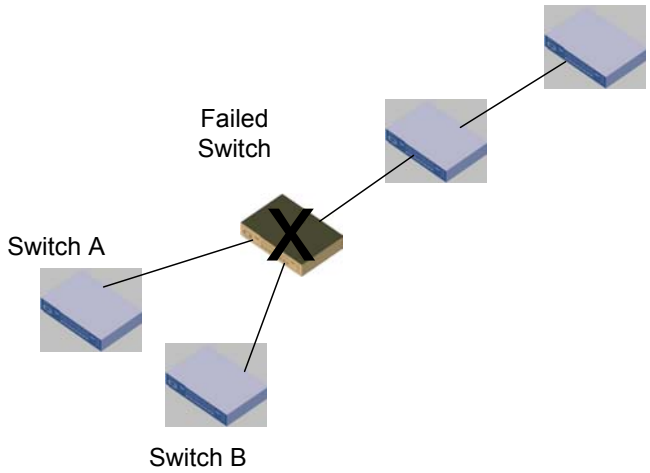

Figure 11-2. In a Core/Edge Topology there is resilience to Fabric failures.

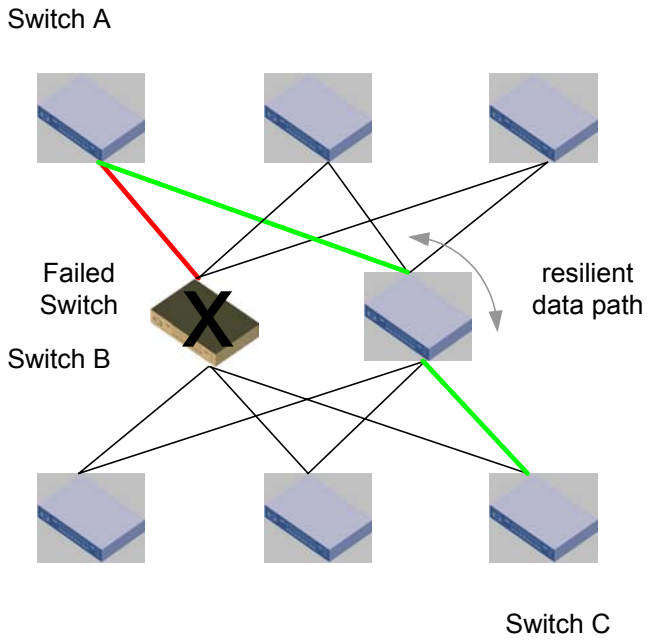

SAN Availability Classifications

Devices attached to a fabric may require highly reliable access to support applications such as storage consolidation, server clustering, high availability, or business continuance operations. Both resilient and non-resilient dual fabrics can be referred to as "redundant fabric SANs." Redundant designs are always recommended for high availability systems, and any large SAN deployment where downtime for the entire SAN could affect hundreds of servers. There are four primary categories of availability in SAN architecture. In order of increasing availability, they are:

Single fabric, non-resilient: All switches are connected to form a single fabric, which contains at least one single point of failure. The Cascade topology is an example of this category of SAN (see Figure 11-1).

Single fabric, resilient: All switches are connected to form a single fabric, but there is no single point of failure that could cause the fabric to segment. Topologies such as core/edge, ring, and full mesh topologies are examples of single, resilient fabrics (see Figure 11-2).

Multi-fabric, non-resilient: In a dual fabric non-resilient SAN, half of the switches are connected to form one fabric, and the other half form a separate fabric. This model can be extended to more than two fabrics if desired. Within each fabric, at least one single point of failure exists. This design can be used in combination with dual-attached hosts and storage devices to keep a solution running even if one fabric fails, or if a rolling upgrade is needed. An example of this type of SAN is a dual fabric SAN built with core/edge fabrics — each fabric only containing a single core.

Multi-fabric, resilient: The most common multi-fabric SAN is the dual fabric SAN. In a dual fabric resilient SAN, half of the switches are connected to form one fabric, and the other half form a separate fabric. This model can be extended to more than two fabrics if desired. No fabric has a single point of failure that could cause the fabric to segment. This design can be used in combination with dual-attached hosts and storage devices to keep an application running even if one entire fabric fails due to operator error, catastrophe, or hardware/software issues. This is the best design approach for high-availability environments. Another key benefit of this design is the ability to take part of the SAN offline for online upgrades or maintenance without affecting production operations on the remaining fabric(s). An example of this type of SAN is a dual fabric SAN built with core/edge fabrics — each fabric containing two or more core switches. Figure 11-3.

Figure 11-3. In a Dual Fabric Core/Edge Topology there is resilience to Core fabric switch failures and Edge fabric switch failures.

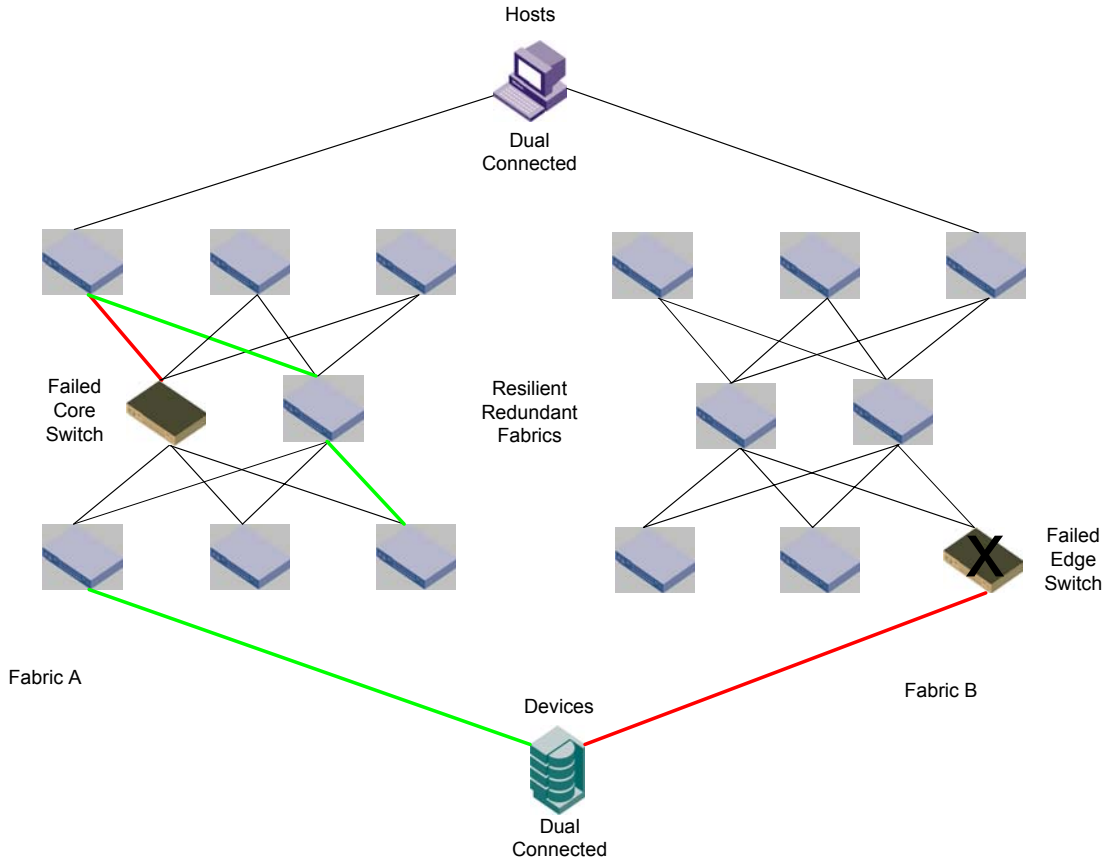

Redundant Fabric SANs

Resilient fabrics and the fault tolerant components that comprise them are very reliable. However, no single fabric can ever truly be a High Availability (HA) solution. The fabric itself is still potentially subject to failures caused by things like operator error, disaster, or hardware/software malfunctions. To account for those categories of error, another level of availability must be used: The redundant fabric SAN. This is sometimes known as a multi-fabric or dual-fabric SAN.

Redundancy in SAN design is the duplication of components up to and including the entire fabric to prevent the failure of the SAN solution. Using a fully redundant fabric makes it possible to have an entire fabric fail as a unit or be taken offline for maintenance without causing downtime for the attached nodes. When describing availability characteristics, what we are concerned with is *path* availability. If a particular link fails, but the path to the data is still there, no downtime is experienced by the users of the SAN. It is possible that a performance impact may occur, but this is a very small event compared to one or many crashed servers. Two or more fabrics must be used in conjunction with multiple HBAs, multiple RAID controllers, and path switch-over software to be effective for those SAN devices that require the highest availability. Figure 11-3 illustrates the ability of redundant fabrics to withstand large-scale failures. Note that tape drives can be effectively utilized in a redundant fabric environment as well – even if they are single attached.

In a redundant SAN architecture, there must be at least two *completely separate* fabrics – just as a high-availability server solution requires at least two completely separate servers. Duplicating components and providing switch-over software is well established as the most effective way to build high availability systems. Similarly, multi-fabric SAN architectures are the best way to achieve high availability in a SAN. In addition to enhancing availability, using redundant fabrics also enhances scalability. Using dual fabrics essentially doubles the maximum size of a SAN. If a fabric is limited by vendor support levels to 34 switches/1200 user ports and a single fabric solution with dual attach devices is utilized, then the SAN is limited to 1200 ports. Twelve hundred dual attach ports is equivalent to 600

devices. However, if a dual fabric with dual attach device solution is utilized, the SAN is capable of supporting 2400 user ports or 1200 devices. Any devices that are dual attached and are capable of supporting an active-active or active-passive dual-path essentially double the potential bandwidth. An active-active dual path means that I/O is capable of using both paths in normal operation. Some devices only support active-passive dual-pathing. With active-passive dual-pathing, the passive path is utilized only when the primary path fails. Some devices, such as tape drives, are not currently capable of supporting multiple paths. It is possible to address this issue by equally distributing tape devices between the redundant fabrics and configuring the backup applications to use an alternate tape drive should an outage on one of the fabrics occur. However, some elements of a tape backup solution, such as Robot control within a single library enclosure, do not currently map well into a redundant fabric environment.

Any single attached devices, such as a tape drive, non-critical storage and hosts can be single-attached, by alternately assigning them between the fabrics. When implementing a logical group of single-attached devices, such as a tape library with multiple tape drives and a robot, ensure that these devices reside on the same fabric, and if possible on the same switch. When deploying redundant fabrics, it is not always necessary to deploy symmetrical fabrics. For example, when using a dual fabric, the first fabric could consist of several interconnected SAN islands, while the second fabric consists of isolated islands,

The IBM eServer BladeCenter architecture provides all of the internal and external capabilities to connect to resilient redundant fabrics. It was designed to provide the maximum resilience to internal and external failures. The IBM eServer BladeCenter has the option to install 2 Brocade SAN Switch Modules into the Switch Module bays at the rear end of the chassis. Each of the Internal Host connections if installed with the Fibre Channel Expansion Cards have redundant connections to the Brocade SAN Switch Modules in bay 3 and in bay 4. The IBM eServer BladeCenter architecture provides dual connected hosts to redundant fabrics right inside the BladeCenter Chassis. The full benefit of this is to extend that redundant resilient internal architecture outside of the box when connecting the IBM eServer BladeCenter to an external SAN fabric.

Guideline — It is recommended to connect the Brocade SAN Switch Module in Bay 3 to one fabric and the Brocade SAN Switch Module in Bay 4 to a separate redundant fabric. This provides the maximum failure protection.

An example of the IBM eServer BladeCenter integrated into a resilient redundant Core/Edge Topology can be seen in Figure 11-4.

Figure 11-4. Brocade SAN Switch Module's in a Dual Fabric Core/Edge Topology.

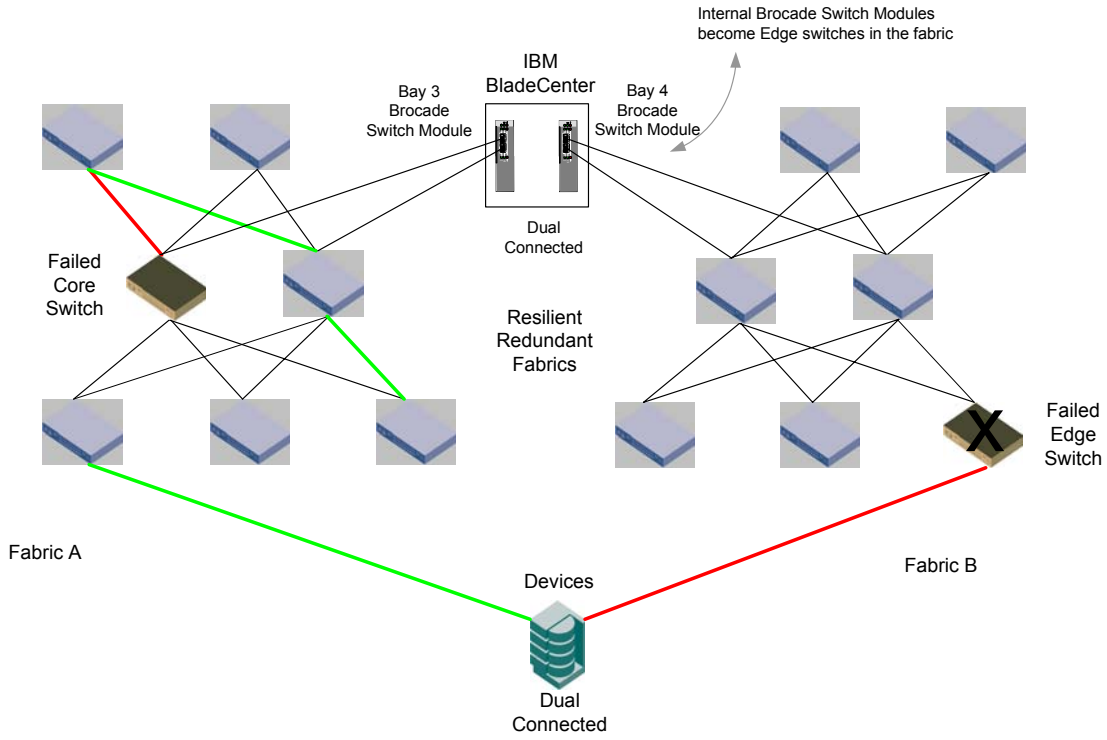

SAN Topologies

The *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0* covers the aspects of fabric topology, such as scalability, performance, and availability in great depth. The flexible fabric architecture of Brocade switches allows arbitrarily complex fabrics to be built when it is necessary to solve complex problems, but also allows simple solutions to be built to solve simple problems.

The following simple topologies are discussed in more detail in that guide:

- Cascade
- Ring
- Full Mesh
- Partial Mesh
- Core/Edge
- Composite Core/Edge

Chapter 12

SAN Scalability

The scalability of a SAN is the size to which that SAN could be expanded without fundamental restructuring. Scalability is so important to SAN design that it is frequently the first criteria used in deciding how to approach the SAN architecture. The designer starts by asking two questions: 1) "how many ports does the SAN need now?" and 2) "how many ports *will the SAN* need in the near future?" The solution is then designed to meet the port count requirements.

SANs should be designed to scale to the largest size that they could reasonably be expected to need to achieve in a reasonable time frame, rather than merely using the requirements at the time of implementation as a target. This will prevent the SAN from being "painted into a corner" and needing to be fundamentally restructured after entering production.

Investment protection is another area that relates to scalability. If an existing switch is replaced with a newer or higher port count switch to increase scalability, it is valuable to reuse the existing switch elsewhere in the fabric. Proper initial planning facilitates this as well.

The core/edge fabric topology is the most frequently deployed topology in cases where scalability needs are great. It is derived from the star topology, which is common in traditional data networks. The core/edge fabric topology (see Figure 11-2) is a similar design, except that the core is normally redundant, and there is typically only one level of edge switches. Some core/edge implementations opt for a single core switch when deploying the fabric in dual fabric SAN architecture. The logic behind this approach is that should a single core fail, the remaining and redundant fabric can maintain SAN operations.

A core/edge topology is scalable from many perspectives. It is possible to use variable size switches in the cores and the edges. The larger the core switch, the larger the fabric can grow. If large cores and edges are utilized, it is possible to build very large fabrics. This concept of scaling a fabric by using variable port count switches is shown in Figure 12-1.

Figure 12-1. Core/Edge topologies showing a starter configuration with 16-port core switches and larger much more scalable configuration with 128-port core switches.

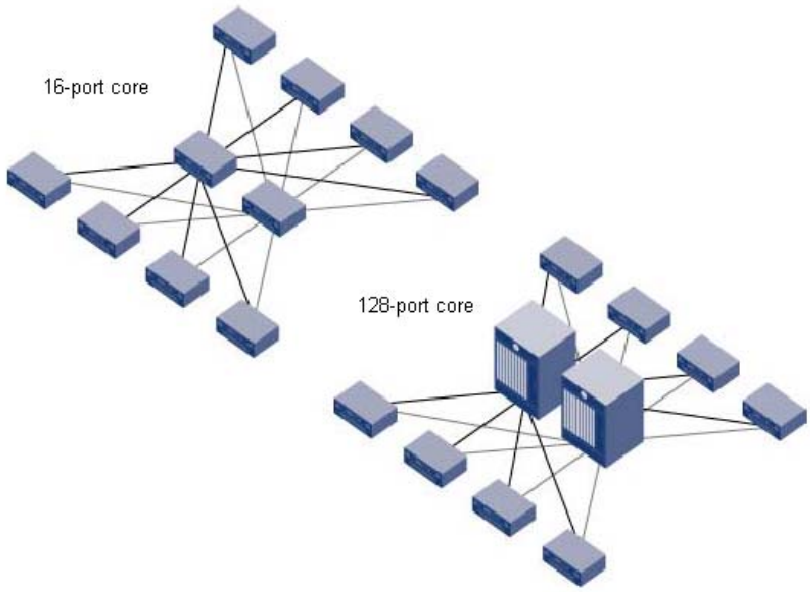

A reasonable network progression might start with 16-port or 32-port core switches and migrate to 64-port or 128-port cores when the scalability limits of the smaller cores are reached. See the *Brocade SAN Migration Guide* for detailed information on how such a migration would be performed. As shown in Figure 12-2, when additional ports are required, the lower port count switches in the core can be replaced with the higher density 64-port or 128-port switches. The lower port count switches can then be redeployed at the edge.

Figure 12-2. Upgrading the core to higher port count switches and redeploying the former core switches to the edge.

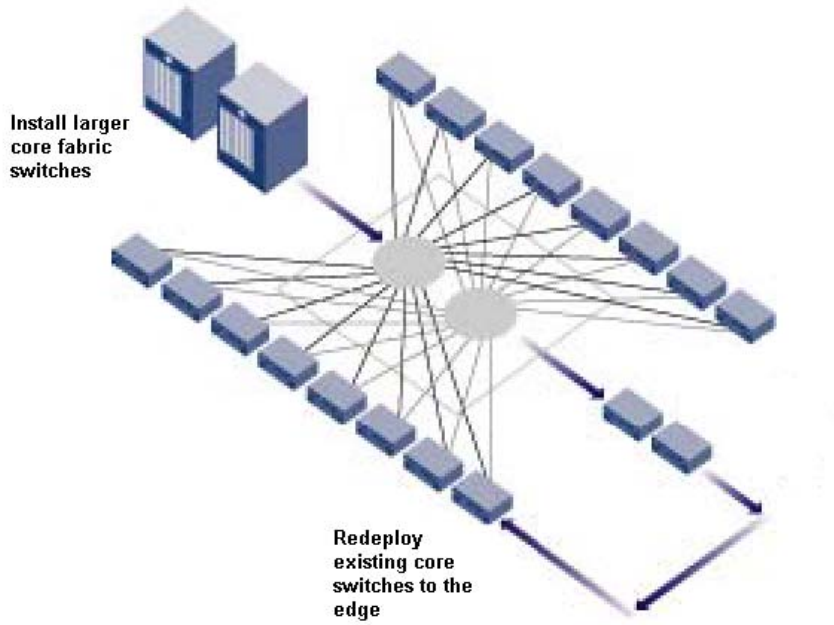

In these diagrams the embedded Brocade SAN Switch Modules can be any of the Edge switches.

Note — The Brocade SAN Switch Module should be included in larger network fabrics as Edge switches and not as Core switches simply because there are not enough empty ports for ISLs to connect other switches to.

The ultimate limitation in a fabric design today, and as defined in the Fibre Channel standards, is a maximum of 239 physical switches, be they 8, 16, 64 or 128 port versions. As a practical matter, no vendor has yet tested networks of this size due to the expense and complexity of implementing such a network. The current practical switch-count limit is fewer than 239 switches, based upon empirical testing.

To determine whether or not a SAN is supported, it will be necessary to work with your switch provider to determine if your SAN design is valid. Important variables that determine the supportability of a particular SAN include the number of switches, version of Fabric OS, the topology, number of ISLs/trunks, number of connected devices, and hop count.

For very large SAN Design configurations please supply your switch provider with the following information:

- The maximum number of switches
- The total number of ports
- The maximum number of SAN devices
- The minimum version of Fabric OS implemented.
- The fabric hop count

Note — The Brocade SAN Switch Module counts as a switch when calculating maximum number of switches in a fabric.

Chapter 13

SAN Performance

This chapter contains the following sections:

- SAN Performance
- IO Profiles

An *over-subscribed* link is one on which multiple devices ***might*** contend for bandwidth. A *congested* link is one on which multiple devices ***actually are*** contending for bandwidth. Traditional data networks have been built with very high levels of over-subscription on links for years. The Internet is probably the best-known example of this.

While not capable of supporting Internet-like over-subscription ratios, real-world SANs can be expected to have several characteristics that enable them to function well, even with over-subscribed links. These characteristics include bursty traffic, shared resources, low peak usage by devices, good locality, and devices that can generate only a small fraction of the I/O as compared to the available bandwidth. Most networks have all of these characteristics to some degree. Moreover, organizations can often realize substantial cost savings by deliberately designing a SAN with a certain amount of over-subscription.

When performance service levels are critical and the bandwidth requirements are high, lower over-subscription levels or traffic localization should be targeted.

Today, many devices attached to a SAN are not capable of generating traffic at the full Fibre Channel bandwidth of 100 MB/sec or 200 MB/sec. Figure 13-1 and Figure 13-2 detail a simplified scenario using a handful of devices to explain SAN bandwidth consumption.

Figure 13-1. Low server utilization.

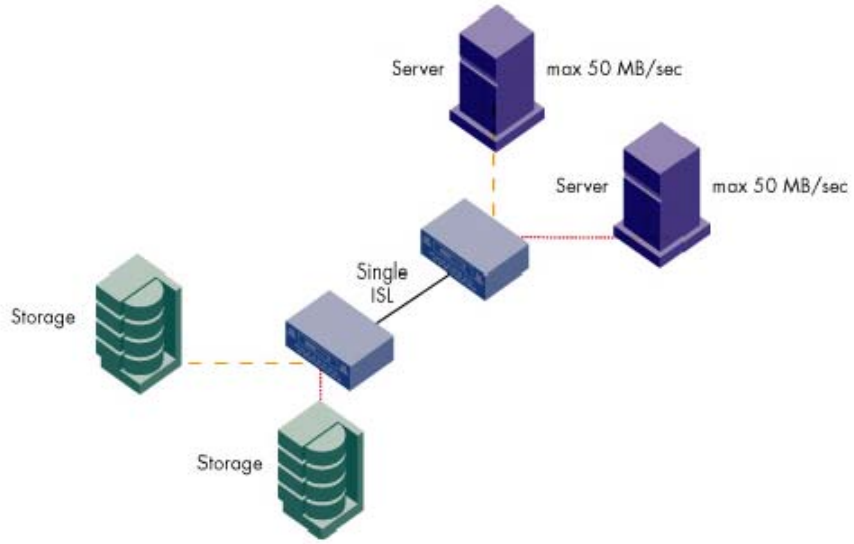

Note that in the ISL in Figure 13-1, the total amount of traffic that is intended to cross between switches never exceeds the 200 MByte/sec capacity of the link (assuming a 2 Gbit/sec ISL). Even if the servers in Figure 13-1 are running at their theoretical maximum performance, there still might be performance bottlenecks with the storage devices. In Figure 13-2, the two servers are accessing a single storage port, so the 2:1 fan-out of the storage port becomes the limiting factor.

Figure 13-2. High-bandwidth consumption with 2 ISLs.

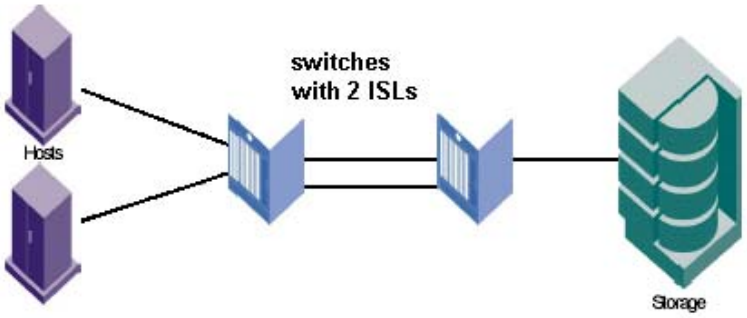

The key to managing bandwidth is capturing or estimating performance requirements and matching these requirements to an appropriately designed SAN. If the servers are capable of generating 400 MB/sec of traffic and there are two 2 Gbit/sec ISLs (see Figure 13-2), the network routes the traffic over both of them, and congestion does not occur. The storage port can operate at only 50 MB/sec; therefore, each server can average only 25 MB/sec. This scenario is common in storage consolidation environments where many servers need to share a single storage port. However, the I/O requirements for most servers can be surprisingly low (1 or 2 MB/sec) and a single storage port can sustain many hosts without overwhelming its I/O capability.

I/O Profiles

Understanding an application's I/O requirements is essential to the SAN design process. An individual I/O can be classified as either a read or a write operation. Although I/O is usually a mixture of reads and writes some applications are strongly biased. For example, video server I/O activity is normally almost 100 percent reads, while video editing cluster I/O may be mostly writes. Tape I/O is primarily write oriented – unless doing a restore.

I/O can further be classified as random or sequential. Examples of random I/O include an e-mail server or an OLTP server. Sequential I/O is characteristic of decision support (such as data warehousing), scientific modeling applications, or backup applications.

The third characteristic of I/O is size, which typically ranges from 2 KB to over 1 MB. Typically, user file systems have smaller I/O sizes, whereas video servers or backups may have very large sizes. The *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0* provides more detail on the application I/O profiles that establish the typical magnitude of application bandwidth consumption.

For SAN design performance purposes, I/O is classified by bandwidth utilization: *light*, *medium*, and *heavy*. It is very important to support test assumptions by gathering actual data when possible. You can gauge the type of I/O activity in your existing environment by using I/O measurement tools such as **iostat** and **sar** (UNIX) or **diskperf** (Microsoft).

Chapter 14

ISL Trunking

This chapter contains the following sections:

- ISL Trunking
- ISL Over-subscription ratios
- Recommended ISL Over-subscription ratios

Brocade ISL Trunking is a feature that enables traffic to be optimally shared across available inter-switch links (ISLs) while preserving in-order delivery. A trunk group logically joins two, three, or four ISLs into one logical ISL (on the Brocade SAN Switch Module the limit is two ISLs in one trunk group, as there are only two external ports on a Brocade SAN Switch Module). Use of trunking can minimize or eliminate congestion in the SAN because trunking optimizes ISL utilization. The use of trunking minimizes the effort of managing a SAN since ISLs are now managed as a group instead of individually and optimizes FSPF performance as FSPF does not have to compute as many routes.

The ISLs of the Brocade SAN Switch Module in the IBM eServer BladeCenter can be trunked and will provide additional performance for the 14 internal Server Blades of the IBM eServer BladeCenter. It is recommended when using trunking to incorporate the use of this feature into a dual redundant fabric configuration. This provides the best availability of the fabric if failures should occur. An example of the use of Trunking in a dual redundant fabric configuration can be seen in Figure 14-1. This diagram expands upon the concepts learned in the previous discussions on fabric resilience and redundancy.

Figure 14-1. Utilizing ISL Trunking in Brocade SAN Switch Module in a Dual Fabric Core/Edge Topology to increase performance.

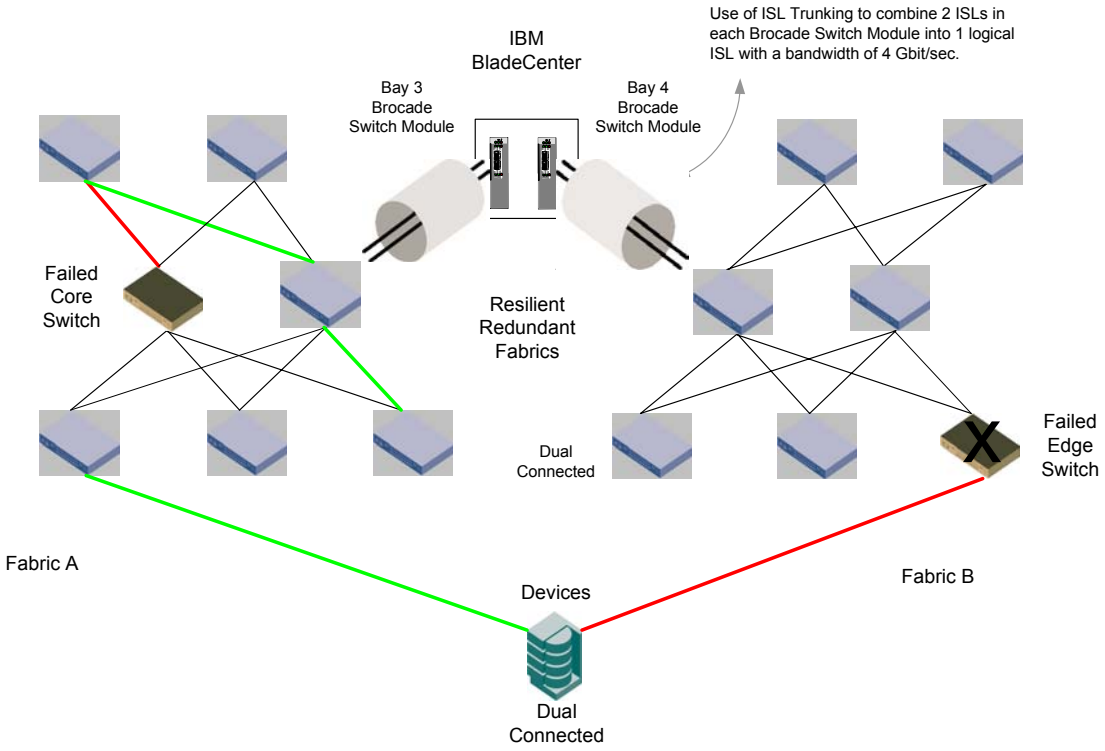

ISL Over-subscription Ratios

When designing a SAN, it is important to understand the performance boundaries such as storage fan-out ratios and storage performance. While any SAN device that connects to a SAN at 2 Gbit/sec is theoretically capable of 2 Gbit/sec, in reality, that device is most likely capable of a much lower performance. If a device truly is capable of generating 2 Gbit/sec of I/O, then the principles of locality should be applied or sufficient bandwidth should be provisioned for the ISLs. A very popular SAN application is storage consolidation, where many hosts share a storage device or port.

Several popular storage vendors target an average of a 6:1 fan-out. This means that on average six hosts are sharing a single storage port. If there were 32 storage ports in a fabric, then one would expect to find an average of 192 hosts. Even if every host requires 1 Gbit/sec or 2 Gbit/sec of bandwidth, the storage devices in the fabric are only capable of delivering 32 Gbit/sec (1Gbit/sec ports) or 64 Gbit/sec (2 Gbit/sec ports). This equates to 3-6 MB/sec per host. While some ports in the fabric may require maximal bandwidth, not all ports require sustained maximal bandwidth and rarely, if ever, do these ports require maximal bandwidth simultaneously.

When all ports operate at the same speed, ISL over-subscription is the ratio of device, or data input ports that might drive I/O between switches to the number of ISLs over which the traffic could cross.

For the IBM eServer BladeCenter architecture there can be a maximum of 14 internal Server connections that can go over a maximum of 2 ISLs per Brocade SAN Switch Module. The maximum over-subscription ratio on the Brocade SAN Switch Module is then 14 host device ports to 2 ISLs. Which can be abbreviated to14:2. This can further be reduced as a fraction to a 7:1 over-subscription ratio.

The basic over-subscription formula is “ISL Over-Subscription = Number of Nodes: Number of ISLs”, or Io=Nn:Ni. This is reduced as a fraction so that Ni=1.

If all of these hosts tried to simultaneously use the ISLs at full speed in a sustained manner — even if the hosts were accessing different storage devices — each would receive only about one-seventh of the potential bandwidth available. Figure 14-2 shows the ISL Over-subscription ratio for the IBM eServer BladeCenter architecture.

Figure 14-2. IBM eServer BladeCenter over-subscription ratio.

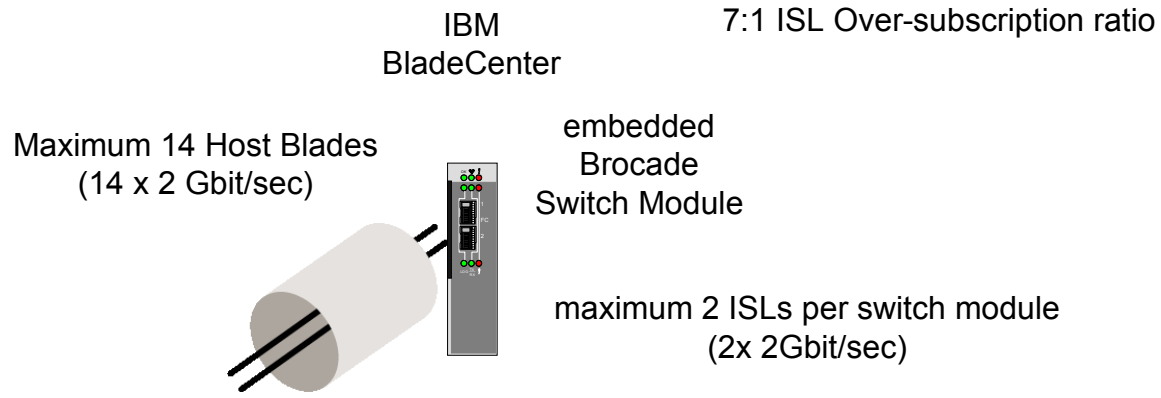

Recommended ISL Over-subscription Ratios

ISL over-subscription ratios apply in practice to Core/Edge fabrics. The calculations for ISL over-subscription ratios for a Core/Edge fabric are simple and straightforward, while these same calculations become more complex and less pertinent for other topologies, such as a ring topology. ISL over-subscription ratios principally apply to edge switches, as the role of a core switch is to connect other switches. Connecting devices to a core switch is supported and makes sense for particular scenarios, such as when there are excess ports available on the core switch or for performance

purposes. When devices are connected to the core switch, the number of ISLs/trunks is usually equal to or greater than the number of devices and the devices actually are under-subscribed or at a 1:1 ISL over-subscription ratio – minimizing the value of this metric for core switches. The ISL over-subscription ratio does become more meaningful for a core switch when devices are connected to the core and there are more devices than ISLs/trunks.

A 7:1 ISL over-subscription ratio is aligned with an industry average of 6:1 fan-out. The trend in the storage industry is that the hosts to storage ratios are increasing, as is the performance of storage devices. A 7:1 ISL over subscription ratio should be targeted in SAN designs, with the ISL over-subscription ratio being adjusted higher or lower to meet particular performance requirements.

The higher the ISL over-subscription ratio, the lower the performance and conversely, the lower the ISL over-subscription ratio, the higher the performance. An ISL over-subscription ratio of 3:1 results in high performance and fewer available ports while an ISL over-subscription ratio of 15:1 results in lower potential performance and more available user ports. The practical boundaries for ISL over-subscription ratios is 3:1 for high performance SANs and 15:1 where lower performance is sufficient. The factors that influence this position include resiliency, number of ports available on a Brocade switch, and industry host to storage fan-out ratio.

The IBM eServer BladeCenter architecture approaches the best of both worlds by having an ISL over-subscription ratio that lies in the mean of performance and available host connections.

In summary the ISL attachment strategies for the Brocade SAN Switch Module can be characterized into the following four choices :

- If you are connecting to a single switch. ISL Trunk for performance. Figure 14-3.
- If you are connecting to a single bladed Director you have two options. Figure 14-4:
    - Split the ISLs across blades for availability
    - Trunk the ISLs for performance
- If you are attaching to a fabric with a single core switch. ISL Trunk for performance. Figure 14-5
- If you are attaching to a fabric with dual core switches you have two options. Figure 14-6 :
    - Split the ISLs for availability
    - Trunk the ISLs for performance

Figure 14-3. Connecting to a single switch.

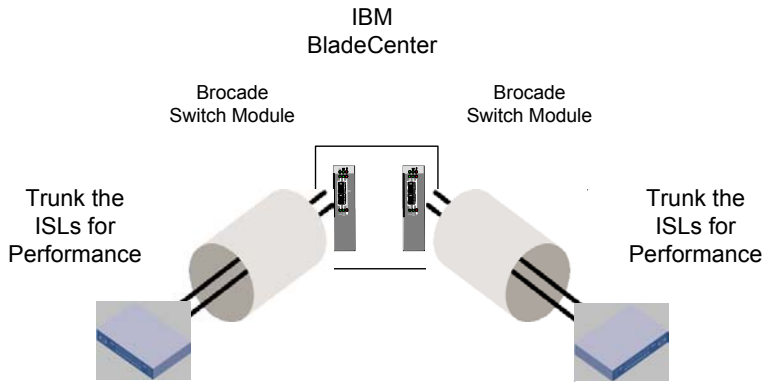

Figure 14-4. Connecting to a single Director.

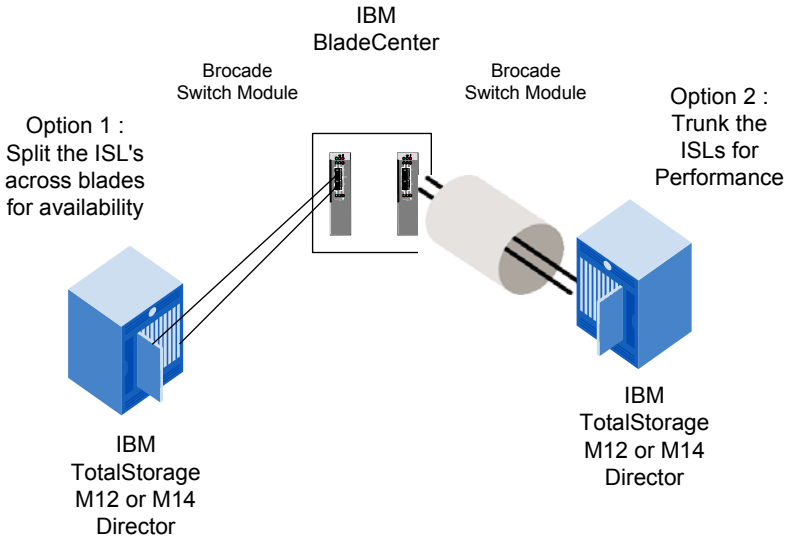

Figure 14-5. Connecting to a fabric with a single core.

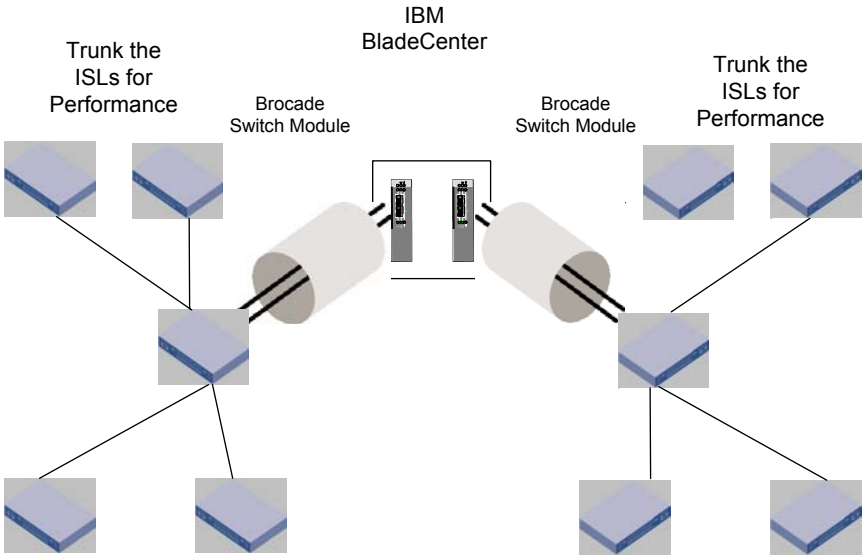

Figure 14-6. Connecting to a fabric with a dual core.

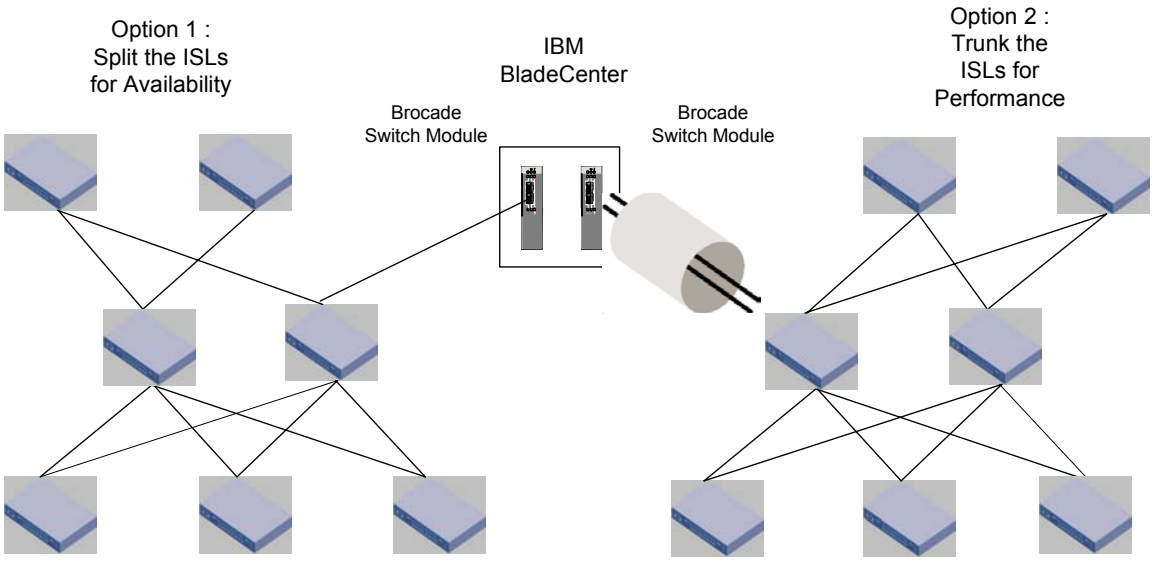

Chapter 15

Architecting SANs With SilkWorm Switches

This chapter contains the following sections:

- Architecting SANs with SilkWorm Switches
- Device Attachment Strategies
- Brocade SAN Switch Module Design Considerations

Many SAN related elements such as device attachment strategies, platform specific topics, switch location in the fabric, Zoning, Advanced Security (Secure Fabric OS), Extended Fabrics, and supportability have some bearing on the design of a SAN. It is these more in depth topics and the usage of SilkWorm switches in a SAN architecture that are discussed in this section.

The Core/Edge topology is identified as a reference topology for the guidelines presented in this chapter. The Core/Edge topology is preferred for scalable, available, and high performance fabrics for a number of reasons. With a Core/Edge topology, it is easy to satisfy SAN functional requirements. Given a diverse set of requirements: performance, locality, connectivity, and scalability, the Core/Edge topology provides the most flexible architecture to address these overall requirements.

The Core/Edge fabric is a variation on the well-established "star" topology popular in Ethernet LANs. There are a few differences. Because Fibre Channel uses a routing protocol (i.e. FSPF) with load sharing, Fibre Channel fabrics can take full advantage of multiple Core switches. In an Ethernet network, multiple switches at the center of a star would usually act in an active/passive backup relationship, using a Spanning Tree Protocol or some variation.

These differences make multi-core fabrics very popular, since it is possible to easily scale the fabric's bandwidth by adding Core elements. In addition, the requirements of a Core fabric switch are more stringent than those of the center switch in an Ethernet star. Due to the properties of Fibre Channel, the acceptable performance and reliability characteristics are very high.

The introduction of trunking further increases the effectiveness of a Core/Edge fabric due to more efficient utilization of the ISLs and lessened management requirements. Some Core/Edge implementations opt for a single Core switch when deploying the fabric in a dual-fabric SAN architecture (see Figure 15-1). The logic behind this approach is that should a single Core fail, the second fabric in the SAN can maintain operations. In a resilient Core/Edge fabric two or more switches reside in the center of the fabric (the Core) and interconnect a number of other switches (the edge).

Figure 15-1. Single Core/Edge fabric vs. Dual Core/Edge fabric.

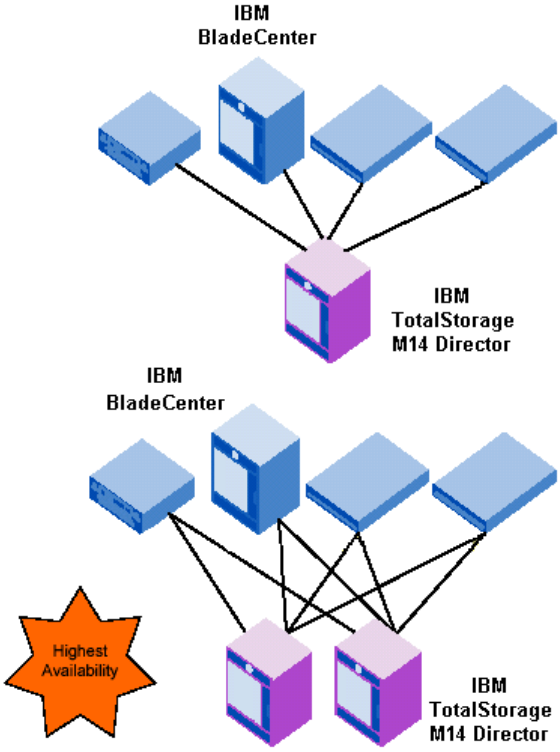

Switches that reside in the middle of the fabric are referred to as core switches. The switches that are interconnected by the core switches are referred to as edge switches. The simple form of the core/edge fabric has two or more core elements, each of which consists of a single switch. In a simple core, the core switches do not connect with each other. Edge switches in a core/edge fabric also do not connect to each other. They only connect to the core switch.

Devices such as hosts and storage are attached to free ports on the edge switches. These ports are referred to as edge ports or user ports. Free ports on the core switches should usually be reserved for additional edge switches when using 16-port switches and can connect SAN devices for higher port count switches. The scalability of a core/edge fabric is reduced when a device is attached to a core.

Note

A Core/Edge fabric is typically built with two or more core switches, but can be built with a single core switch if that fabric is used in a dual fabric SAN -- the redundant fabric maintains the SAN availability should the single core switch become unavailable.

A Core/Edge topology can be built with a variety of switch platforms, such as the SilkWorm 2000 series, 3016, 3200, 3250, 3800, 3850, 3900, 12000, and 24000. The type of switch in a fabric does have some bearing on the practical as well as supported size of a fabric. A SilkWorm 24000 is used as the core in Figure 15-1, as the 128-ports per domain support enables the size of the fabric to grow to several thousand ports by connecting edge switches. The edge can be built with a variety of switch platforms. The IBM eServer BladeCenter is a perfect fit for an Edge switch in this type of configuration.

A key benefit of the core/edge topology is the use of FSPF, which automatically distributes the load across all paths equally. In fact, all edge-to-edge paths are equal in a true core/edge topology. There are two or more paths between any two edge switches in a resilient core/edge topology. Because of this, core/edge fabrics have very good performance under varying to zero locality conditions.

Additional benefits of the Core/Edge topology:

- Well-tested and reliable
- Widely deployed in production environments
- Simple and easy to understand
- Able to solve most design problems, fits well with many SAN solutions, and is an effective choice when design requirements are not well known
- Easy to grow without downtime or disconnection of links and devices
- Pay as you grow
- Flexible
- Capable of exhibiting stellar performance, with full utilization of FSPF load sharing and resiliency features
- Conducive to performance analysis. Because the Core/Edge topology is symmetrical, it is a straightforward process to identify performance issues. Every device has an equivalent path to any other device and the same available bandwidth between any two devices. To identify a SAN performance issue it is only necessary to monitor the core switches. With other topologies, this is not the case.
- The potential to scale to thousands of ports (using high port count switches)

The guidelines established in this chapter can also apply to other topologies, such as a full mesh or ring. It is up to the reader to interpret the guidelines in this section for topologies other than Core/Edge. Regardless of topology chosen, a redundant fabric (i.e. dual fabric) SAN is recommended.

Device Attachment Strategies

The *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0* covers device attachment strategies in great depth. The IBM eServer BladeCenter architecture simplifies the device connection strategy of the Host connections to the fabric. How switches connect to other switches and how devices connect to those switches significantly influences the performance and availability of a SAN. Easy to understand and consistent device attachment strategies also simplify the operation and maintenance of a SAN. For more detail on other device connections strategies refer to the above-mentioned DDM Guide. These include:

- Trunk and ISL Connections
- Edge Switch ISL/Trunk Connections
- Core Switch and Standalone ISL/Trunk and Device Connections
- Attaching SAN Devices for Availability
- Connecting Devices to the Core
- Low Locality Device Attachment
- Platform Specific Design Considerations
    - SilkWorm 2000 Series, 32x0 and 38x0 Switches
    - SilkWorm 24000, 12000 and 3900

Brocade SAN Switch Module Design Considerations

Since its inception, Brocade has been the leader in developing technology and testing practices to expand the limits of Fibre Channel fabrics. As fabrics increase in size, the numbers of switches, inter-switch links (ISLs), and edge devices increase rapidly. This in turn increases the demand on the fundamental computing resources of the Control Processor in each switch, as it must rapidly complete tasks such as processing Zoning configuration updates distributed by other switches, analyzing and distributing RSCNs, responding to Name Server queries from hosts logging into the switch, etc. Customers planning to build very large fabrics (approx. more than 1000 ports) should plan on implementing them solely with the Brocade 2 Gbit/sec switch family as there are limitations in the memory and processor power of the legacy SilkWorm 1 Gbit/sec switches.

The Brocade SAN Switch Module utilizes a different and more powerful control processor than the legacy SilkWorm 2000 series, 3200 and 3800 switches. Additionally, the Brocade SAN Switch Module is configured with more memory than the SilkWorm 2000 and 3200/3800 switches. These additional resources position the Brocade SAN Switch Module to function well in larger fabrics.

Hot code activation (HCA) is one of the features in Fabric OS 4.1 and later versions (including v4.2.1). The defining characteristic of hot code activation is that while a new firmware image is being activated on a switch there is no disruption of end-to-end data flow between the hosts and storage devices. No disruption means no dropped frames, no retries, and no time-outs. The ASICs on the switch continue to process frames while the new firmware is being activated. Hosts that are logged into their targets will never be aware that anything has happened.

The Brocade SAN Switch Module, with only one Control Processor, must shutdown and reboot Fabric OS as part of the HCA process. This process completes in less than 60 seconds, with the critical fabric services available in about 50 seconds. Again, although there is a window when fabric services are unavailable, there is no disruption to end-to-end data flow between SAN devices.

Because it does take longer to do hot code activation on the Brocade SAN Switch Module as compared to the SilkWorm 24000 or SilkWorm 12000 (which both have redundant CPs), the switches directly linked to the Brocade SAN Switch Module need to be tolerant of the 60 seconds when no fabric services from the Brocade SAN Switch Module, will be seen. Brocade has made the necessary modifications to Fabric OS v3.1.0 and v2.6.1 and later versions to extend the time-out values so that link and fabric re-configuration is avoided. For this reason, it is strongly recommended that customers deploy neighboring switches (i.e. immediately connected via an E-port) to the Brocade SAN Switch Module with Fabric OS v2.6.1 (or later) on the SilkWorm 2000s series, Fabric OS v3.1.0 (or later) on the SilkWorm 3200/3800, or Fabric OS v4.1.0 (or later) on the SilkWorm 3250/3850/3900/12000/24000.

During the Brocade SAN Switch Module hot code activation, if earlier releases (Fabric OS v2.6.0x, Fabric OS v3.0.2x, or v4.0.x) are deployed on the neighboring switches to a Brocade SAN Switch Module, a time-out will occur on these neighboring switches resulting in a fabric re-configuration.

The Brocade SAN Switch Module HCA procedure should always be performed when the fabric is stable. However, recovery actions will be taken upon completion of the HCA to ensure that no RSCNs, or fabric configuration changes are missed. If any new hosts or targets were added to the Brocade SAN Switch Module switch during the HCA reboot time, then the initial FLOGI will time out. After the reboot, the switch will reset that port to cause the FLOGI to happen again. If any device is added elsewhere in the fabric, the RSCN will be delivered after the reboot completes. Finally, if any E-port cables are pulled, or the fabric rebuilds for any reason, then after the HCA reboot completes, the Brocade SAN Switch Module will cause the fabric to rebuild again so that it may participate in the fabric.

Guideline	To prevent unnecessary disruptions in the fabric when firmware is activated on a Brocade SAN Switch Module, it is recommended that any switches directly connected to the Brocade SAN Switch Module use Fabric OS v2.6.1, v3.1.0, or v4.1.0 (or subsequent versions).

Chapter 16

Zoning Design

This chapter contains the following sections:

- Zoning Design Considerations and Guidelines
- Zoning and Scalability
- Zoning Database Size

Zoning Design Considerations & Guidelines

Zoning is an important element of a secure and healthy SAN. Zoning does have an impact on SAN designs. For an overview of how zoning works refer to the *Brocade Zoning User's Guide* and the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0* these documents also provide guidelines for implementing zoning. This section highlights key elements of zoning that relate to a SAN design.

Zoning and Scalability

Zoning optimizes fabric services, such as RSCN distribution and name server response, and limits unnecessary device discovery. With the new zoning and related name server changes in Fabric OS 3.1 and 4.1 and subsequent releases, zoning becomes necessary for the proper functioning of large fabrics. For instance, the distribution of RSCNs (registered state change notifications) is reduced to only devices affected by a zone change. In prior releases of Fabric OS 3.x, 4.x, and all versions of Fabric OS 2.x, a zone activation (for example, executing the command **`cfgEnable`**) resulted in an RSCN being distributed to all devices – regardless of whether these devices were affected by a zone change. Additionally, not using zoning results in unnecessary delays during device discovery for some hosts, especially when a host pointlessly authenticates with hundreds of devices. These delays can last minutes, pause ongoing I/O, and cause unpredictable behavior on a host. Use of zoning on the switches limits the number of devices visible to a host and eliminates this host-based scalability problem.

Guideline — The implementation of zoning is recommended for any SAN and especially critical for any large fabric since zoning is fundamental to the functioning of multi-hundred port fabrics.

With Fabric OS v3.1/v4.1 and later, zoning changes cause different RSCN (Registered State Change Notification) behavior. In Fabric OS v3.1/v4.1 and later, when zone changes are enabled or disabled, fabric RSCNs are only sent to devices that completed an SCR (State Change Registration) and that are in the affected zones. In all Fabric OS v2.x releases, the locally connected devices that completed an SCR will receive these RSCNs, regardless if the device is affected by a zone change.

With a mixed fabric, the devices in the zones that are affected, as well as all devices local to the Fabric OS v2.x switches, receive an RSCN. The RSCN filtering of a device is handled by the Name Server of the switch to which it is attached. The Fabric OS of the switch that originates a zoning change is irrelevant.

Guideline — Make certain devices that require RCSN suppression are directly attached to a switch running Fabric OS version 3.1 or 4.1 or higher.

Zoning Database Size

Zoning consumes a finite amount of processing and memory resources. As the number of devices in a SAN grows, so do the demands on these same resources. The zoning implementation is optimized to minimize processing resources and leverage ASIC capabilities as much as possible. The zoning database size for SilkWorm 2000 series, 3200, and 3800 switches is 96 KB and 128 KB for SilkWorm 3016, 3250, 3850, 3900, 12000, and 24000 switches. To check the size of a zone database, use the command **`cfgSize`**. A switch with a zoning database size limit of 96 KB limits the size of the zoning database for the whole fabric – even if a SilkWorm 3016, 3250, 3850, 3900, 12000, or 24000 switch is present in the fabric. As the size of a SAN grows, it is important to monitor the zoning database. Typically, the zone database size needs to be of concern as the size of a SAN exceeds several hundred ports. The size of an alias name, zone name, or configuration name is limited to 64 characters for Fabric OS versions 2.6.1, 3.1, and Fabric OS 4.1 or later releases. While it is possible to create 64-character zone, alias, or configuration names, doing so consumes more memory than a shorter name. Additionally, shorter names are easier to remember and less prone to typing errors. Be wary of sacrificing meaning for shortness. See the Whitepaper *Zoning Implementation Strategies For Brocade SAN Fabrics* for effective guidance for naming aliases, zones, and configurations. The variable size of zone objects makes it very difficult to state guidelines as a number of zone entries or alias. A zone database size is similar to disk storage. The usage is not measured so much by how many files are located on the storage, but by the amount of space taken up by the files. It is recommended to review the zoning configuration of a fabric periodically for unused alias, zoning, and configuration entries and to then delete these unnecessary entries. Unnecessary alias, zone, and configuration entries frequently result from the merging of fabrics or the addition of a switch with predefined zones into an existing fabric.

Note	The maximum zoning database size for SilkWorm 2000 series, 3200, and 3800 switches is 96 KB and 128 KB for the Brocade SAN Switch Module switch. A switch with a zoning database size limit of 96 KB limits the size of the zoning database for the whole fabric to 96 KB.
Guideline	Limit the name of an alias, zone or configuration to as few characters as possible while maintaining meaning of that name. Target 16-characters or less for an alias, zone, or configuration name
Guideline	For SANs that exceed several hundred ports, monitor the size of the fabric-zoning database with the CLI command **`cfgSize`**
Guideline	Routinely review a zoning configuration to identify unused aliases, zones, and configurations and then remove these unused entries

Chapter 17

Security Design

Due to the complex nature of Security this Guide cannot address all of the design implications that must be addressed in an Advanced Security (Secure Fabric OS) environment. For comprehensive information on how security works refer to the *Brocade Secure Fabric OS User's Guide* and the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0,* these documents provide the necessary guidelines for implementing security.

Note

A secured fabric must be *entirely* secured and all switches in a secured fabric must run a version of Fabric OS that supports security and these switches must be licensed to run security.

Note

The following minimum version requirements apply for fabrics that need secure mode enabled:

- Fabric OS version 2.6.2 (or higher) with the Security license - for secure fabrics containing a SilkWorm 2000 series switch and any of the following SilkWorm switches: 3016, 3200, 3250, 3800, 3850, 3900, 12000, 24000.
- SilkWorm 3200 and 3800 switches - Secure Fabric OS version 3.1 (or higher) with the Security license
- SilkWorm 3900 and 12000 switches - Secure Fabric OS version 4.1 (or higher) with the Security license
- SilkWorm 3250, 3850, and 24000 switches - Secure Fabric OS version 4.2 (or higher) with the Security license
- Brocade SAN Switch Module - Advanced Security (Secure Fabric OS) version 4.2.1 (or higher) with the Security license

Note

The Brocade SAN Switch Module switch has a different default username than "admin," which exists on all other SilkWorm switch products. Due to this change, a new command, **`userrename`**, must be used to rename the default "USERID" user account to "admin" before connecting the Brocade SAN Switch Module to a secure fabric made up of other Brocade SilkWorm switches.

SAN Deployment

This section discusses fundamental SAN Deployment concepts and associated guidelines and checklists for deploying an effective SAN design.

This section contains the following chapters:

- SAN Deployment Overview
- Planning
- Staging
- Validation
- Maintenance

Chapter 18

SAN Deployment Overview

Once the SAN is designed using sound principles as detailed in the *SAN Design* section, it needs to be deployed. The process of SAN deployment is more than plugging in cables, turning on the power and setting IP addresses. In fact, there are four distinct phases all SAN deployments go through. Briefly, these four phases are planning, staging, validation and maintenance. The definition and benefits of doing each of them properly are summarized below:

1. **Planning -** Preparing for the staging of the SAN and related equipment is important; therefore, a chapter is dedicated to this subject. Proper planning allows for estimation of time and effort and provides justification for resources. A good plan provides a means of measuring progress and greatly assists in avoiding potential pitfalls that prevent timely project completion.

2. **Staging -** After the planning phase the SAN needs to be put together. Staging covers everything from uncrating and racking the switches to configuring Brocade Fabric OS and the applications that will run on the hosts, storage and other devices that are attached to the SAN.

3. **Validation** - Once staged, the entire SAN configuration needs to be tested and validated to confirm it is ready for production. The tests should verify device connections, check for the SAN robustness, and most importantly, test the application availability under varying failure sceneries.

4. **Maintenance -** Once the SAN transitions to an operational state, changes are likely to occur, such as the addition of hosts or storage. This may require more switches if all user ports are allocated. The Fabric OS or other firmware and software may need to be upgraded. Maintaining the SAN is all about the day-to-day activities that keep the SAN running smoothly and efficiently.

Like any complex project, there are many different ways of deploying a SAN. Even though there is a level of complexity, there are still general guidelines and practices that should be followed. Many guidelines in this Section are in the form of checklists. Checklists provide the essential SAN deployment information to ease the transition from the planning stage to production. All checklists can be used as is, or modified to fit specific SAN environments. Taken as a whole, the guidelines and practices discussed within each checklist will allow for sound decisions; optimizing the IT investment once in production. The checklists are not meant to be all-inclusive. Other methods of performing the activities discussed are possible and encouraged.

It is crucial that the production SAN be supportable, simple to maintain, and easily scaled. Proper planning and documentation is critical to make these goals attainable. Effective up front planning simplifies the actual staging phase. Node location will be known, storage requirements met and cabling simplified. Proper documentation functions much like a map, allowing the quick identification of devices and configuration settings. This reduces the potential downtime, whether it is scheduled or not. Since these benefits, and others, are so important, there is emphasis on putting together an effective planning strategy.

Chapter 19

Planning

Having an effective plan prior to the staging of the equipment is critical for overall SAN deployment success. This success is measured in many forms. The greatest benefit of a good plan is that it gets the SAN deployed on time and within budget. Doing this ensures the ROI to be realized in the shortest time possible.

Planning is all about understanding the requirements and allocating necessary resources. Proper planning does take extra up-front effort and cost. For higher port-count SANs, it is critical that there be at least one person who is accountable for developing and driving it. With a good plan, progress towards completion can be measured, the right persons identified, the roles and responsibilities defined, the site and SAN documented, and SAN resources can be efficiently utilized. With proper planning, less effort is needed to maintain the infrastructure. When the staging requirements are known, obstacles that may impede progress can be avoided.

This chapter provides some guidelines on what essential information is needed for putting an effective plan together. Checklists are used throughout this chapter to provide a framework for gathering requirements to meet specific site needs.

This chapter contains the following sections:

- Planning Documentation
- Zone Planning
- Security Planning
- Extended Fabrics Planning
- Fabric OS Upgrade Planning
- LAN Planning

Planning Documentation

This section will provide some guidelines regarding documenting the SAN. Knowing what documentation is needed allows for planning IP addresses, switch domain numbers, what ports should be used for ISLs, etc. Once created, having the documentation readily available allows change management of all components. In addition, being organized with updated documentation saves time and effort when referencing equipment for service calls. Even while not being serviced, when information is needed, it can be referenced quickly and easily. Here is a checklist that provides a recommendation as to what documentation should be created before and during the staging of the equipment. For templates to assist with your planning refer to the Appendices of the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0.*

Documentation Checklist

1. Get an Equipment Binder
2. Logical Design Diagram
3. Switch Spreadsheet
4. ISL Port Map
5. Device Spreadsheet
6. Label All Cables
7. SAN Verification Test Plan Requirements
8. SAN Verification Test Plan

Guideline	Set unique domain numbers for each switch in the SAN. This allows for simpler merging of fabrics.

Guideline	As a convention, consider setting the domain ID of each switch to the last octet of its IP address. Be aware that the highest allowed domain number is 239.

Zone Planning

Zoning allows the hosts to access specific storage devices on the SAN. For those SANs with multiple OS platforms, zoning allows for OS separation and co-existence. With no zoning defined on the SAN, any device can see any other device. This is the default setting. Once zoning is in place, all devices must be members of a defined zone. Those devices that are not will be blind to all others. This section will provide some guidelines as to zoning plan definition. For additional information reference the *Brocade Zoning User's Guide.* More detail is provided in the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0*.

Zoning requires careful thought and planning. Armed with the documentation created in the previous section, and understanding the requirements, allows the creation of a zoning plan. Creativity is important here as there is no one "correct" zoning configuration for a given SAN fabric configuration. In general, follow any specific zoning recommendations provided by the switch vendor.

Zoning Plan Checklist

1. Gather the list of host and storage devices to be zoned from the device spreadsheet.

2. Define the storage requirements for each host based upon software application requirements.

3. Adhere to recommended storage device configurations such as LUN masking, LUN security, and other specific features supported by the vendor.

4. Consider specific host requirements for storage value-added feature sets such as Server Free backup, LUN snapshots, or LUN mirroring over distance.

Some key points when planning for zoning:

- Clearly understand the storage requirements for each host. This means understanding specifically what storage is needed for the software application that will be in production. To understand this, the number and size of LUN presentations on each storage array Fibre Channel port must be clearly defined.
- Be sure to adhere to recommended storage configurations by the switch vendor for LUN masking and other storage specific features. Some vendors may recommend using a separate HBA for tape devices. For this case, define the zoning configuration to fence off tape devices from any other HBAs, which see the disk storage, in the same host.
- Keep in mind the different OS platforms, backup application requirements and the number of paths to each LUN, which may drive the zoning plan. There may be specific host requirements for storage value-added feature sets such as Server Free backup, LUN snapshots, or LUN mirroring over distance.
- As a general rule, have overlapping zones in all cases. An overlapping zone has the HBAs share one or more storage ports, but with the HBAs separate from each other. This is sometimes referred to as a single initiator zoning.

Guideline	For a large number of zone configurations, generally over 15, use Fabric Manager or Web Tools. These tools vastly simplify the zoning implementation. WWNs of devices and the ports on the switch they are attached to are also seen automatically in the GUI, so validation of connectivity can be done simultaneously
Guideline	If possible, use persistent binding on the host. This will provide consistent controller, target and LUN numbers for each storage LUN. Backup applications are especially sensitive, as these numbers map directly to the backup application device identities.
Guideline	If using World Wide Names (WWN), zone by port World Wide Name rather than node World Wide Name. This is because a port World Wide Name uniquely identifies a port to which a target is attached. Some Multi-pathing software may get confused and not be able to discover targets properly. This is especially true when using multi-port HBAs.
Guideline	Be aware of mixing different HBA vendors in a single zone. Each vendor HBA responds differently to RSCNs, a method to notify an HBA for device discovery, and may cause one of the HBAs to lose the zoned device.
Guideline	Use single initiator zones that have one HBA per zone.
Guideline	Separate HBAs from each other for clustered hosts. Allow each HBA to see the same storage but not each other. Once again, RSCNs, may cause the clustered host HBA to lose the storage array.

Security Planning

Due to the complex nature of Security this Guide cannot address all of the deployment implications that must be addressed in an Advanced Security (Secure Fabric OS) environment. For comprehensive information on how security works refer to the *Brocade Secure Fabric OS User's Guide* and the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0* these documents provide the necessary guidelines for implementing security.

A few details are covered here to help you get started.

- Security Measures
- Advanced Security (Secure Fabric OS)

Security Measures

There are some SAN Security measures that should be in place before implementing Advanced Security (Secure Fabric OS). Here are some guidelines in the form of a checklist to assist with the planning process. These steps can be taken to provide some initial restrictions on accessing the SAN and to provide some control over change management. For

maximum SAN Security, these measures should be used in conjunction with Advanced Security (Secure Fabric OS). Advanced Security (Secure Fabric OS) provides a single point of management and policies that allow complete control over what switches, devices and management stations are allowed to access the SAN. More detail is provided in the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0*.

Security Measures Checklist

1. Prevent Physical Access
2. Prevent Remote Access through IP security measures
3. Hard Zone the devices
4. Lock Down E-port creation with **`portCfgEport`**

Advanced Security (Secure Fabric OS)

Planning for SAN security and change management is important. Organizations understand that the data managed in their SAN environment is often highly sensitive and must have controlled access properly to ensure confidentiality, integrity, and availability. A compromise in any of these areas could have unintended consequences, resulting in the loss of proprietary information, capital, or other core business resources. Proper SAN Security planning with Advanced Security (Secure Fabric OS) mitigates these risks by ensuring proper SAN security access controls are in place and enforced. Because SAN security is, by its own right a separate subject, a comprehensive treatment will not be discussed in this document. In order to be effective in the implementation of Advanced Security (Secure Fabric OS), there are two assumptions that are made for the duration of the discussion: 1) Significant non-SFOS security measures are already in place and 2) proper security practices exist within the IT infrastructure.

Advanced Security (Secure Fabric OS) Preparation Checklist

1. Obtain and read the Advanced Security (Secure Fabric OS) documentation.
2. Be safe. Backup switch configurations with **configupload**. Prior information such as zoning will be wiped out when a switch or fabric is allowed to join an Advanced Security (Secure Fabric OS) enabled fabric.
3. Verify PKI Objects Exist (All Brocade SAN Switch Modules are shipped with the PKI objects from the factory). This is required for Advanced Security (Secure Fabric OS) implementation.
    - **pkishow** (Fabric OS 4.1)
    - **configshow "pki"** (Fabric OS 3.1/2.6.1)
4. If the PKI objects do not exist, obtain the PKICert tool to setup the fabric for secure mode. This is required for older switches. This utility runs on Windows and Solaris only. This can be obtained from the following IBM Web Site link the for Advanced Security (Secure Fabric OS) Upgrades :

   http://www-1.ibm.com/support/docview.wss?rs=0&context=HW200&context=SWJ00&q=2109%2Bssg1*&uid=ssg1S1001653&loc=en_US&cs=utf-8&lang=
5. Download and Install Brocade SecTelnet and Secure Shell client (SSH) Security Software Utilities
6. Verify Fabric OS Version. Update as required.
7. Install Security and Zoning Licenses on all switches in the SAN. This is required for Advanced Security (Secure Fabric OS).
8. Schedule downtime when enabling secure mode. A reboot of each fabric in the SAN is required as the firmware update is disruptive.
9. Recommended: If introducing the Brocade SAN Switch Module into an existing Fabric OS 2.x and 3.x environment consider changing the setting of the Core PID on all switches running Fabric OS 2.x and 3.x.

Caution	Before changing the Core PID parameter please read the recommendations and details for setting the Core PID format, in the *Fabric OS Procedures Guide* and consult the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0* for more information on the deployment of these settings.

Guideline	Setting the Core PID at the initial staging phase, on Fabric OS 2.x or 3.x based switches, will allow for a seamless introduction of a switch running Fabric OS 4.x into the SAN fabric.

Now that the SAN is prepared for Advanced Security (Secure Fabric OS) (SFOS), here are some guidelines for planning the implementation. Proper planning is essential to create the most secure environment possible. It is highly recommended that the checklist of activities in used as a basis for any specific plan, which should be followed rigorously.

Advanced Security (Secure Fabric OS) Implementation Checklist

1. Create a SFOS switch and device list for SFOS policies. This list should contain hostnames, switch names, IP addresses, and WWNs
2. Plan the FCS placements for each switch. Select a Primary and Backup FCS switch.
3. Covertly mark FCS Switches. Use a small physical mark so that FCS switches are easily located.
4. Determine policy requirements for each device and host.
5. Select the SFOS management hosts. Brocade SecTelnet and a Secure Shell (SSH) client may be required.
6. Perform a final review of all configuration selections. Verify all changes have been included.
7. (Optional) Disable Telnet Daemon on switches running Fabric OS 4.1or later.
8. Set the switch Recovery Password and Boot Password (Fabric OS 4.1or later).
9. Rename the admin-level ID “USERID” to the Brocade-specific default of "admin" and the user-level ID to the Brocade-specific default of “user” before enabling security. Otherwise, the switch will not be allowed in the secure fabric. This is covered in Chapter 8 of this manual.
10. Enable SFOS and verify its operation.
11. Backup Primary FCS switch configurations with **`configupload`**.

Warning	Each switch in the fabric needs to be rebooted to activate Advanced Security (Secure Fabric OS). Be aware that enabling Advanced Security (Secure Fabric OS) will reboot BOTH CPs simultaneously on the SilkWorm 12000 and 24000. It is recommended to schedule downtime on single fabric SANs before enabling Advanced Security (Secure Fabric OS).

Warning	The Primary FCS databases are not automatically backed up when secure mode is disabled. Data will be lost if the FCS switch is not backed up before the command **`secmodedisable`** is done on the Primary FCS switch. To backup the Advanced Security (Secure Fabric OS) data, use the command **`configupload`** on the Primary FCS.

Guideline	Consider a locked closet to physically secure the Primary and Backup FCS switches and/or the management station.
Note	SecTelnet or SSH* must be used to administer the Primary FCS switch when running Advanced Security (Secure Fabric OS). Brocade SecTelnet requires digital certificates are installed on each switch to be administered. This utility only encrypts the passwords sent over the LAN, all other commands etc., are sent as clear text.
Note	In a fabric that contains SilkWorm 2000 series switches, the maximum security DB size is limited to 32 KB, with 16 KB active. In a fabric containing Brocade SAN Switch Module switches, the security DB size maximum is 128 KB, with 64 KB active. For all fabrics, the maximum number of DCC policies is limited to 620 at this time.

Note : * Secure Shell (SSH) is a standards based secure method for accessing SilkWorm switches running Fabric OS 4.1 or later. Any SSH client that supports version 2 of the protocol can be used. There are literally hundreds of freeware SSH clients available that have this capability. On the switch side, SSH is only supported on Fabric OS 4.1 or higher. Two popular clients have been tested, Putty and Fsecure.

Extended Fabrics Planning

This section will provide the information for planning connections of Brocade SilkWorm fabrics over longer distances. One common reason is for data replication, which provides site redundancy. In this way, if one site goes down due to a disaster, the data can be recovered and brought online in minutes rather than days. Another typical use is consolidated remote data archival to tape.

This Guide cannot address all of the deployment implications that must be addressed for Extended Fabric environments. For comprehensive information on how Extended Fabrics works refer to the *Brocade Fabric OS Features Guide* and the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0* these documents provide the necessary guidelines for implementing Extended Fabrics. For additional information refer to Appendix B, *Long-distance Technologies for Storage Area Networks* in the DDM listed above.

A few details are covered here to help you get started.

SAN long distance connectivity may involve single mode fiber or through more sophisticated network equipment that allows for greater line availability. For Metro Area Networks (MAN), those distances up to 100 Km, DWDM equipment is generally used. Longer distances that go up to thousands of kilometers generally require FC protocol conversion. This means a Wide Area Network (WAN) transport method is required. Different equipment, such as an ATM long haul switch, may be required. These types of devices are out of scope for this document.

Configuring a port for Extended Fabrics can be performed using the **`portCfgLongDistance`** CLI command or through Web Tools. Specify the port and Extended Fabrics level as arguments to the **`portCfgLongDistance`** command.

Fabric OS v4.2.1 contains an additional optional parameter, “VC Translation Link Initialization”, to the **`portCfgLongDistance`** CLI command. When set to “1”, this parameter indicates that enhanced link reset protocol should be used on the port. The default value for this parameter is 0 and is compatible with earlier Fabric OS v3.0.x implementations. For optimal performance, specify “1” when E-Port links are between switches with Fabric OS v3.0.2 and greater, or Fabric OS v4.0.2 and greater. Specify “0”, or nothing, when connecting a switch with Fabric OS v3.0.2 or above switch to previous releases of Fabric OS.

If an Extended Fabrics port is to be configured on a SilkWorm 2000 series switch, the fabric-wide long distance parameter **`fabric.ops.mode.longDistance`** must be set to “1” in the configure CLI menu. Configuration of this parameter requires the switch to be disabled. This parameter must be set on all switches within the fabric. Also one

must set the appropriate Extended Fabrics mode for each long-distance port using the **portCfgLongDistance** CLI command.

When configuring Extended Fabrics on Brocade SAN Switch Module series switches, only port level configuration is necessary, and the **fabric.ops.mode.longDistance** must be set to "0", which is the default value.

Brocade supports Extended Fabrics ports between same series switches (i.e. SilkWorm 2000 series to SilkWorm 2000 series switches) as well as SilkWorm 3000 series switches to SilkWorm 12000 switches. Directly connecting a SilkWorm 2000 series switch to a SilkWorm 3000/12000 series switch is unsupported when using Extended Fabrics. For a mixed SilkWorm 2000 and SilkWorm 3000/12000 fabric, where the long-distance ports are between SilkWorm 2000 series switches, the fabric-wide parameter fabric.ops.mode.longDistance must be set to a value of 1 on all switches within the fabric. For mixed fabric configurations where long-distance ports are located between SilkWorm 3000 and/or SilkWorm 12000 series switches, the fabric-wide long distance parameter is not required.

Brocade Security features are supported on Extended Fabric links. This includes connections over dark fiber and DWDM networks.

Note — Please refer to Chapter 6 for more specific information on the limitations of Extended Fabric with the Brocade SAN Switch Module.

Fabric OS Upgrade Planning

This section will provide some high level guidelines when defining a strategy for performing a firmware upgrade of Fabric OS on an existing Brocade fabric. This activity typically happens in the maintenance phase of a SAN deployment. For the purposes of this discussion, it is assumed that no new switches will be added or removed from the existing SAN infrastructure. For that case, and more details, please see the *Brocade SAN Migration Guide.* For detailed instructions on all upgrades, refer to the *Brocade Fabric OS Procedures Guide* for the specific version of Fabric OS used on the switch. Be sure to upgrade to the firmware qualified or recommended by the switch provider.

Each Brocade-based SAN is unique. This is due to the wide variety of OS platforms, HBAs and storage arrays that may be attached. This uniqueness means that each upgrade needs to be carefully planned to minimize the risk of unscheduled downtime. Scheduled downtime is normally required for single fabric SANs. For dual fabric SANs, upgrade one fabric at a time.

Fabric OS 4.2.1, only supports the FTP protocol for firmware upgrades.

If you wish to use the CLI , the Brocade SAN Switch Module allows two admin telnet sessions. Use one for **firmwaredownload** and the other for **firmwaredownloadstatus**. **Firmwaredownloadstatus** is a handy command that shows a log of each upgrade phase. When complete, use **firmwareshow** to display the firmware version on each compact flash partition.

The Firmware download procedure can also be completed using the Web Tools and Fabric Manager GUI interfaces. Fabric Manager has the added advantage of being able to download to multiple switches concurrently.

Fabric OS Upgrade Planning Checklist

1. Analyze the potential risks and impact to each device on the SAN.
2. In order to maximize fall backwards capability, preserve each fabric switch configuration with **configupload**.
3. Use Fabric Manager for larger multi-fabric SANs.
4. Verify the upgrade version is supportable.
5. Gather the documentation and readme notes for the firmware release.
6. Schedule downtime for single fabric updates.
7. For dual fabrics, update one fabric at a time.

LAN Planning

There are a few guidelines to consider when attaching the fabric to the corporate LAN infrastructure. In general, it is highly recommended to configure a separate VLAN for each fabric. If at all possible, avoid the use of proxy servers from the SAN management stations outside the local subnet. In fact, it is recommended to use a management station on the same VLAN. For detailed guidelines on how to connect Brocade switches to the corporate network, please refer to the technical white paper titled: *LAN Guidelines for Brocade SilkWorm Switches.*

Note

When taking the LAN Guidelines for Brocade SilkWorm Switches into account it is important to note that all IP Ethernet communication to the Brocade SAN Switch Module is actually forwarded from the IBM eServer BladeCenter Management Module.

Chapter 20

Staging

Once the plan is complete and the resources readied, the new SAN can be built and prepared for production. Staging the SAN is more than uncrating, racking and installing the Brocade switches and other components. Staging also includes configuring the firmware and software of each component. Guidelines on how to accomplish these tasks are presented in this chapter in the form of checklists.

This chapter assumes that a new Core/Edge fabric with "clean" Brocade switches will be staged. A "clean" switch has no defined or active zoning configuration and has all default settings. It is assumed throughout this chapter that a Core/Edge topology has been chosen as part of the design criteria. The guidelines in this chapter may need to be modified for other topologies. This chapter will not explicitly cover an existing SAN fabric migration to other Brocade switch platforms. Please refer to the *SAN Migration Guide* for recommended guidelines and procedures.

Staging a new SAN Fabric requires essentially two tasks. The first task is uncrating, racking, cabling and providing power to the Brocade switches. The second task is configuring the Brocade Fabric Operating System firmware. The focus of this chapter will be on the commands and Web Tools views used in staging Brocade SAN fabrics with Fabric OS 4.2.1. New high-level troubleshooting functions are also introduced within these Fabric OS releases and will be discussed at a high level in this chapter.

This chapter will not provide recommendations on racking, cabling or power installation for the Brocade SilkWorm Fabric Switch Family. Guidelines for these tasks are provided in the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0.* Nor will this chapter focus on the devices, such as HBAs and storage targets, attached to the Brocade SAN Fabric.

Please see the *Brocade Fabric Manager User's Guide* and the *Brocade Advanced Web Tools Administrator's Guide* for detailed instructions on usage.

Note

For an online list of commands, use the *help* command. To view the online command reference, use the command help <command>. For example, **`help switchshow`** will provide the online reference for the **`switchshow`** command. For more information please refer to the *Brocade Fabric OS Reference Guide*.

This chapter contains the following sections:

- Preparing the Switches for the SAN Fabric
- Switch Staging Steps
- Preparing the Fabric Configuration
- Profiling the SAN

Preparing the Switches for the SAN Fabric

There are many possible ways to configure the switches that make up a fabric. Such as the configuration of a switch is dependant upon its role. During the planning phase of deployment, the following questions should be addressed:

- Will it be the principal switch?
- Is it a core or edge switch?
- Is long distance required?
- Are the correct license keys installed?

Once the answers are known, the goal of this section is to provide guidance on preparing Brocade Fabric OS for production in a multi-switch SAN fabric. There are two major steps. First, prepare each switch for attaching to the corporate LAN infrastructure and joining a fabric. The second step is to do the fabric wide configuration, this being primarily zoning. Each step will have a separate checklist. As with the rest of the document, this section first provides "bare bones" guidelines. Examples will be used throughout. Some guidelines will be considered optional and will be noted as such. All commands discussed in this section are available to the admin user.

Brocade Switch Preparation Checklist

Gather Planning Information (Switch Spreadsheet). Templates are provided in the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0*.

1. Set IP Address
2. Login and Change Passwords
3. Check Fabric OS version
4. Check Switch Status
5. Set the Date and Time
    a. Optional: Set the Time Zone
6. Set the Switch Name
7. Set the Domain ID Number
    a. Optional: Set the Core PID parameter on Fabric OS 2.x or 3.x switches if required
    b. Optional: Set the Extended Edge PID Format if required
    c. Optional: Set up a preferred principal switch in the fabric (Fabric OS 4.1 and above only)
8. Add Fabric OS licenses
9. Optional: Modify Port Configurations if required
    a. Optional: Setup ports for Extended Fabrics if required
10. Optional: Name devices with **`portName`**
11. Set Telnet Session Timeout Value
    a. Optional: Disabling the telnet daemon when secure mode is enabled (Fabric OS 4.1 only)
12. Optional: Customize Monitoring Features
    a. Switch Status Policy
    b. Track Changes
    c. SNMP Traps
    d. Fabric Watch
13. Check Environmental Status
14. Baseline and backup the switch configuration

Before getting started, gather the switch spreadsheet put together during the planning phase. This shows the planned IP addresses and domain names as well as the switch roles (Core Switch, Edge switch, Management Switch, etc.) Once the IP addresses and domain numbers are known, it is just a matter of executing the appropriate commands to set these values on the associated switch.

Switch Staging Steps

1. Set IP Address

The Ethernet IP Address, Ethernet Subnet Mask, and Gateway IP address should not be configured using local mechanisms on the Brocade SAN Switch Module, such as the **`ipaddrset`** CLI command or Advanced Web Tools.

The values must be configured using the IBM eServer BladeCenter Management Module, because all IP Ethernet access to the switch module itself is forwarded through the Management Module. If the switch module's IP address information is changed without changing the Management Module configuration, then telnet access to the switch may be lost.

This restriction does not apply to the Fibre Channel IP address and the Fibre Channel subnet mask (also referred to as the "in-band IP address" and "subnet mask"). These can still be configured using any of the standard switch management mechanisms.

In the IO Modules Task tab click on Management to manage the Brocade SAN Switch Module's IP Address from the IBM eServer BladeCenter Management Module. In this example the Current IP Address is 192.168.196.121. This example shows how to change that to 10.64.210.25. Follow Figures 20-1 to 20-4.

Note	By default the IBM eServer BladeCenter Management Module assigns the Brocade SAN Switch Module in I/O Bay 3 with 192.168.70.129 and the Switch Module in I/O Bay 4 with 192.168.70.130. [The default IP Address for the IBM eServer BladeCenter Management Module is 192.168.70.125.]

Figure 20-1. IP Address configuration using the IBM eServer BladeCenter Management Module. Step 1. View current configuration.

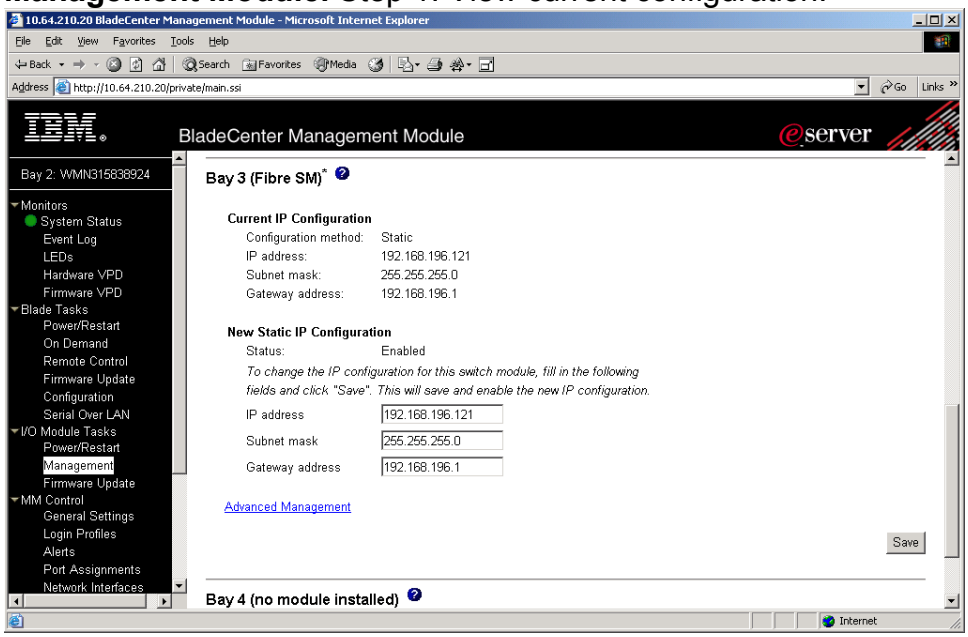

Figure 20-2. IP Address configuration using the IBM eServer BladeCenter Management Module. Step 2. Change the “new static IP address configuration”, then Click the “save” bottom on the bottom right.

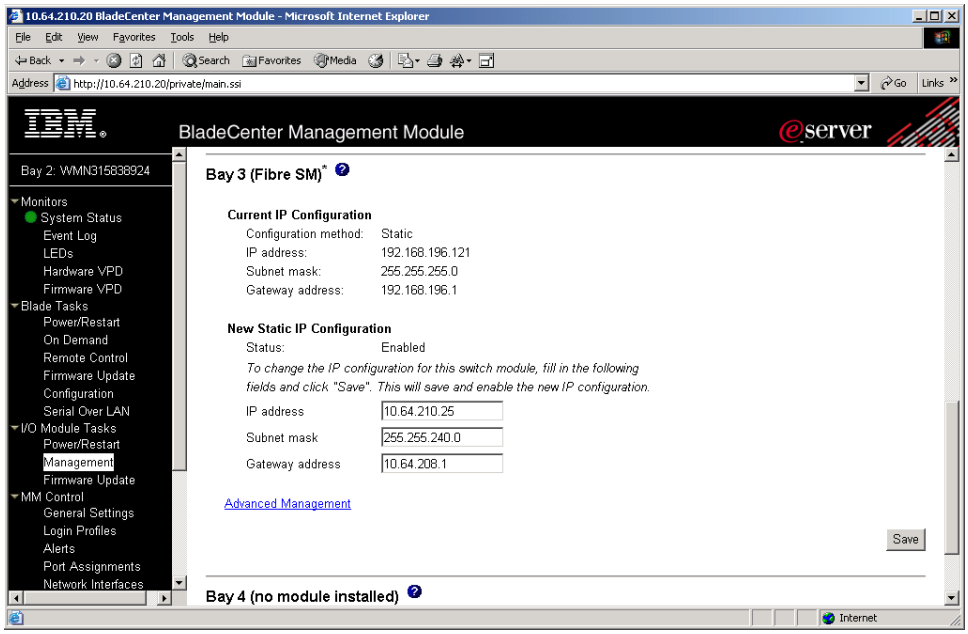

Figure 20-3. IP Address configuration using the IBM eServer BladeCenter Management Module. Step 3. Wait for the Management module to update the new information.

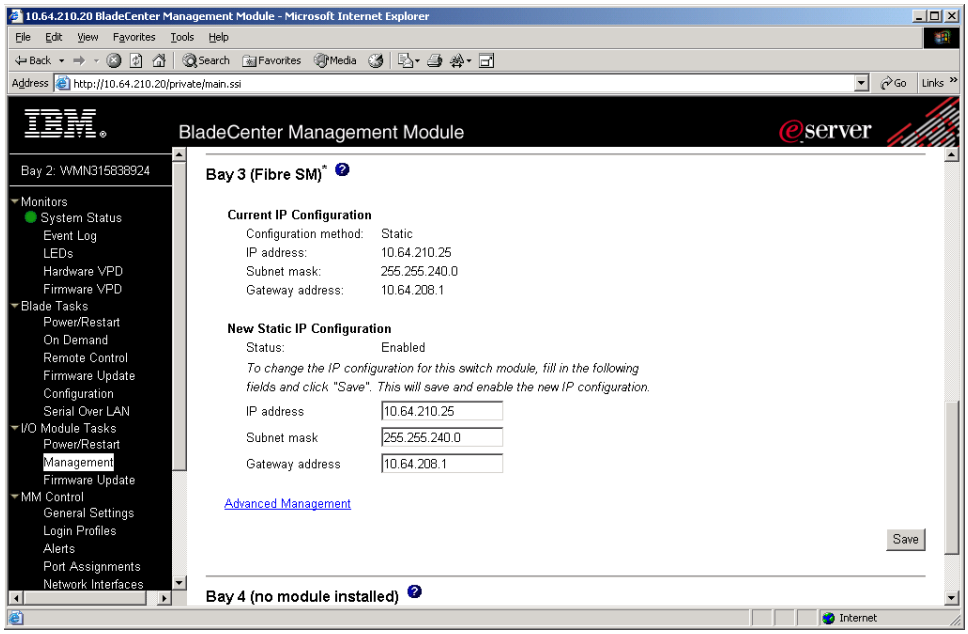

Figure 20-4. IP Address configuration using the IBM eServer BladeCenter Management Module. Step 4. If you wish to maintain this configuration across resets of the Switch Module then ensure the “Preserve new IP configuration on all resets” is “Enabled” for this Switch Module.

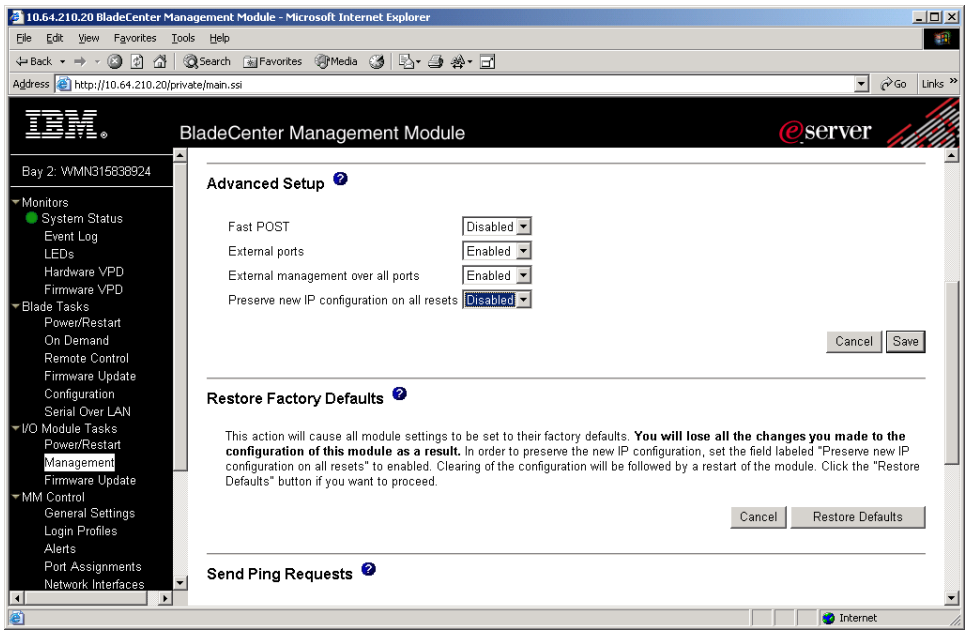

2. Login and Change Passwords

On the first telnet Login to the Brocade SAN Switch Module, the user will be challenged to change passwords for three of the four accounts on the system (root, factory and user). The four accounts are: root, factory, USERID and user. This step can be skipped, by using the CONTROL-C key combination, however the challenge will continue at every new login until all of the accounts passwords have been changed. Store the passwords in a safe location. See Figure 20-5 for an example.

Note — By default the Brocade SAN Switch Module admin user account login is “USERID” with a password of “PASSW0RD”. The “0” is the number zero and not the letter “O”.

Figure 20-5. Login change password Challenge.

```
brocadessm login: USERID
Password:
Please change your passwords now.
Use Control-C to exit or press 'Enter' key to proceed.

Warning:  Access to  the Root  and Factory accounts may be required  for
proper  support  of  the switch.  Please  ensure  the Root  and  Factory
passwords are  documented in a secure location.  Recovery of a lost Root
or Factory password will result in fabric downtime.

for user - root
Changing password for root
Enter new password:
Re-type new password:
passwd: all authentication tokens updated successfully
Please change your passwords now.
for user - factory
Changing password for factory
Enter new password:
Re-type new password:
passwd: all authentication tokens updated successfully
Please change your passwords now.
for user - user
Changing password for user
Enter new password:
Re-type new password:
passwd: all authentication tokens updated successfully
Saving passwords to stable storage.
Passwords saved to stable storage successfully
brocadessm:USERID>
```

If you forget your passwords the IBM eServer BladeCenter Management Module has the capability to reset the Brocade SAN Switch Module back to factory defaults. The “Restore Factory Defaults” option in the Advanced Management section of the I/O Module Management Task will reset the passwords back to the factory default (USERID/PASSW0RD) setting and allow the user to change the other 3 on the first login. Figure 20-6

Figure 20-6. Login change password Challenge.

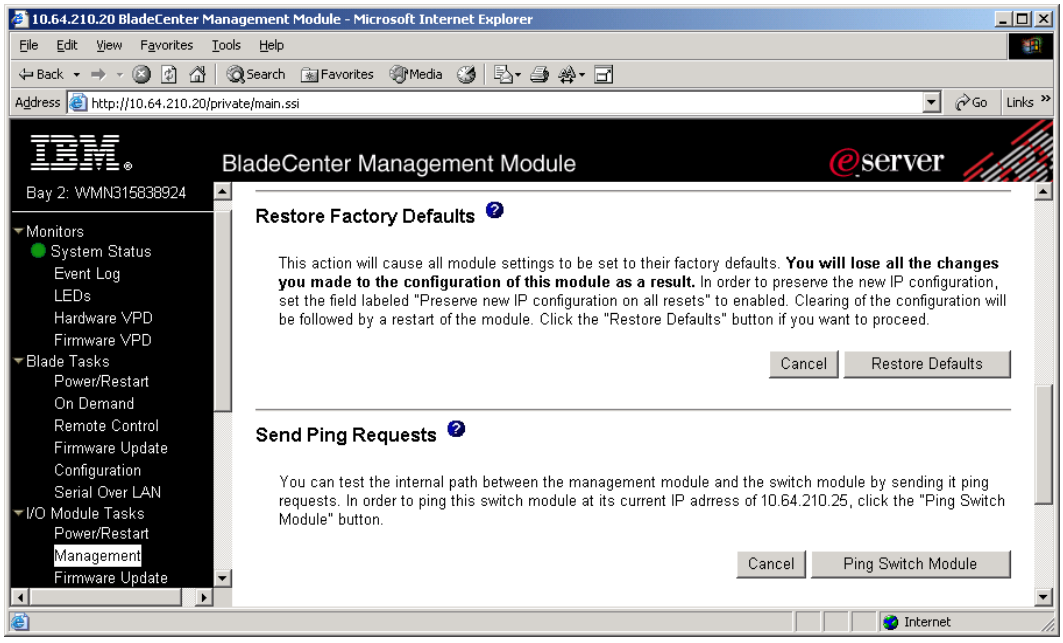

Caution — Setting the Brocade SAN Switch Module to factory defaults should be done with caution. If the Brocade SAN Switch Module is already part of an existing fabric this action may cause fabric disruption. Perform this during a maintenance window and ensure you have captured a configuration file using the configupload CLI command or from the Web Tools interface.

3. Check Fabric OS version

The recommended CLI commands to check the Fabric OS version for the Brocade SAN Switch Module are: **`version`** and **`firmwareshow`**. The **`firmwareshow`** command will provide additional information on the two banks of firmware stored in the compact flash for redundancy. The Fabric OS version can also be found in the Web Tools and Fabric Manager interfaces. See Figures 20-7 to 20-9.

Figure 20-7. Version and firmwareshow output.

```
brocadessm:USERID> version
Kernel:     2.4.19
Fabric OS:  v4.2.1
Made on:    Thu Apr 8 19:45:43 2004
Flash:      Tue Apr 13 19:07:48 2004
BootProm:   4.1.0
brocadessm:USERID>
brocadessm:USERID> firmwareshow
Primary partition:     v4.2.1
Secondary Partition:   v4.2.1
brocadessm:USERID>
```

Figure 20-8. Fabric OS Version in Web Tools.

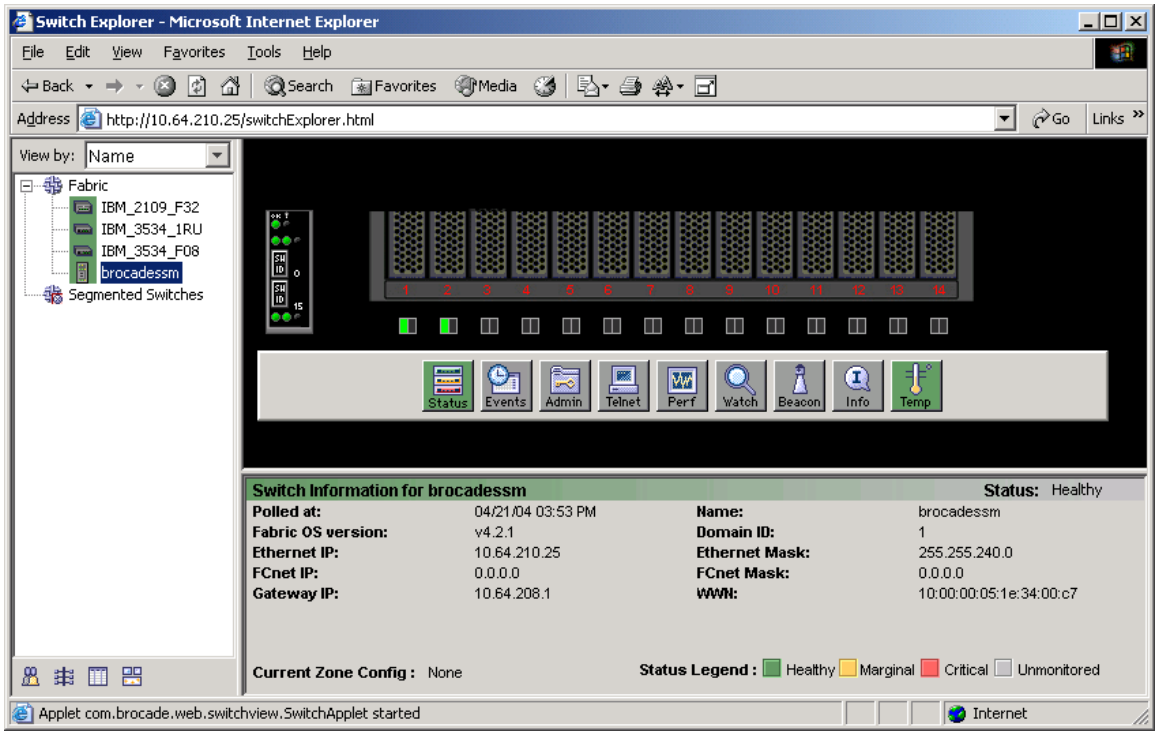

Figure 20-9. Fabric OS Version in Fabric Manager.

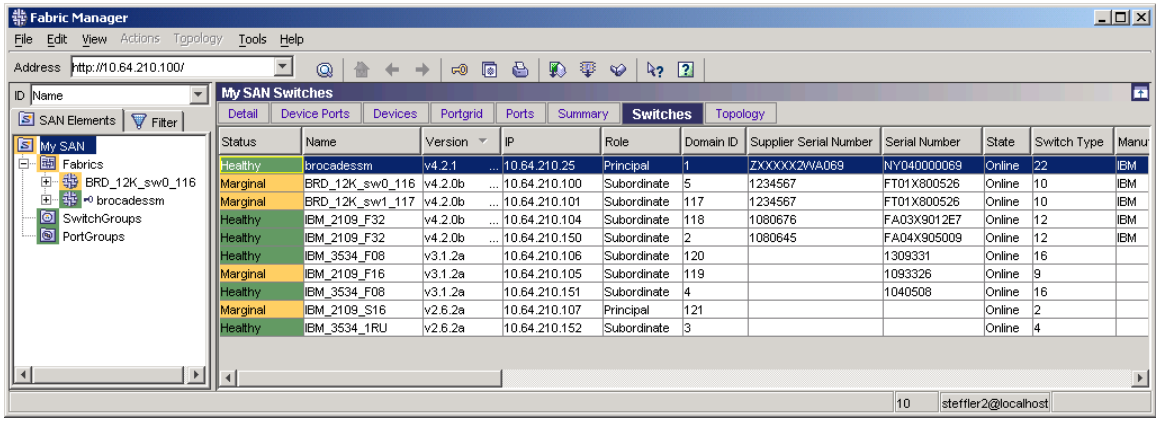

4. Check Switch Status

The recommended CLI command to check the Brocade SAN Switch Module status is: **`switchStatusShow`**. The Switch status can also be found in the Web Tools and Fabric Manager interfaces. See Figure 20-10 for the Web Tools view (Click on the Status button from the main Web Tools Window)

Figure 20-10. Switch Status in Web Tools.

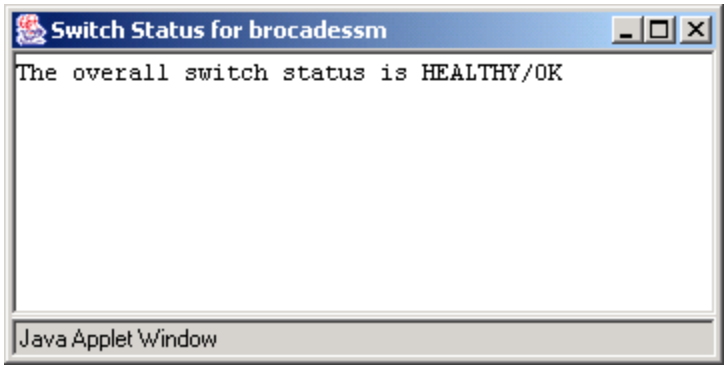

5. Set the Date and Time

To set the date and time use the **date** command. Figures 20-11 and 20-12 have examples that set the date and time using the CLI and Fabric Manager. Using **date** with no arguments provides the current day and time. Alternatively Fabric Manager can be used to push the current time on the server where Fabric Manager is installed onto a group of switches in a fabric. This can be found in the Actions Menu under "Set Time".

Figure 20-11. Command Line. Set Date and Time.

```
brocadessm:USERID> date "0331160204"
Sunday March 22 16:02:00 UTC 2004
brocadessm:USERID>
```

Figure 20-12. Fabric Manager. Set Date and Time.

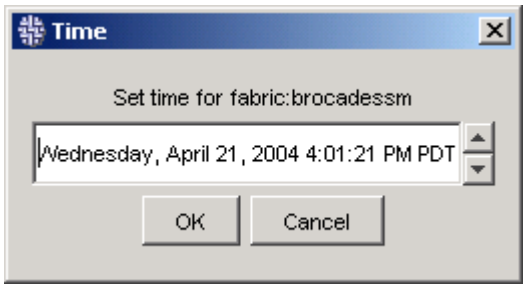

6. Set the Switch Name

Use **switchname** to set the name of the switch. Setting a name helps to quickly identify a switch and its role within the fabric. Figure 20-13 below sets the name of the switch to be FabA-C1-117, since it has been designated to be the first core switch in the fabric and the last octet of its IP Address is "117". (The default switch name for all Brocade SAN Switch Modules is "brocadessm".) Note the prompt changes to reflect the new switch name. You can also perform this task from Fabric Manager by right-clicking on the switch icon and choosing the "rename" option from the drop down list.

Figure 20-13. Set Switch Name.

```
brocadessm:USERID> switchname "FabA-C1-117"
FabA-C1-117:USERID>
```

7. Set the Domain ID Number

While it is not necessary, it is highly recommended to set a domain ID. If no domain is set, the switch will automatically derive a domain ID as part of a fabric rebuild. To automatically obtain a domain from the fabric, it is necessary that the switch connect to the fabric in a disabled state and then be enabled once the connection is complete. By setting the domain ID the domain will become deterministic on subsequent power cycles or reboots.

Setting the domain ID number of the switch applies to all versions of Fabric OS. The domain ID number is the first 8-bits of the 24-bit port ID (PID). The default domain number is 1 for all switches. The following steps provide a high level outline of the procedure, for more in depth information please refer to the *Fabric OS Procedures Guide*.

1. In the Switch Information tab of the Web Tools Admin interface click on the "Disable" button in the "Switch Status" field, and then click the "Apply" button. If using the CLI : **`switchdisable`**
2. Choose a new Domain ID (from 1 to 239) and place it in the Domain ID field of the "Name and ID" area and then click the "Apply" button. If using the CLI you will need to use the **`configure`** menu driven command.
3. In the Switch Information tab of the Web Tools Admin interface click on the "Enable" button in the "Switch Status" field, and then click the "Apply" button. If using the CLI : **`switch enable`**.

The Web Tools interface is shown in Figure 20-14

Figure 20-14. Domain ID configuration in Web Tools.

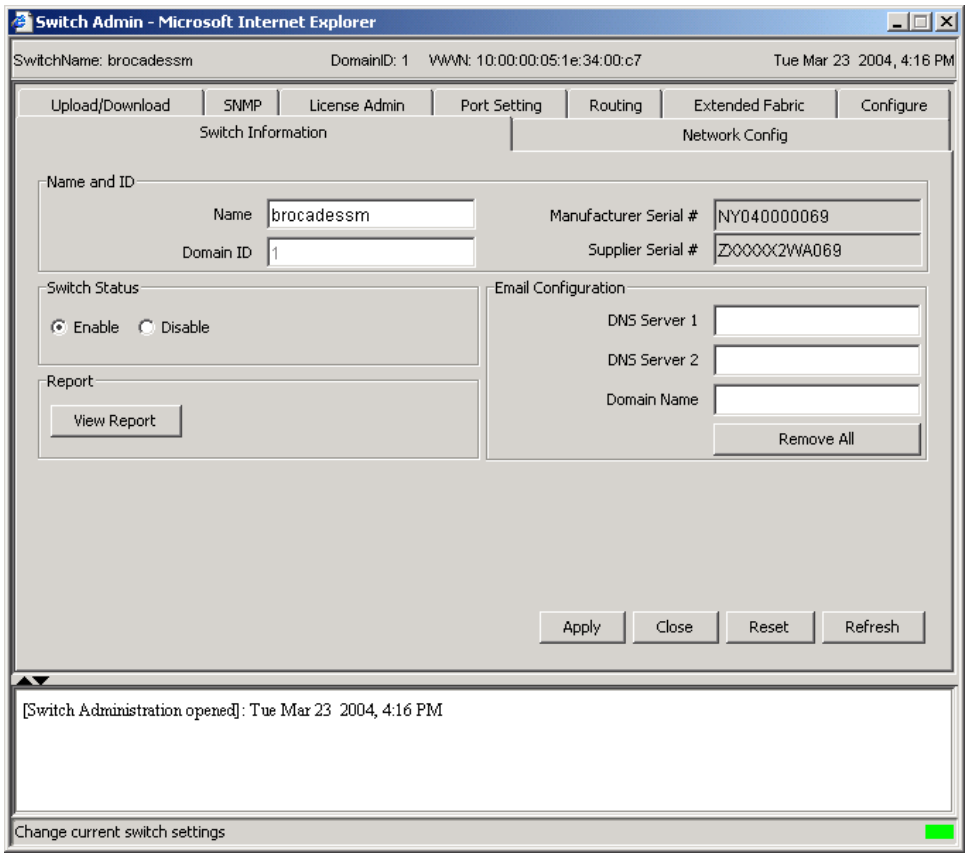

7 a & b. (Optional) Core PID and Extended Edge PID Formats

If introducing the Brocade SAN Switch Module into an existing Fabric OS 2.x and 3.x environment consider changing the setting of the Core PID on all switches running Fabric OS 2.x and 3.x.

The Brocade SAN Switch Module has the Core PID format ON (set to "1") by default. Fabric OS 4.2 or greater only supports PID Formats 1 and 2 (2 is the Extended Edge Format). If these switches are introduced to a fabric that has switches running Fabric OS 2.x or 3.x with the default Core PID format setting of 0, the Brocade SAN Switch Module will segment. When introducing the Brocade SAN Switch Module to an existing fabric of legacy switches with the Core PID format of "0", it is recommended to change the legacy switches Core PID parameter to "**1**". This is a fabric wide parameter.

Caution — Before changing the Core PID parameter please read the recommendations and details for setting the Core PID format, in the *Fabric OS Procedures Guide* and consult the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0* for more information on the deployment of these settings.

7c. (Optional) Principle switch

The principal switch is responsible for handing out domain IDs to the rest of the fabric upon a fabric build. In some cases it may be desirable to hard set a switch to always be the principal switch. In Fabric OS 4.1 or greater only, a preferred principal switch can be selected. Please read the recommendations and details for setting the Core Principle switch, in the *Fabric OS Procedures Guide* and consult the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0* for more information on the deployment of this feature.

8. Add Fabric OS Licenses

Follow the instructions given by your switch provider. This normally entails using the paper-pak instructions and going to a web site (brocade.com) and entering a transaction key number.

1. In the "Info" tab of the Web Tools main interface look in the LicenseID field. If using the CLI: **`licenseidshow`**. Figure 20-15.
2. In the "License Admin" tab of the Web Tools Admin interface copy the new license into the "New License Key" field, and then click the "Add" button. If using the CLI : **`licenseadd`** (i.e. `brocadessm:admin> licenseadd "SeQedReQRSbfRfeB"`). Figure 20-16.
3. Verify the new licenses are listed in the "License Admin" Tab. You may need to click on the "Refresh" button. If using the CLI : **`licenseshow.`** Figure 20-17.

Figure 20-15. LicenseID information is found in the WebTools "Info" Tab.

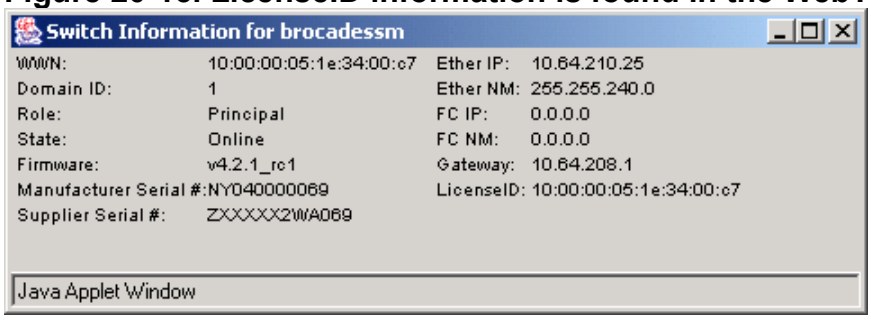

Figure 20-16. Adding a License using WebTools.

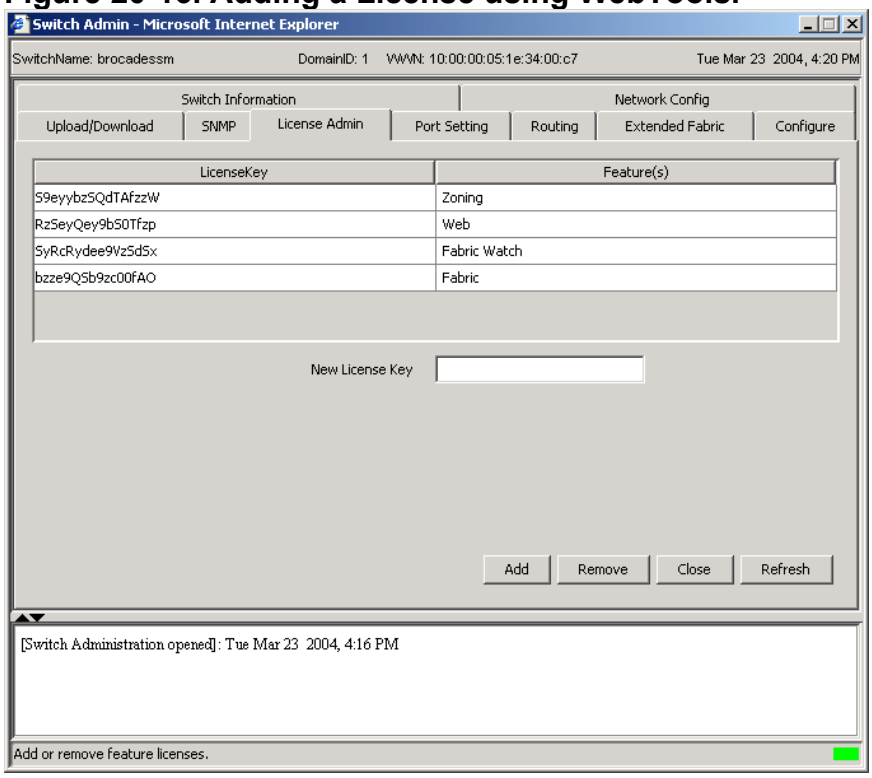

Figure 20-17. New Licenses listed in WebTools.

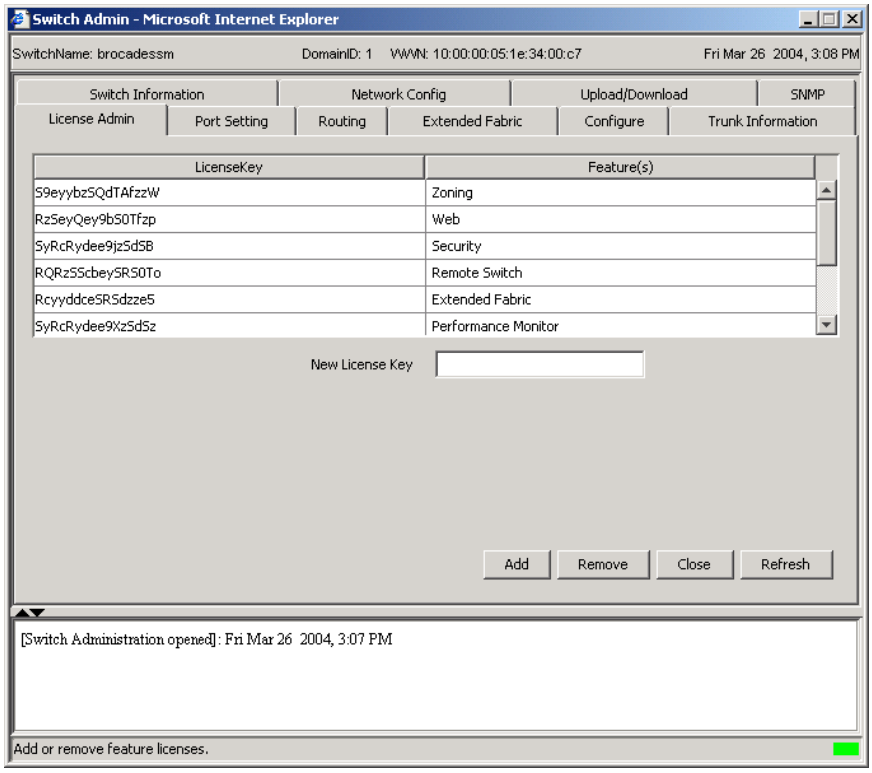

9. (Optional) Port Configurations

It is recommended to maintain the existing Port Configurations settings on the Brocade SAN Switch Module. The Internal Port Configuration Settings for Ports 1-14 should never be changed. If for some reason some of the external Port Configuration settings need to be changed they can be changed from the WebTools interface. However, the CLI interface provides a more comprehensive set of port configuration commands that are not possible through the Web Tools interface. The *Fabric OS Reference Manual* has a comprehensive list.

In the "Port Setting" tab of the Web Tools Admin interface there are several columns, which can be modified. If using the CLI: **portCfgShow** shows the current settings for the ports. Figure 20-18 and 20-19.

Figure 20-18. Port Name information is found in the WebTools "Port Setting" Tab.

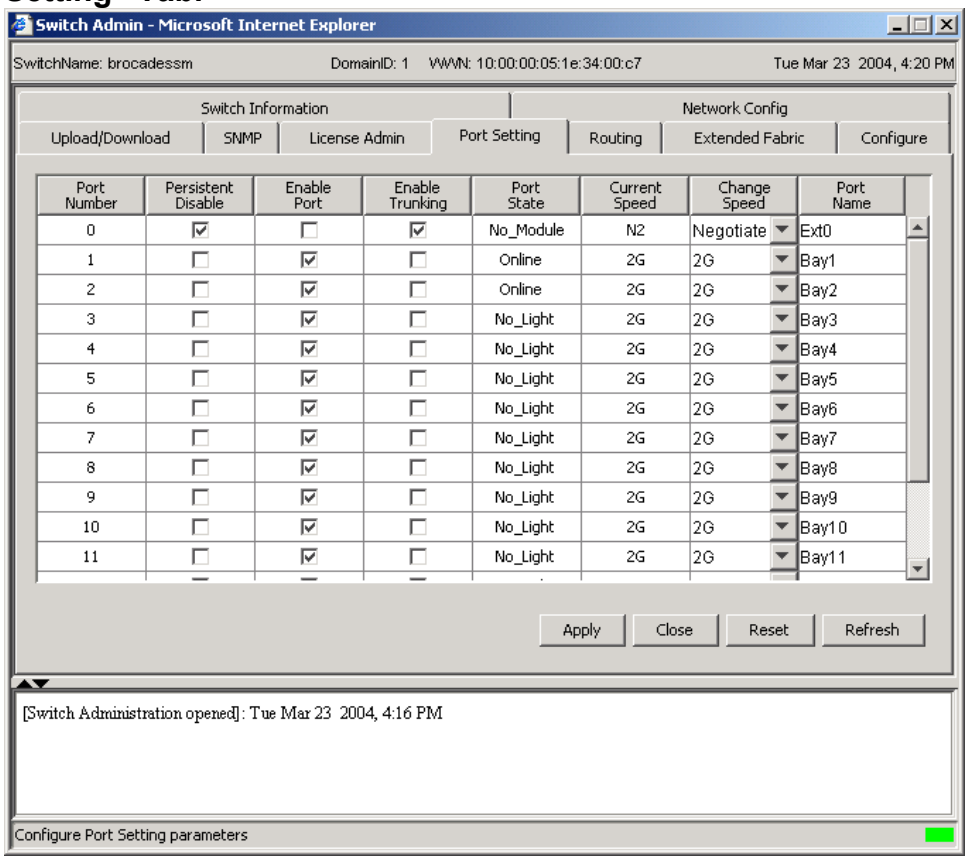

Figure 20-19. portCfgShow information from the CLI.

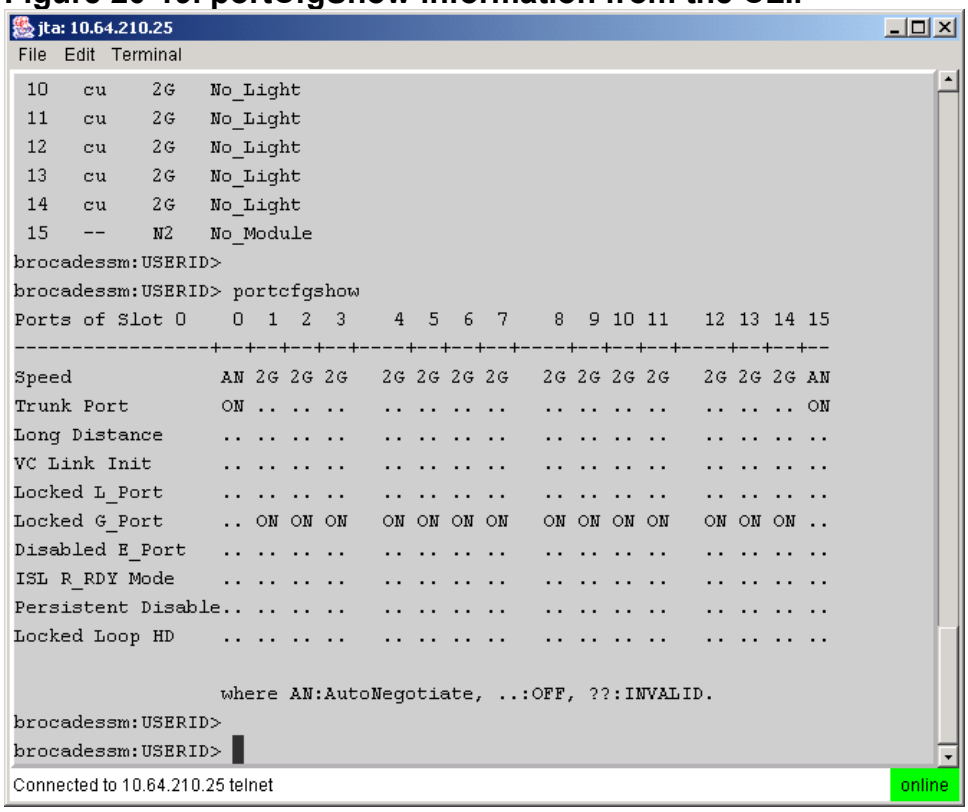

10. (Optional) Name Devices with PortName

In Fabric OS 4.2.1 there is a command, **`portname`**, that lets the administrator label a port. Use **`portname`** to label port 26 on a Brocade SAN Switch Module as shown.

In the "Port Setting" tab of the Web Tools Admin interface there is a column called "Port Name". The Port Name can be modified from this location. If using the CLI: **`portName`**. Figure 20-20.

Figure 20-20. Port Name information is found in the WebTools "Port Setting" Tab.

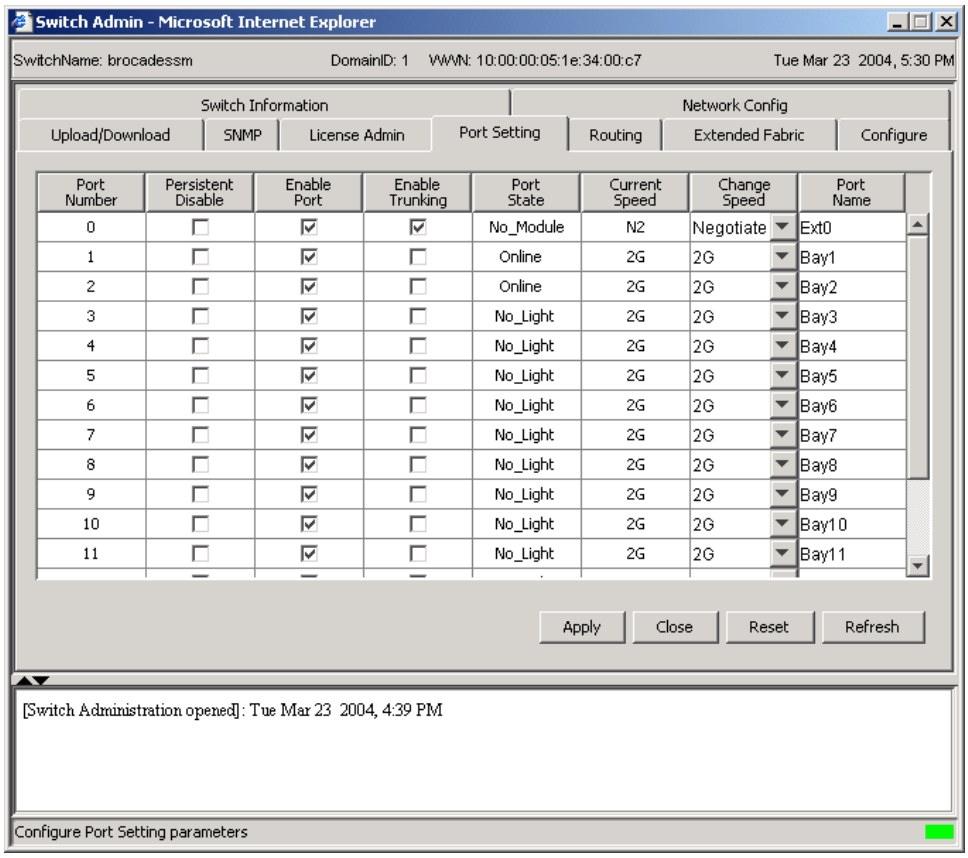

Guideline	The Brocade SAN Switch Module comes pre-configured with PortNames that match the internal Server Blade Bays and External Port identifiers. It is best to append to the existing Port Names with additional information that might be pertinent for that port.

Note	Port names can be cleared on port-by-port basis with **`portCfgDefault`**. The **`portShow`** command displays port name in the first line of the output. **`switchShow`** does NOT display the portName. The port name label is persistent across switch reboots and power cycles.

11. (Optional) Set Telnet Session Timeout Value

It is recommended to at least maintain the default admin telnet session timeout value, which is set to 10 minutes. This ensures security best practices and prevents a telnet session from locking up access to a switch. The Brocade SAN Switch Module's telnet timeout default is 10 minutes. If you wish to change this value use the **`timeout`** CLI command.

12. (Optional) Customize Monitoring Features

Although optional, it is recommended to make you aware of the following features :

- **`SwitchStatusPolicySet`** which monitors the status of several key health parameters
- **`TrackChangesSet`** which tracks the occurrence of an administrator logging into the switch or logging out of a switch and whether a switches configuration has changed
- Configuration of the Switch Module's SNMP agent so that you can incorporate the switch into your existing SNMP management framework
- Fabric Watch tuning for your particular environment.

For guidelines on setting these refer to *SAN Management* section later on in this Guide. Additional information can be found in the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0* as well.

13. Environmental Status

There are several commands that are recommended to check the switch environmental status. The recommended commands on the Brocade SAN Switch Module are: **`uptime`**, **`tempshow`**, **`sensorshow`**, and **`switchstatushow`**. **`Switchstatusshow`** displays any triggered messages set by **`switchstatuspolicyset`**. **`Switchstatuspolicyset`**, discussed in the SAN Management section, allows the environmental settings that trigger warning messages to be customized.

14. Baseline and Backup the Switch

Once all the settings are complete, use `configupload` to backup the switch configuration. This uploads all of the configuration parameters that define the switch configuration. This is highly recommended in the event the original settings need to be restored. If minor changes to the configuration are required, just edit the configuration file and re-download the configuration to the switch. One further benefit is that a "standard" switch configuration can be defined and uploaded to all of the remaining switches in the fabric. This is known as baselining.

Guideline	Define a "golden" switch. For larger SAN fabrics, or for the staging of many smaller fabrics, use **`configupload`** to baseline the "golden" switch. The "golden" switch configuration can be downloaded with Fabric Manager to all others in the fabric. Do baselines by Fabric OS version and platform (i.e. have a separate configdownload file for the Brocade SAN Switch Modules). Doing this helps when it comes time for change management.

20

Preparing the Fabric Configuration

Once each switch has been prepared, the SAN fabric can be built and configured. A checklist provides some essential high-level guidelines and task order. Most of what needs to be done is the zoning of devices. Once these steps are complete, the staging is essentially complete. This section only covers the high level outline of that checklist, please consult the *Brocade SilkWorm Design, Deployment, and Management Guide SAN DDM Version 3.0* for detailed step-by-step procedures.

SAN Fabric Configuration Checklist

1. Gather Planning Documentation (ISL Map, Device Spreadsheet, Logical Design Diagram)
2. Cable ISLs. Check the fabric using the commands: **`islshow`**, **`trunkshow`**
3. Cable Host and Storage Devices. Check using the commands: **`switchshow`** and **`nsshow`**
4. Use documentation to label all cables.
5. Check consistency of Routing Settings for DLS and IOD on each switch in the fabric *
6. Optional: Create a Dummy Zone to prevent device access.

   The default zoning configuration allows any device to see any other device. If devices are plugged in and there is an initial desire to have all of them locked out until you can create proper zones, define a dummy-zoning configuration. To do this, create a zone configuration with an unused port. This example shows the four steps that creates a dummy zone for switch domain 101 and port 1. With this complete, no devices are allowed to access any other device.

   ```
   brocadessm:USERID> zonecreate "dummyzone","101,1"
   brocadessm:USERID> cfgcreate "dummycfg","dummyzone"
   brocadessm:USERID> cfgenable "dummycfg"
   brocadessm:USERID> cfgsave
   ```

7. Setup an NTP Time Service for the Fabric. Use the Principal or Primary FCS Switch only, to connect to the external NTP time server.
8. Validate the fabric by comparing the entries in the Name Server with the devices attached to the fabric. (i.e. compare the number of entries in "**`nsallshow`**" with the number of devices attached to the switches.
9. Pre-Zoning Check. Verify hosts and storage are in the fabric that are to be zoned. On a core switch use the command: **`nscamshow`**
10. Implement Zoning.
11. Optional: Implement Security
12. Profile each SAN fabric.

* Routing Settings (DLS and IOD)

Dynamic Load-Sharing (DLS)

Routing is generally based on the incoming port and the destination domain. This means that all the traffic coming in from a port (either an E_Port or an Fx_Port) directed to the same remote domain is routed through the same output E_Port.

If DLS is turned off (using **`dlsReset`**), load sharing is performed only at boot time or when an Fx_Port comes up. This is the factory default for the Brocade SAN Switch Module and all IBM TotalStorage SAN switches.

By disabling DLS, the possibility of dropped frames is eliminated every time a change in the fabric occurs. A change in the fabric is defined as an E_Port going up or down or an Fx_Port going up or down.

If DLS is turned on (using **dlsSet**), when there are multiple equivalent paths to a remote switch, traffic is shared among all the paths. Load sharing is recomputed when a switch is booted up or every time a change in the fabric occurs. With DLS enabled, traffic on existing ISL ports might be affected when one or more new ISLs are added between the same two switches. Specifically, adding the new ISL might result in dropped frames as routes are adjusted to take advantage of the bandwidth provided by the new ISL.

Enabling DLS optimizes fabric routing. For example, if an Fx_Port goes down, another Fx_Port might be rerouted from one E_Port to a different E_Port. The switch minimizes the number of routing changes, but some are necessary to achieve optimal load sharing. This is the factory default on all Brocade branded switches.

Note Follow the guidance of your SAN support vendor for the configuration of this setting.

Viewing and changing this parameter.

1. Enter the **dlsshow** command to view the current DLS setting.

One of the following messages appears:

- DLS is set. (This message means that the DLS option is turned on. Load sharing is reconfigured with every change in the fabric.)
- DLS is not set. (This message means that the DLS option is turned off. Load sharing is only reconfigured when the switch is rebooted or an Fx_Port comes up.)

2. Enter the **dlsset** command to enable Dynamic Load Sharing.
3. Enter the **dlsReset** command to disable Dynamic Load Sharing.

Example

```
switch:admin> dlsShow
DLS is not set
switch:admin> configshow "route.s"
route.stickyRoutes:1

switch:admin> dlsSet
Committing configuration...done.
switch:admin> dlsShow
DLS is set
switch:admin> configshow "route.s"
route.stickyRoutes:0
```

In Order Delivery (IOD)

The IOD parameter, enforces in-order delivery of frames during a fabric topology change. In a stable fabric, frames are always delivered in order, even when the traffic between switches is shared among multiple paths. However, when topology changes occur in the fabric (for instance, a link goes down), traffic is rerouted around the failure and some frames might be delivered out of order. The setting of this parameter ensures that frames are not delivered out-of-order, even during fabric topology changes.

If IOD is turned on (using **iodSet**), frames are always delivered in order even during fabric topology changes. This is the factory default for the Brocade SAN Switch Module and all IBM TotalStorage SAN switches.

If IOD is turned off (using **dlsReset**), frames may be delivered out-of order during fabric topology changes. The default behavior for Brocade branded switches is for the IOD option to be off, this enables fast rerouting after a fabric topology change.

Note	Follow the guidance of your SAN support vendor for the configuration of this setting.

Viewing and changing this parameter.

1. Enter the **iodShow** command to view the current IOD setting.

One of the following messages appears:

- IOD is set. (Enables the in-order delivery (IOD) option. Enforces in-order delivery of frames during a fabric topology change.)
- IOD is not set. (Turns off the in-order delivery (IOD) option. This command may cause out-of-order delivery of frames during fabric topology changes.)

2. Enter the **iodSet** command to enable In-Order Delivery.
3. Enter the **iodReset** command to disable In-Order Delivery.

Example

```
switch:admin> iodshow
IOD is set
switch:admin> configshow "route.d"
route.delayReroute:1

switch:admin> iodreset
Committing configuration...done.
switch:admin> iodshow
IOD is not set
switch:admin> configshow "route.d"
route.delayReroute:0
```

Profiling The SAN

Once the SAN fabric has been built and zoning has been configured, it is good practice to capture a profile of the fabric. The “show” commands are ideal for this. This list can be used for Brocade SAN Switch Module switches, which run Fabric OS 4.2.1or greater.

These commands can be easily scripted using the Brocade API Scripting Toolkit. As changes take place in the fabric, consider periodically updating the profile information. This will proactively simplify change management and provide a living documentation set that can be referenced when technical support is required.

Brocade SAN Switch Module Profiling Commands

```
version
firmwareshow
chassisshow
licensidshow
licenseshow
ipaddrshow
switchshow
portcfgshow
configshow
agtcfgshow
snmpmibcapset
cfgshow
nsallshow
nsshow
nscamshow
fabricshow
islshow
trunkshow
tempshow
sensorshow
sfpshow
errdump
```

Chapter 21

Validation

Once the SAN is staged, it is recommended to verify its functionality and robustness before going into production. While less important for the entry-level environment, validation becomes critical for SANs with higher port counts. While all pertinent tests for a particular implementation will not be discussed, this section should provide a framework for a customized validation plan. Sample procedures in this section will demonstrate how to check the SAN stability and High Availability (HA). While not covered, it is important to validate the Secure FOS environment as well. If at all possible, it is a good idea to do these tests with generated I/O, preferably with the application up and running.

These tests are meant to be used as guidelines and a proof point that the SAN is operating properly before it is put into production. This section will focus on validating a Core/Edge SAN. If Core/Edge is not used, all tests in this section can be tailored for other fabric topologies. Separate tests will be required for the particular application in use in the SAN.

This chapter contains the following sections:

- Sample Script to Generate I/O
- Sample Validation Recommendations
    - Fabric Stability Validation
    - High Availability ISL Failure Simulation
    - High Availability Switch Failure Simulation

Sample Script to Generate I/O

With no application, it is possible to generate I/O. If using UNIX hosts the following sample script can be used. For Windows hosts, use an I/O tool such as IOmeter. This script creates and writes to a file and then does continuous reads of it. The path and size in blocks need to be specified. As an example, **`sbtest /file01 1000`** will create a file of size 1000 blocks in /file01 and once created, it will do successive reads until terminated. More I/O can be generated on a single host using multiple instances of this script running in the background. **`Portperfshow`** is a helpful CLI command that can be used to quickly check expected I/O performance and functionality. (The default Performance Monitoring Graphs in Web Tools will work well too). The sample script is shown below, Figure 21-1.

Figure 21-1. Sample I/O script for UNIX validation

```
#!/bin/sh
PATH=$1
SIZE=$2
COUNT=`/usr/bin/expr $SIZE \* 2`
TMPFILE="$PATH/sbtest.$$"
RUN=0
echo "Building test file ($TMPFILE)..."
/usr/bin/dd if=/dev/zero of=$TMPFILE bs=512k count=$COUNT > /dev/null 2>&1
echo "Done."
while [ 0 -eq 0 ];
do
DATE=`/usr/bin/date`
echo "Run #: $RUN Timestamp: $DATE"
/usr/bin/dd if=$TMPFILE of=/dev/null bs=512k count=$COUNT >/dev/null 2>&1
RUN=`/usr/bin/expr $RUN + 1`
Done
```

Sample Validation Recommendations

SAN Fabric Validation Checklist

1. Fabric Stability Validation
2. High Availability ISL Failure Simulation
3. High Availability Switch Failure Simulation

Fabric Stability Validation

For validating the stability of the fabric, run the host application software for a period of time. If available, use simulation tools provided by the software application vendor. Run the sample script above for approximately72 hours to check for I/O stability and observe for any problems. Switch event logs should be setup to capture any problems in the fabric. In addition, a host syslogd can be setup to capture event data. No problems should be observed. If there are, troubleshoot them at this time.

Guideline — Check the switch event logs for any issues with the fabric while validating its stability. Setup the syslogd system log on a host and configure to capture fabric event messages.

High Availability ISL Failure Simulation

Take a case where all hosts are attached to one edge switch. From this switch, one trunk with two ISLs is attached to a core switch to form a trunk group with 4 Gbit/sec of available bandwidth. This trunk group contains an ISL trunk master and a second ISL as a trunk member. A maximum of 400 MB/sec should be generated across them to provide maximum ISL stressing. These tests simulate failures on an edge and a core switch. Hot code load should also be tested before going into production.

Case 1 Member ISL Cable Pull

With I/O, remove the non-trunk master ISL for ten seconds and then replace it. Repeat this test three times. I/O should continue without any effect. Some frames may be "lost" when the cable is pulled. If this is the case, the host will log messages indicating that the I/O is being retried.

Case 2 Trunk Master ISL Cable Pull

With I/O, remove the Trunk master ISL for ten seconds and then replace it. Repeat this test three times. I/O traffic may be paused briefly on the second fabric, however the I/O should not timeout on the host.

High Availability Switch Failure Simulation

Case 1 Edge Switch Failure

Initiate a **`fastboot`** on a switch module with no I/O to simulate edge switch failure. Repeat this test three times in succession. The switch should fully recover after the **`fastboot`** and the fabric should rebuild. Other devices not associated with the rebooted switch should continue to perform I/O. Verify this with the **`fabricshow`** command.

Case 2 Core Switch Failure

Initiate a **`fastboot`** on a core switch. Repeat this test three times in succession. The switch should full recovery after the **`fastboot`** and the fabric should rebuild. Verify this with the **`fabricshow`** command.

Case 3 CP Failover (SilkWorm 12000/24000)

Initiate an **`hafailover`** on an active CP with I/O running in the fabric. Full recovery without an I/O pause should occur in all cases.

Case 4 Non-Disruptive Code Load (Fabric OS 4.1 or higher only)

Initiate a firmwaredownload on a switch with Fabric OS 4.1 or higher. Full recovery without an I/O pause should occur in all cases.

Chapter 22

Maintenance

After the Brocade SAN fabrics have been staged and validated it is important to maintain them for proper and continued operation. This chapter will provide guidelines on firmware upgrades, troubleshooting commands for gathering important information about the SAN.

This chapter contains the following sections:

- Fabric OS Upgrades
- Using Fabric OS Troubleshooting Tools
- Support Show Command Groups
- Persistently Disabling a Switch or Port
- PathInfo
- Support for Boot Over SAN

Fabric OS Upgrades

This section will provide some high level guidelines when upgrading firmware on the Brocade SAN Switch Module from Fabric 4.2.1 to a later release that contains new features or bug fixes. This section will cover Web Tools firmwaredownloads on a single switch (CLI commands will be mentioned). For upgrading multiple switches, Brocade Fabric Manager is suggested.

In Fabric OS 4.x, **`firmwaredownload`** only allows the FTP protocol to be used for upgrades. When upgrading from Fabric OS 4.2.1 to a newer release on the Brocade SAN Switch Module, there is only one processor and the firmware must failover to itself. In this case, the Fabric OS Linux kernel and other processes that run in user space must be stopped and restarted through a **`fastboot`**. By default, a **`firmwarecommit`** is launched after the switch is rebooted. This process runs in the background and copies the new firmware from the flash memory primary partition (just upgraded) to the backup partition.

`Firmwaredownloadstatus` can be used on another telnet session to check the upgrade progress. This essentially plays back the **`firmwaredownload`** log. The only message that will be received during the upgrade is "Firmwaredownload has started." After the **`fastboot`**, which happens automatically, run **`firmwaredownloadstatus`**. The total upgrade time is about 10 minutes, including the required **`firmwarecommit`**. This process backs up the primary flash partition to a secondary one. Although it takes about 13 minutes to complete the upgrade on a Brocade SAN Switch Module, it will be non-disruptive (Non-disruptive means that data traffic continues to flow during the failover or firmware upgrade). The fabric state will be saved in non-volatile memory, and after the reboot, all of the fabric services start up first. The characteristics of firmware activation on the Brocade SAN Switch Module are that data flow (Reads, Writes) is not delayed at all. The Fibre Channel ports and ASICs remain programmed and data continues to flow. The Fabric Services (Name Server, FLOGI, etc.) are paused between 30-50 seconds, while the switch goes through its **`fastboot`** process and the processor comes back with a new image.

1. In the "Upload/Download" tab of the Web Tools Admin interface there is a function button called "Firmware Download". This button must be chosen. The protocol drop down menu only allows the "ftp" service to be chosen for this switch type. The FTP server information must be placed into the following fields : User Name, Host IP, Password and File Name. The Fabric OS 4.2.1 code is typically bundled into a ".zip" file for Windows servers or a ".tar.gz" file for UNIX servers. These compressed files will need to be uncompressed in an appropriate directory on the FTP server. Once unbundled there are a number of files but the user must only indicate a hidden symbolic link file called "release.plist" in the File Name for download. If using the CLI: **`firmwaredownload`**. An example is provided in Figure 22-1. A confirmation pop-up that comes up after initiating is shown in Figure 22-2. The Fabric Manager interface is shown in Figure 22-3. For more information on Firmware Downloads read the *Brocade Fabric OS Procedures Guide* and the *Brocade Fabric OS Reference Manual*.

Figure 22-1. Firmwaredownload using WebTools.

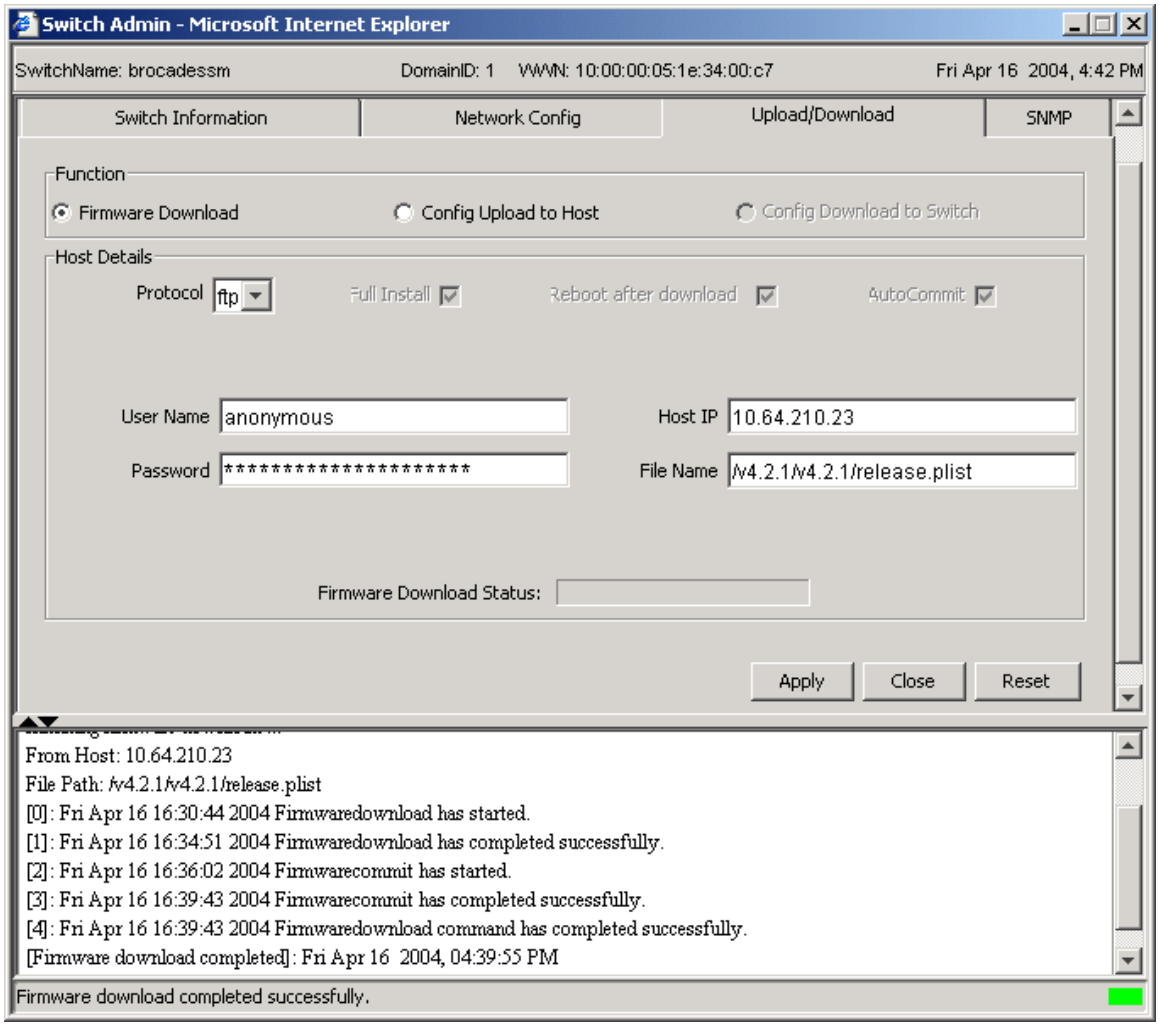

Figure 22-2. WebTools Firmware Download Confirmation pop-up.

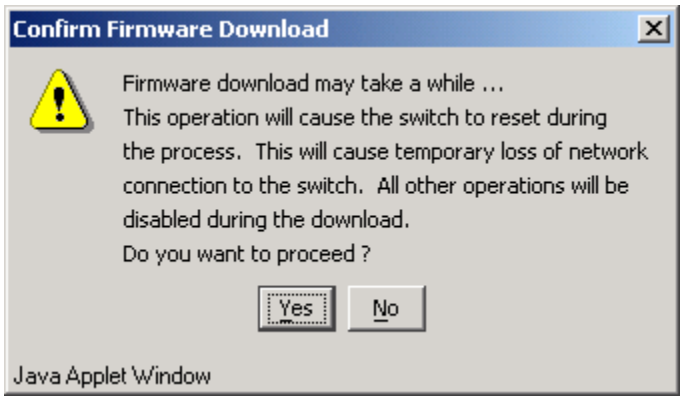

Figure 22-3. Firmwaredownload using Fabric Manager.

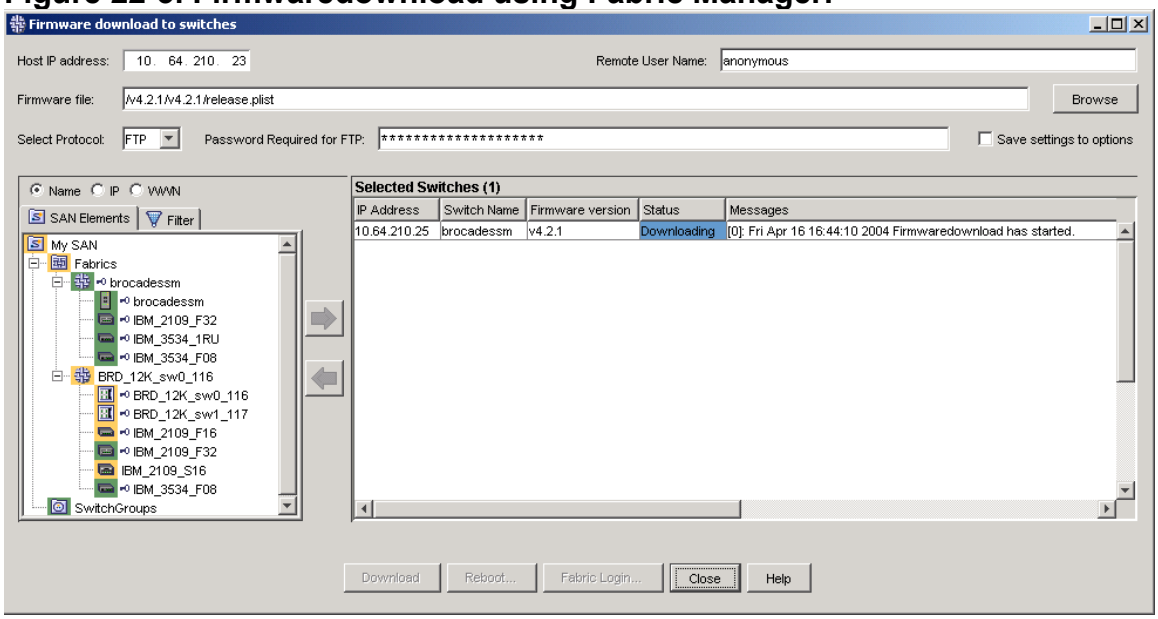

Using Fabric OS Troubleshooting Tools

This section will just provide an overview of Brocade Fabric OS troubleshooting tools. Some examples will be given to demonstrate the practical use. For details and examples for setting up persistent error logs and the uses of **portlogdump**, please see the *Brocade Fabric OS Procedures Guide*.

Support Show Command Groups

Supportshow is an invaluable troubleshooting tool. It essentially provides a one time event capture of the entire system at the point in time the command is run. It includes many of the typical "show" commands that experienced Fabric OS users are familiar with and much more output so that the Support organization can resolve any issue without having to go back for more data. **Supportshow** can be run with any combination of 11 command groups. This makes **Supportshow** more flexible and easier to capture the desired information.

If required for support reasons, **supportshow** can always be reconfigured to display more information. As an example of a command group, here are the commands for the FC Fabric command group number 4.

```
fabricShow
islShow
trunkShow
topologyShow
faShow
qlShow
cfgShow
fabStatsShow
fabLogDump
```

This section will provide some guidelines on setting up **supportshow** command groups and make a recommendation on which ones to use. By default, eight command groups are enabled for **supportshow**. These are shown by **supportshowcfgshow** output below (admin/USERID level command). However, changes to the **supportshow** configuration can only be made as the root user.

```
brocadessm:USERID> supportshowcfgshow
 os         enabled
 exception  enabled
 port       enabled
 fabric     enabled
 services   enabled
 security   enabled
 network    enabled
 portlog    enabled
 system     enabled
 extend     disabled
 filter     disabled
 perfmon    disabled
brocadessm:USERID>
```

If you have the root password consider configuring **supportshow** to disable the following groups. For general information about the fabric and the switch **supportshow** is running on, these groups are really all that is needed. For most SAN administrative troubleshooting cases, the data provided by the remaining groups will do.

```
brocadessm:root> supportshowcfgdisable "os"
Config update Succeeded
brocadessm:root> supportshowcfgdisable "port"
Config update Succeeded
brocadessm:root> supportshowcfgdisable "security"
Config update Succeeded
brocadessm:root> supportshowcfgdisable "portlog"
Config update Succeeded
brocadessm:root> supportshowcfgdisable "network"
Config update Succeeded
```

If using Advanced Security (Secure Fabric OS), do not disable the security command group as shown. Finally, verify the setting with **supportshowcfgshow**. The output that should be seen is shown in the following example.

```
brocadessm :root> supportshowcfgshow
os           disabled
exception    enabled
port         disabled
fabric       enabled
services     enabled
security     disabled
network      disabled
portlog      disabled
system       enabled
extend       disabled
filter       disabled
perfmon      disabled
brocadessm:root>
```

For all **supportShow** output, no matter how its configured, there are certain commands that will always be executed. These are, in the order of execution : **date**, **version** and **supportshowcfgshow**. Enabling and disabling is persistent except for the filter and extended groups. This is not a concern as these two command groups are rarely used. Two command groups have only one member. The exception group has **errdump**. The portlog group only has **portlogdump**.

The **supportshowcfgshow** commands :

supportShowCfgShow - Displays list of command groups and whether they are enabled.
supportShowCfgEnable - Allows root user to enable a single command group
supportShowCfgDisable - Allows root user to disable a single command group

Persistently Disabling a Switch or Port

There are four commands to allow for persistently disabling ports or switches. When configured, the state of the switch or port will remain disabled through power cycles or reboots. Reasons why this may be done include:

- There may be a bad SFP or switch, which causes fabric instability. These may need to be brought down temporarily until a replacement is found.
- Unused ports may be persistently disabled for security concerns. When in this state, no device or switch will be allowed to join the fabric on that port.
- To prevent operator error when devices or switches are connected to the wrong port.

To disable a switch persistently, use **`switchCfgPersistentDisable`**. After a few moments, the switch is disabled. To verify, use **`switchshow`** to display the current state. Note that *SwitchRole* is now *Disabled (Persistent)*. This indicates that the command has taken effect. The **`portcfgshow`** still shows all ports as not disabled. This is fine, since the switch as a whole is disabled.

To re-enable the switch persistently, use **`switchCfgPersistentEnable`**. Note that the fabric must re-configure after it is brought back online.

Note	If the switch is re-enabled with `switchenable` it will be enabled temporarily. The next power cycle, **reboot** or **fastboot** will cause the switch to be disabled, persistently. This is because the state of the switch is now stored in the flash non-volatile memory.

Use **`portCfgPersistentDisable`** to persistently disable a port. Use **`portcfgshow`** to check the status. To re-enable a port persistently, use **`portcfgpersistentenable`**. To temporarily enabled a port, use **`portenable`**.

Note	**`PortCfgDefault`** will turn off persistent disabling of a port. **`PortCfgDefault`** will also set all other port settings to default values.

Warning	When a **`portcfgpersistentdisable`** is done on an enabled E-port, a fabric reconfiguration may occur. This is the same behavior as the **`portdisable`** command. Persistently disabling groups of ports is not supported. Each port must be persistently disabled individually.

Guideline	Do not use the **`portCfgDefault`** command on the Brocade SAN Switch Module. Use individual port commands to configure ports. The Brocade SAN Switch Module is shipped with particular defaults to ensure it works properly as an embedded system where known deterministic devices will be attached to the Internal 14 ports.

Note	The IBM eServer BladeCenter Management Module performs the **`portCfgPersistentDisable`** if the “Disable” External Ports has been configured in the Advanced Setup Menu of the IO Module Tasks for that IO bay. Figure 22-4.

Figure 22-4. IBM eServer BladeCenter Management Module External Ports Setting.

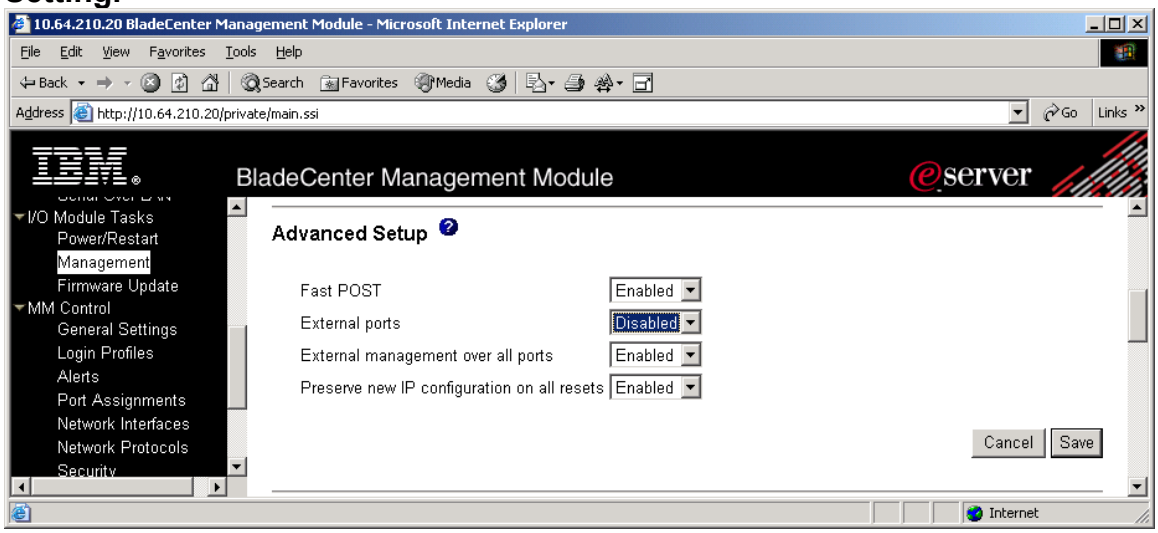

Pathinfo

`Pathinfo` is new command available in Fabric OS 2.6.1, 3.1.2 and 4.2/4.2.1.

The **`pathinfo`** command provides *traceroute* functionality for the SAN. It allows the paths to be discovered and viewed in a usable form. The previous method of tracing paths used **`urouteshow`** and was complicated to analyze.

The fabric topology information is actually known at every switch, however FSPF routing is known only locally; there is no global, end-to-end routing table. In this context, end-to-end refers to a source switch with connected E-Ports and a destination switch also with connected E-Ports and connected devices. In general, each switch in the fabric contains a table of routes that enable frames to be forwarded to adjoining switches. If desired, an end-to-end routing table can be constructed by logging into every switch and reading the local routing tables. **`Pathinfo`** simplifies this task by providing a flexible command line and an interactive menu. This interactive menu is displayed by default if no arguments are supplied.

PathInfo provides additional routing information including.

- Destination port state
- Link statistics for every hop from source to destination
- Link utilization for each hop from source to destination

PathInfo Examples

As designed, **`pathinfo`** is intended to gather information on a specific data stream, not the entire fabric. The simplest example uses a destination domain as an argument. When used, this will provide routing information from the embedded port on the local switch to the embedded port on the remote switch domain. The example below shows the command and expected output for the destination domain of 32. Each hop is shown along the way to the destination.

```
ess031:admin> pathinfo 32, -1, 13
Target port is Embedded

Hop  In Port Domain ID (Name)   Out Port     BW      Cost
------------------------------------------------------------
0       E       31      (ess031)        9       1G      1000
1       23      93      (ess093)        5       2G      500
2       6       32      (ESS032)        E       -       -
value = 3 = 0x3
```

Another common usage is to gather the routing information from a source port on a local switch to a destination port on a remote switch. This requires three arguments; the source domain, the source port and the destination port. To use the embedded port, use **-1**. The example, below uses a source domain of 32, it's embedded port and a destination port of 13. To determine more complex pathing, it is recommended to use a diagram of the fabric with labeled ISL connections and device attachments. This allows arguments to be quickly determined. If the appropriate information is **not** entered properly, **pathinfo** may provide error messages.

```
ess031:admin> pathinfo 32, -1, 13
Target port is F_Port

Hop  In Port Domain ID (Name)    Out Port    BW      Cost
------------------------------------------------------------
0       E       31      (ess031)        9       1G      1000
1       23      93      (ess093)        5       2G      500
2       6       32      (ESS032)        13      1G      -
value = 3 = 0x3
```

An example of this is shown next. If no devices are attached to the destination port being used as the argument, the error message displayed is target port not active. Normally this does not mean that no frames are being passed to the device or that it is not online, it simply means no device is attached. If a device is attached, this may mean the device is not online.

```
ess031:admin> pathinfo 32, -1, 14
Target port not active

Hop   In Port Domain ID (Name)    Out Port   BW      Cost
------------------------------------------------------------
0       E       31      (ess031)        9       1G      1000
1       23      93      (ess093)        5       2G      500
2       6       32      (ESS032)        14      1G      -
```

Guideline

For determining complex path information quickly and easily, create or use a diagram of the SAN fabric with labeled ISL connections and device locations. This simplifies the process of gathering the appropriate **pathinfo** arguments. Use -1 as the argument for an embedded switch port.

This next example shows how to display the return path. Use the **–r** option as shown. Depending on the fabric topology, it may not be the same as the outbound path. The following example shows that the reverse path in this case is equivalent.

```
ess031:admin> pathinfo "-r", 32, -1, 13
Target port is F_Port

Hop In Port Domain ID (Name) Out Port BW Cost
------------------------------------------------------------
0       E       31      (ess031)        9       1G      1000
1       23      93      (ess093)        5       2G      500
2       6       32      (ESS032)        13      1G      -

Reverse path

3       13      32      (ESS032)        6       2G      500
4       5       93      (ess093)        23      1G      1000
5       9       31      (ess031)        E       -       -

value = 5 = 0x5
```

If no parameters are given, the PathInfo command works in interactive mode. The menu is displayed much like the configure command. Interactive mode allows more parameters to be specified. The choices are shown below with default values.

- max hops (default = 25)
- domain (required)
- source port (default = embedded)
- destination port (default = embedded)
- enable basic statistics (default = no)
- enable extended statistics (default = no)
- trace reverse path (default = no)
- source route (default = no)
- strict source routing (may be specified if source route enabled)
- timeout (default = 5 seconds)

The interactive mode menu is shown with various choices being made are shown in the following output:

```
ess031:admin> pathinfo

Max hops: (1..127) [25]
Domain: (1..239) [-1] 32
Source port: (0..15) [-1]
Destination port: (0..255) [-1] 13
Basic stats (yes, y, no, n): [no] y
Extended stats (yes, y, no, n): [no] y
Trace reverse path (yes, y, no, n): [no] y
Source route (yes, y, no, n): [no] n
Timeout: (1..30) [5]
```

Pathinfo can also be used in advanced mode to display basic statistics. The following example shows partial output with extended statistics for one hop.

```
Hop  In Port Domain ID (Name)     Out Port   BW      Cost
------------------------------------------------------------
1       23      93      (ess093)       5       2G      500

Port                    23                     5
                Tx              Rx      Tx              Rx
                ------------------------------------------------
F/s (1s)        0               0       0               0
F/s (64s)       0               0       0               0
Words           9270            8003    9296            9741
Frames          435             450     544             535
Errors          -               0       -               0
```

Support for Boot Over SAN

Contact IBM Support for supported configurations and instructions for configuring the IBM eServer BladeCenter Fibre Channel Expansion Card as the boot device. For more information on the Fibre Channel Expansion Cards refer to the latest IBM eServer BladeCenter Fibre Channel Expansion Card BIOS/driver Readme and the *IBM eServer BladeCenter Fibre Channel Expansion Card Installation and User's Guide* available at :

http://www.ibm.com/pc/support

In addition certain storage vendors may have statements of support for Remote Boot environments, please contact them directly.

SAN Management

This section discusses fundamental SAN Management concepts and associated guidelines.

This section contains the following chapters:

- SAN Management Overview
- Brocade SAN Switch Module Management Tools
- SNMP
- Fabric Watch
- Advanced Performance Monitoring

Chapter 23

SAN Management Overview

The trend of fabrics increasing in switch count, and geographically separate locations of fabrics and switches, has created the need for centralized, secure, and cost effective SAN management. A SAN consists of switches, host systems and storage devices. Switch vendors have provided several command and web based tools to manage switch configurations. A few major storage providers have integrated a basic level of switch management utilities into their existing management software packages. However, their primary focus remains on managing storage, not the fabric. As SAN fabrics become larger, managing a fabric by accessing each individual switch from a command line interface can be an inefficient and time consuming process. SAN management software must be capable of performing complex tasks including configuring, maintaining, monitoring and troubleshooting a SAN in a simple and effective manner. Although, there are many SAN management software packages in the market place, they are somewhat limited in providing SAN level management functionality. Today's SAN Management software applications must include:

- SAN Security
- Graphical display of SAN devices
- Easy access to all switches of the fabric to maintain consistent configurations
- Real time self-monitoring status and advance warning
- Real time critical fabric event detection and notification
- Fault isolation and corrective action
- Analyze long-term behavioral trends

SAN security is a crucial issue that must be addressed by ensuring that only a select group of authorized users are given access to make SAN configuration changes in a controlled manner. SAN access can be secured by implementing secure policies at the system as well as the Fabric OS level.

Graphical display of a SAN provides a quick view of its components and their relationship to one another. The device type and in some cases the status of a device should be determined by the device icon coloring scheme. It also simplifies the SAN device access process utilizing the point and click method.

Maintaining a consistent configuration within the fabric, and across fabrics, requires performing a variety of tasks on more than one switch in a fabric. For example, upgrading firmware can be a time consuming task if a large number of switches are involved. An efficient method of upgrading firmware is to discover the switches by logging into a single switch and then initiating the firmware download procedure on a selected group of switches.

A real time self-monitoring status and advanced warning can be very helpful to maintain the health of the SAN by avoiding costly down time in an enterprise environment. In a redundant system, the failure of a FRU may not be necessarily a disastrous event, however the failure does increase the vulnerability of the overall system. An advanced warning helps the administrator plan corrective action in time to restore redundancy.

A fabric event critical in nature can disrupt the fabric operation. For example, a failure of an Inter Switch Link (ISL) between two switches may result in fabric segmentation. When a faulty condition is self-detected and immediate notification is sent by an appropriate method(s), corrective action may be initiated before the event escalates. Thus a self-monitoring system can assist a SAN administrator minimize disruption time.

Diagnostic utilities are desirable to identify a failing component and possible root cause of the failure. A diagnostic utility should be able to run in the background without disturbing the fabric operation. When implemented, it can minimize the maintenance time significantly.

Analyzing behavior patterns on a switch port for an extended period can expose a deficiency and/or anomaly affecting the overall throughput performance of the SAN. Detailed information must be gathered over a period of time, using an application, which can probe the switch on a periodic basis.

SAN Management Scope

The SAN Management administrative scope can be broken down into three distinct areas; Fabric, Switch and Port. Fabric administration tasks typically involve parameters that are applied to many switches at once. Switch tasks generally involve environmental monitoring. Port level tasks generally involve performance monitoring, setting up port level definitions etc. Of course, there is overlap between these areas.

The following list itemizes SAN Management at the Fabric, Switch and Port level and the various administrative tasks associated with them. In each area, a recommended Brocade management tool for these tasks is shown.

Fabric Administration (Fabric Manager)

- SAN Discovery and Topology Display
- SAN Logical Grouping (Switch Groups and Port Groups)
- Fabric Login (Global access)
- Switch/Port Administration
- At-A-Glance views of switch and device information in the fabric
- Fabric Change Management
- ISL Checking
- Fabric Merge Checking (checks for Zoning and Switch Configuration parameter conflicts)
- Web Tools Device Management launcher
- Telnet (Sec-Telnet) launcher
- Consolidated Event Monitoring
- Switch Firmware Download
- Switch Sequence Reboot
- Switch Configuration Download/Upload (Base-lining)
- Complete Fabric Backup and Compare
- Zoning Management (through WebTools)
- Name Server Displays (through WebTools)
- Advanced Security (Secure Fabric OS) Policy Management
- Set Time on all switches in a Fabric
- Call Home support
- Switch License Management
- Firmware Download to FDMI capable Host Bus Adapter

Switch Administration (Web Tools or CLI)

- Switch information
- Network configuration
- Firmwaredownload
- Configuration upload/ download
- SNMP setup
- License Administration
- Fabric parameter configuration
- Routing, DLS and IOD setup
- Zoning Management
- Name Server Displays

Port Administration (Web Tools or CLI)

- Port Configuration
- Extended Fabric configuration
- Trunking information

Chapter 24

Brocade SAN Switch Module Management Tools

Brocade's approach to SAN Management is fabric centric. The main emphasis is to simplify the entire SAN management process by simplifying the three levels of management for Brocade switches from a centralized management server. Brocade has developed and refined a rich command line interface for advanced users and a simple to use Web Tools interface for new users and advanced users alike. Licensed fabric health monitoring tools, including Brocade Fabric Watch and Advanced Performance Monitoring, assist SAN administrators with SAN monitoring and improving the operating environment. As SAN configurations grow larger and more complex, managing them using multiple management tools has become very challenging. Brocade Fabric Manager was introduced to assist users, by simplifying the entire management process by providing centralized management. Brocade Fabric Manger is capable of managing complex tasks at the fabric, switch, and port levels. In addition, access to Brocade Web Tools, Fabric Watch, and Advanced Performance Monitoring is made available via the Fabric Manager interface. This allows for complete SAN management from a single user interface.

Brocade SAN Switch Module management and monitoring products include:

- IBM eServer BladeCenter's Management Module
- Command Line (Telnet)
- Web Tools (HTTP)
- Fabric Manager
- SNMP
- Fabric Watch
- Advanced Performance Monitoring
- Brocade Fabric Access API

The Brocade CLI and WebTools interfaces have been featured throughout this Guide. The following chapters will focus on SNMP, Fabric Watch, and Advanced Performance Monitoring.

Chapter 25

SNMP

One of the standard methods for monitoring and managing a network device is through Simple Network Management Protocol SNMP). It is a universally accepted protocol that is portable, lightweight, and is widely deployed. SNMP allows an administrator to monitor the health and performance of countless devices locally or remotely via an Ethernet port. Enterprise management software, like IBM Tivoli Netview and IBM Tivoli Storage Area Network Manager (ITSANM) can monitor thousands of devices in an enterprise and have extended their support to manage Brocade SANs. These SNMP Management stations can be run separately by the SAN administrator so, that they can get alerts, trend performance and capture details of error status on switches separately. Within the SNMP model, a manageable network consists of one or more management stations and a collection of agent systems that are also known as network elements. The manager communicates with the agent using SNMP Protocol. Brocade currently supports SNMP version 1 (SNMPv1).

The fabric switches must be properly configured either from the command line interface or using Brocade Web Tools. Web Tool can also be accessed from the Fabric Manager Admin menu. A switch configuration setup consists of:

- Brocade SNMP Configuration
- Brocade SNMP Agent Setup
- Brocade SNMP Trap Setup
- Brocade Track Changes Setup

Brocade SNMP Configuration

To configure the SNMP agent software on a Brocade switch no special license key is needed. It is included in the base functionality of the switch. The commands used to configure the SNMP Agent are "**`agtcfgShow`**", "**`agtcfgSet`**" and "**`agtcfgDefault`**". Once the SNMP agent software is configured, SNMP information can be automatically sent to an SNMP management station.

The SNMP trap software on a Brocade switch also requires no special license. The commands used to configure whether SNMP traps are sent out is "**`snmpmibcapset`**". This feature allows the user to choose which MIB elements will be used to send SNMP trap information.

Brocade SNMP Agent Setup

A default **`agtcfgShow`** output is included in Figure 25-1. The most notable aspects of configuring the Agent are included in the following steps :

1. To begin, the *sysDescr*, *sysLocation* and *sysContact* information fields can be changed for the specific environment that the switches are installed in.
2. *swEventTrapLevel* indicates which severity level messages will be sent to the SNMP Manager. By default, this value is set at 0, implying that no swEventTrap is sent. Possible values are :
   0 - none
   1 - critical
   2 - error
   3 - warning
   4 - informational
   5 – debug
   Note : To configure the SNMP Manager to receive all messages this would be changed to a "5".

Guideline	You might want to set the *swEventTrapLevel* to something in the middle say a "3", which will capture all the critical, error and warning messages. If you set it to 4 or 5 you will be inundated with messages, which may be appropriate if you are troubleshooting an issue in the SAN.

3. *authTraps* does some authentication checking if a request comes in from a community that is not known to the agent.
4. The next section of the SNMP agent configuration is where the Trap Recipient IP addresses and community strings are entered. The user should configure the IP addresses and community strings that are on their SNMP management station.
5. The final section is an SNMP Access Control List (ACL). This provides a way for the administrator to restrict SNMP get/set operations to certain hosts/IP addresses. This would be used for enhanced management security in the Storage Area Network.
6. Figure 25-2 shows an example of setting the *swEventTrapLevel* to 5 so that all Trap events are sent to the SNMP manager and configuring the *authTraps* parameter so that authentication-warning messages get sent to the SNMP manager. It also shows the user how to set the IP Address of the Trap recipient.
7. Figure 25-3 shows the resultant configuration settings from command performed in step 2.

These configurations can also be performed from Web Tools Interface through the Admin interface. See Figure 25-4.

Figure 25-1. default agtcfgshow for the Brocade SAN Switch Module

```
brocadessm1:USERID> agtcfgshow
          sysDescr = Brocade SAN Switch Module for IBM eServer BladeCenter
       sysLocation = End User Premise.
        sysContact = Field Support.
  swEventTrapLevel = 0
         authTraps = 0 (OFF)

SNMPv1 community and trap recipient configuration:
  Community 1: Secret C0de (rw)
    No trap recipient configured yet
  Community 2: OrigEquipMfr (rw)
    No trap recipient configured yet
  Community 3: private (rw)
    No trap recipient configured yet
  Community 4: public (ro)
    No trap recipient configured yet
  Community 5: common (ro)
    No trap recipient configured yet
  Community 6: FibreChannel (ro)
    No trap recipient configured yet

SNMP access list configuration:
Entry 0:  No access host configured yet
Entry 1:  No access host configured yet
Entry 2:  No access host configured yet
Entry 3:  No access host configured yet
Entry 4:  No access host configured yet
Entry 5:  No access host configured yet
brocadessm1:USERID>
```

Figure 25-2. agtcfgset command usage

```
brocadessm:USERID> agtcfgset

Customizing MIB-II system variables ...

At each prompt, do one of the following:
  o <Return> to accept current value,
  o enter the appropriate new value,
  o <Control-D> to skip the rest of configuration, or
  o <Control-C> to cancel any change.

To correct any input mistake:
<Backspace> erases the previous character,
<Control-U> erases the whole line,
sysDescr: [Brocade SAN Switch Module for IBM eServer BladeCenter]
sysLocation: [End User Premise.]
sysContact: [Field Support.]
swEventTrapLevel: (0..5) [0] 5
authTrapsEnabled (true, t, false, f): [false] t

SNMP community and trap recipient configuration:
Community (rw): [Secret C0de]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Community (rw): [OrigEquipMfr]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Community (rw): [private]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Community (ro): [public]
Trap Recipient's IP address in dot notation: [0.0.0.0] 192.168.2.10
Community (ro): [common]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Community (ro): [FibreChannel]
Trap Recipient's IP address in dot notation: [0.0.0.0]

SNMP access list configuration:
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
          sysDescr = Brocade SAN Switch Module for IBM eServer
BladeCenter
       sysLocation = End User Premise.
        sysContact = Field Support.
  swEventTrapLevel = 5
        authTraps = 1 (ON)
```

continued

```
SNMPv1 community and trap recipient configuration:
  Community 1: Secret C0de (rw)
    No trap recipient configured yet
  Community 2: OrigEquipMfr (rw)
    No trap recipient configured yet
  Community 3: private (rw)
    No trap recipient configured yet
  Community 4: public (ro)
    Trap recipient: 192.168.2.10
  Community 5: common (ro)
    No trap recipient configured yet
  Community 6: FibreChannel (ro)
    No trap recipient configured yet

SNMP access list configuration:
Entry 0:  No access host configured yet
Entry 1:  No access host configured yet
Entry 2:  No access host configured yet
Entry 3:  No access host configured yet
Entry 4:  No access host configured yet
Entry 5:  No access host configured yet
Committing configuration...done.
Brocadessm:USERID>
```

Figure 25-3. modified agtcfgshow settings

```
brocadessm:USERID> agtcfgshow
          sysDescr = Brocade SAN Switch Module for IBM eServer BladeCenter
       sysLocation = End User Premise.
        sysContact = Field Support.
  swEventTrapLevel = 5
         authTraps = 1 (ON)

SNMPv1 community and trap recipient configuration:
  Community 1: Secret C0de (rw)
    No trap recipient configured yet
  Community 2: OrigEquipMfr (rw)
    No trap recipient configured yet
  Community 3: private (rw)
    No trap recipient configured yet
  Community 4: public (ro)
    Trap recipient: 192.168.2.10
  Community 5: common (ro)
    No trap recipient configured yet
  Community 6: FibreChannel (ro)
    No trap recipient configured yet

SNMP access list configuration:
Entry 0:  No access host configured yet
Entry 1:  No access host configured yet
Entry 2:  No access host configured yet
Entry 3:  No access host configured yet
Entry 4:  No access host configured yet
Entry 5:  No access host configured yet
brocadessm:USERID>
```

Figure 25-4. SNMP Agent Web Tools configuration.

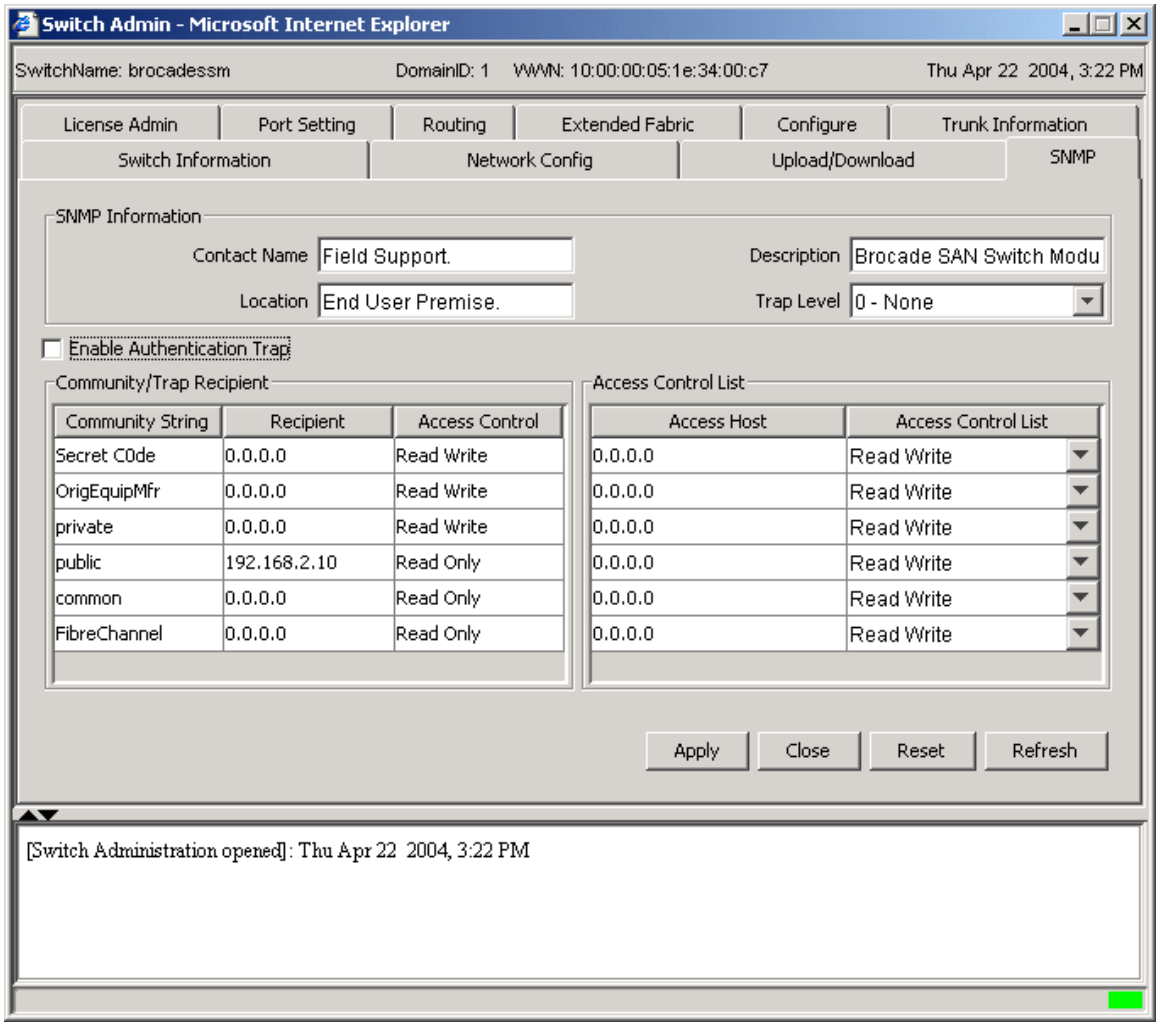

Brocade SNMP Trap Setup

An unsolicited message that goes to the SNMP management station from the SNMP agent on the switch is called a *trap*. Brocade switches send traps out on UDP port 162 only.

The Brocade SNMP Trap configuration allows the user to choose which MIB trap elements will be used to send information to the SNMP management station. There are two main MIB Trap choices for the user to configure. The user can choose to use the Brocade specific MIB Trap, which is associated with the Brocade specific SilkWorm (SW) MIB, or the user can choose to use the Fibre Alliance MIB Trap, which is associated with the Fibre Alliance (FA) MIB. The Brocade MIB was created specifically for monitoring SilkWorm switches. The Fibre Alliance MIB was created by a consortium of companies to manage any type of SAN switch and SAN devices from any company. If you wish to use both, it should be noted that some of the information that is sent is the same for both MIBs. If both Traps are chosen the user may receive multiples of the same traps (duplicates of the same information). The user can also turn off the FA-MIB itself completely. There is no choice to turn off the SW-MIB, it is always enabled.

The user also has the choice of enabling or disabling two additional MIBs and their associated TRAPs the FICON-MIB, HA-MIB, FICON-TRAP and HA-TRAP. The FICON MIB and TRAP are only required for FICON environments and the HA MIB and TRAP are only necessary for the Enterprise Class Switches and Directors (SilkWorm 12000 and SilkWorm 24000). The SW-EXTTRAP is used to provide additional information to a Trap as it includes the *swSsn* (Software Serial Number) to be sent as a part of the Brocade SW Traps. It is also used in conjunction with the SilkWorm 6400 integrated fabrics to provide the detailed group information for a particular trap. The SilkWorm 6400 is no longer sold.

For a description of the Brocade MIB files, naming conventions and loading instructions refer to the *Brocade MIB Reference Manual*. The Brocade SilkWorm (SW) MIB for the Fabric OS 4.2.1 release is named; "SW_v5_1.mib", and can be found in the un-compressed firmware directory called "mibs". The Brocade SW-TRAP information is located within this MIB. The Fibre Alliance SNMP MIB and Traps can be found in the "FA_v3_0.mib" file in the same location.

All of the Brocade MIBs for this release are located in the "mibs" directory of the uncompressed v4.2.1 firmware files.

To configure the Trap options use the "**`snmpmibcapset`**" command. This command is not available from the WebTools interface.

1. Figure 25-5 shows the default values for the **`snmpmibcapset`** telnet command.

Figure 25-5. default snmpmibcapset

```
brocadessm:USERID> snmpmibcapset
The SNMP Mib/Trap Capability has been set to support
FE-MIB SW-MIB FA-MIB FICON-MIB HA-MIB SW-TRAP FA-TRAP FICON-TRAP HA-TRAP
FA-MIB (yes, y, no, n): [yes]
FICON-MIB (yes, y, no, n): [yes]
HA-MIB (yes, y, no, n): [yes]
SW-TRAP (yes, y, no, n): [yes]
FA-TRAP (yes, y, no, n): [yes]
SW-EXTTRAP (yes, y, no, n): [no]
FICON-TRAP (yes, y, no, n): [yes]
HA-TRAP (yes, y, no, n): [yes]
no change
brocadessm:USERID>
```

2. Figure 25-6 shows an example of turning the Brocade SW-TRAP information on, the FICON MIB/TRAP information off, the HA MIB/TRAP information off, and the Fibre Alliance FA-TRAP information off.

Figure 25-6. Brocade SilkWorm (SW) Trap configuration with snmpmibcapset

```
brocadessm:USERID> snmpmibcapset
The SNMP Mib/Trap Capability has been set to support
FE-MIB SW-MIB FA-MIB FICON-MIB HA-MIB SW-TRAP FA-TRAP FICON-TRAP HA-TRAP
FA-MIB (yes, y, no, n): [yes] n
FICON-MIB (yes, y, no, n): [yes] n
HA-MIB (yes, y, no, n): [yes] n
SW-TRAP (yes, y, no, n): [yes]
FA-TRAP (yes, y, no, n): [yes] n
SW-EXTTRAP (yes, y, no, n): [no]
FICON-TRAP (yes, y, no, n): [yes] n
HA-TRAP (yes, y, no, n): [yes] n
brocadessm:USERID>
```

3. Figure 25-7 shows an example of turning the Fibre Alliance specific Trap information on and the Brocade Trap information off.

Figure 25-7. Fibre Alliance (FA) Trap configuration with snmpmibcapset

```
brocadessm:USERID> snmpmibcapset
The SNMP Mib/Trap Capability has been set to support
FE-MIB SW-MIB SW-TRAP
FA-MIB (yes, y, no, n): [no] y
FICON-MIB (yes, y, no, n): [no]
HA-MIB (yes, y, no, n): [no]
SW-TRAP (yes, y, no, n): [yes] n
FA-TRAP (yes, y, no, n): [no] y
SW-EXTTRAP (yes, y, no, n): [no]
FICON-TRAP (yes, y, no, n): [no]
HA-TRAP (yes, y, no, n): [no]
brocadessm:USERID>
```

Brocade SilkWorm (SW) Trap Detail

The Brocade SW-TRAP information can be obtained by opening up the SW MIB file "SW_v5_1.mib". Please refer to the *Brocade MIB Reference Manual* for more detail.

1. swFault - An *swFault* trap is generated whenever the diagnostics detects a fault with the switch.
   - Variables : *swDiagResult*
     swSsn – The SSN of the switch which sent this trap
     swGroupName – Specific to the SilkWorm 6400
     swGroupTyp e– Specific to the SilkWorm 6400
     swGroupMemPos – Specific to the SilkWorm 6400
2. swSensorScn - A *swSensorScn* trap is generated whenever an environment sensor changes its operational state.
   - Variables : *swSensorStatus* – The current status of the sensor
     swSensorIndex – The index of the sensor in the *sensorTable*
     swSensorType – The type (temperature, fan etc.) of the sensor
     swSensorValue – The reading from the sensor
     swSensorInfo – The type and number of the sensor in text format
     swSsn – The SSN of the switch which sent this trap
     swGroupName
     swGroupType
     swGroupMemPos

3. swFCPortScn - A *swFCPortScn* trap is generated whenever an FC_Port changes its operational state.
   - Variables : *swFCPortOpStatus* – The operational status of the Port (offline, online etc.)
     swFCPortIndex – The index of the port in the *swFCPortTable* (where index = port number + 1)
     swFCPortName – The port name identified by the portName command.
     swSsn – The SSN of the switch sending this trap.
     swGroupName
     swGroupType
     swGroupMemPos
4. swEventTrap - This trap is generated when an event whose level is at or below *swEventTrapLevel* occurs.
   Note : This acts like a filter to control the traps that may be sent by the agent whenever an event is written to the Error/Event Log. If this is set to 0, then the agent will send no event traps. If this is set to 5 then the agent will receive all levels of event traps.
   - Variables : *swEventIndex* – The index of the event in the *swEventTable*
     swEventTimeInfo – The time at which event happened
     swEventLevel – The severity level of the event
     swEventRepeatCount – The number of times this event was repeated
     swEventDescr – The description of the event
     swSsn – The SSN of the switch that sent this trap
     swGroupName
     swGroupType
     swGroupMemPos
5. swFabricWatchTrap - This trap is sent by Fabric Watch to notify of a Fabric Watch event..
   Note: Fabric Watch can further be configured to control the traps sent to an SNMP manager for each Class/Area. This feature acts as an additional filter for Fabric Watch traps sent out by the agent.
   - Variables : *swFwClassAreaIndex* – The Class/Area of the threshold
     swFwThresholdIndex – The index of the threshold in the *swFwThresholdTable*
     swFwName – The name of the threshold
     swFwLabel – The label of the threshold
     swFwLastEventVal – The last event value of the threshold
     swFwLastEventTime – The last event time of the threshold
     swFwLastEvent – The last event type of the threshold
     swFwLastState – The last event state of the threshold
     swSsn – The SSN of the switch that sent the trap.
     swGroupName
     swGroupType
     swGroupMemPos
6. swTrackChangesTrap - This trap gets sent when track changes is set to ON and is configured to send SNMP traps. A trap is sent whenever somebody logs in to the switch or logs out of the switch or when there are switch configuration changes.
   - Variables : *swTrackChangesInfo* – The description of the track changes event (config change, login, logout etc.)
     swSsn – The SSN of the switch that sent the trap.
     swGroupName
     swGroupType
     swGroupMemPos

Refer to the following table for the six traps defined in the SW-MIB, when they occur, and how to configure the trap, if possible.

Name	Specific	When	Configure
swFault	1	During boot, if diagnostics fail	Always on
swSensorScn	2	Sensor state change	Always on
swPortSsn	3	Port changes state	Always on
swEventTrap	4	Switch event	Command: **`agtCfgSet`** Variable: *swEventLevel*
swFabricWatch	5	Threshold reached	Command: **`fwConfigure`**
swTrackChanges	6	Login/Logout or Config change	Command: **`trackChangesSet`**

Note	The *swEventTrap*, can be configured using the **`agtCfgSet`** command. The other traps can not be configured independently and the SNMP management station will receive the other traps even if *swEventTrap* is turned off.

Fibre Alliance (FA) Trap Detail

The Fibre Alliance SNMP Traps can be found in the "FA_v3_0.mib" file, which provides Traps on the following information. Please refer to the *Brocade MIB Reference Manual* for more detail.

1. connUnitStatusChange - The overall status of the connectivity unit has changed.
   - Variables : *connUnitStatus* – status of the connection unit
     connUnitState - state of the connection unit
2. connUnitDeletedTrap - A connUnit has been deleted from this agent.
   - Variables : *connUnitID* – id of the connection unit
3. connUnitEventTrap - The connectivity unit has generated an event.
   - Variables : *connUnitEventId* - Internal event ID ranging between 0 and *connUnitMaxEvents* (now obsolete)
     connUnitEventType - The type of this event
     connUnitEventObject - This is used with the *connUnitEventType* to identify which object the event refers to.
     connUnitEventDescr - The description of the event
4. connUnitSensorStatusChange - The status of the sensor associated with the connectivity unit has changed.
   - Variables : *connUnitSensorStatus* - The status indicated by the sensor
5. connUnitPortStatusChange - The status of the sensor associated with the connectivity unit has changed.
   - Variables : *connUnitPortStatus* - An overall protocol status for the port
     connUnitPortState - The user selected state of the port hardware

Chapter 26

Fabric Watch

This section provides information to guide advance users to customize and administer Fabric Watch on the Brocade SAN Switch Module. The primary purpose of implementing Fabric Watch is to monitor the health status of the fabric elements by continuously ensuring that they are operating within the specified threshold boundaries. In the event a safe operating limit of an element is breached, an appropriate message is forwarded to the user by one or more pre-selected methods. The severity state appears in the message to indicate the urgency of the event to assist the administrator to take appropriate action. Fabric Watch accomplishes this in three steps by:

1. Measuring values

2. Comparing against threshold boundary limits

3. Event generation and notification

The Fabric Watch classes of elements are predefined for Brocade SilkWorm switches. The threshold levels for these element classes are provided in a default configuration. The default configuration saves valuable configuration setup time for a new user. Advanced users have the capability to fine-tune these thresholds to their unique fabric environment by choosing a customized configuration

The following information will be covered in this section:

- Fabric Watch Configuration
- Fabric Watch Setup
- Track Changes Setup
- Switch Status Policy Setup

Fabric Watch Configuration

Brocade Fabric Watch proactively monitors and reports on the health of the switches and the SAN fabric. Proactive monitoring and notification improves SAN availability. The real-time alerts from Fabric Watch software help SAN managers solve problems before they become costly failures. With Fabric Watch software, SAN managers can place limits, or thresholds, on the behavior of different switch and fabric elements.

A notable feature of Fabric Watch is that Fabric Watch messages can be configured by individual class elements to send events to an SNMP Manager. If the Brocade SW SNMP traps are turned on (as described in the SNMP chapter) you will receive messages from Fabric Watch for the class elements that have their Threshold Alarm Level configured to send SNMP traps. By default the Environmental classes within Fabric Watch are configured to send SNMP traps. For other classes one must configure the appropriate Threshold Alarm Level.

Fabric Watch *elements* are any fabric or switch component that the software monitors. To monitor elements, Fabric Watch categorizes them into *areas*, and further groups areas into *classes*. The following are the list of classes that fabric watch currently monitors and reports on. For more detail on any of these classes please read the most recent Fabric Watch User's Manual.

- ***Environment class*** - The environment class groups areas that deal with the physical environment inside a switch. Specifically, the environment class encompasses the ambient temperature of the switch, the speed of the fans within the switch, and the functionality and presence of power supplies. An environment class alert will alert you to problems or potential problems with temperature, fans and power.
- ***SFP class*** - The SFP class groups areas that monitor the physical aspects of smart transceivers. A SFP class alarm will alert you to faults that indicate that a transceiver may have deteriorated.
    - *SFP* – Small Form Factor Pluggable transceiver used in 2Gbit/sec switches.
      Note: SFP class monitors smart SFP(s) only.
- ***Fabric class*** - The fabric class groups areas that deal with potential problems that may arise between devices or other switches in a fabric. The fabric class includes the monitoring of Inter Switch Links (ISL), fabric reconfigurations, fabric segmentations, domain ID conflicts, zoning changes, and other related changes to a fabric. A fabric class alert will alert you to problems or potential problems with interconnectivity.

- ***Port class*** **-** The port class appears as the following three separate classes:
  - E-Port class
  - F/FL Port (Copper) class
  - F/FL Port (Optical) class
  - Multiple port classes let you set thresholds for different types of ports. Fabric Watch monitors the behavior of the port for state changes, link issues, quality of connection and optimal performance. You can configure separate and unique thresholds for E-Ports and for F/FL-Ports.
- ***Advanced Performance Monitor class*** - The performance monitor class appears as the following three separate classes:
  - ALPA Performance Monitor class
  - End-to-End Performance Monitor class
  - Filter Performance Monitor class
  - The advanced performance monitor classes serve as tuning tools. Advanced Performance monitor classes group areas that track the source and destination of traffic. You can use advanced performance monitor class thresholds and alarms to determine traffic load and flow and to reallocate resources appropriately.
- ***Security class*** - The security class monitors all attempts to breach your SAN security, helping you fine-tune your security measures.
- ***Switch Availability Monitor class*** - The switch availability monitor class monitors the efficiency of all active ports, providing a measure of switch availability. It provides statistics on switch downtime and uptime to help you identify problems with ports. Check the synchronization status of a port by running **`portShow`** and viewing the *portPhys* content.

Fabric Watch Setup

The Fabric Watch feature will already be loaded on your Brocade SAN Switch Module. This can be verified with the "**`licenseshow`**" command. See Figure 26-1.

Figure 26-1. Fabric Watch License.

```
brocadessm:USERID> licenseshow
SzdyQeSzRh0ezRj:
    Web license
    Zoning license
    Fabric license
    Fabric Watch license
```

The commands used to configure Fabric Watch are listed under the Fabric Watch help command "**`fwHelp`**". To determine if all the Fabric Watch classes are turned on use the "**`fwAlarmsFilterShow`**" and "**`fwAlarmsFilterSet`**" commands. By default, the Fabric Watch alarms are enabled on the Brocade SAN Switch Module. Fabric Watch has been pre-programmed with defaults from Brocade. If you know the specific traffic patterns of your SAN or have specific requirements for any of the class variables, Fabric Watch allows you the flexibility to tune the class alarms to your unique Storage environment.

1. To determine if all the Fabric Watch alarm classes are enabled use "**`fwAlarmsFilterShow`**". If the alarms are disabled, enable them using the "**`fwAlarmsFilterSet`**" command as show in Figure 26-2. This should not be necessary on the Brocade SAN Switch Modules as this has already been pre-programmed.

Figure 26-2. Enabling the Fabric Watch Classes.

```
switch:admin> fwalarmsfiltershow
FW: Alarms are disabled
switch:admin>
switch:admin> fwalarmsfilterset 1
Committing configuration...done.
FW: Alarms are enabled
switch:admin>
switch:admin> fwalarmsfiltershow
FW: Alarms are enabled
switch:admin>
```

2. The "*swFabricWatchTrap*" in the file named "SW_v5_1.mib" described in the SNMP section above, receives Fabric Watch notifications of an event, if the Threshold Alarm Level for a particular class includes the SNMP TRAP option. The **fwConfigure** telnet command can be used to configure Fabric Watch to control the traps that Fabric Watch can send. SNMP traps for each Class/Area can be turned ON or OFF. This effectively acts like a filter for Fabric Watch traps sent out by the agent.

3. Figure 26-3, shows an example of the "**fwConfigure**" command. This example shows that the Value of Threshold Alarm Level is "3" for Temperature, in the Environmental Class. This means that a Fabric Watch event is sent to the Error/Event Log and to an SNMP management station. The other possible options are indicated as well:

   1 – Errlog (sent to the Switch Error/Event Log)
   2 – SnmpTrap (sent to an SNMP management station)
   4 – PortLogLock (locks the portlog if an event is triggered)
   8 – RapiTrap (sent to the Brocade API)
   16 – EmailAlert (sent to an Email recipient, configured using "**fwMailCfg**")

The Threshold Alarm Level is configured using a matrix value with a maximum value of 31. Combine any of the option numbers above to choose the value you wish to use. (For example, the value of "3" is Option 1 and Option 2. So, a message gets sent to the Error/Event Log and to an SNMP Trap.)

Figure 26-3. Example of the "fwConfigure command, with Threshold Alarm Levels at "3".

```
brocadessm:USERID> fwconfigure

1  : Environment class
2  : SFP class
3  : Port class
4  : Fabric class
5  : E-Port class
6  : F/FL Port (Copper) class
7  : F/FL Port (Optical) class
8  : Alpa Performance Monitor class
9  : EE Performance Monitor class
10 : Filter Performance Monitor class
11 : Security class
12 : Switch Availability Monitor class
13 : Quit
Select a class => : (1..13) [13] 1
1 : Temperature
2 : Fan
3 : Power Supply
4 : return to previous page
Select an area => : (1..4) [4] 1

Index ThresholdName                    Status       CurVal
          LastEvent               LasteventTime      LastVal       LastState
================================================================================
    1    envTemp001                    enabled         35 C
          inBetween    Fri Apr 16 16:49:38 2004         34 C     Normal
    2    envTemp002                    enabled         47 C
          inBetween    Fri Apr 16 16:49:38 2004         47 C     Normal
    3    envTemp003                    enabled         47 C
          inBetween    Fri Apr 16 16:49:38 2004         47 C     Normal
    4    envTemp004                    enabled         38 C
          inBetween    Fri Apr 16 16:49:38 2004         38 C     Normal

1 : refresh
2 : disable a threshold
3 : enable a threshold
4 : advanced configuration
```

```
5 : return to previous page
Select choice => : (1..5) [5] 4

Index ThresholdName    BehaviorType    BehaviorInt
    1    envTemp001       Triggered              1
    2    envTemp002       Triggered              1
    3    envTemp003       Triggered              1
    4    envTemp004       Triggered              1

Threshold boundary level is set at : Default

                 Default        Custom
      Unit             C             C
 Time base
       Low             0             0
      High            71            71
   BufSize            10            10

Threshold alarm level is set at : Default

Errlog-1, SnmpTrap-2, RapiTrap-8
EmailAlert-16

Valid alarm matrix is 27

                 Default        Custom
   Changed             0             0
  Exceeded             0             0
     Below             3             3
     Above             3             3
 InBetween             3             3

1  : change behavior type             11 : change threshold alarm level
2  : change behavior interval         12 : change changed alarm
3  : change threshold boundary level  13 : change exceeded alarm
4  : change custom unit               14 : change below alarm
5  : change custom time base          15 : change above alarm
6  : change custom low                16 : change inBetween alarm
7  : change custom high               17 : apply threshold alarm changes
8  : change custom buffer             18 : cancel threshold alarm changes
9  : apply threshold boundary changes 19 : return to previous page
10 : cancel threshold boundary changes
Select choice => : (1..19) [19]
```

4. Figure 26-4, shows an example using the "**fwConfigure**" command to change the value of the Threshold Alarm Level to "3" for RX Performance, in the E-Port Class.
   Note : These examples can also be done through the Web Tools interface which simplifies the process. The defaults are shown in orange and the values they are changed to are in violet.

Figure 26-4. Example using "fwConfigure" to change the Threshold Alarm Levels to "3".

```
brocadessm:USERID> fwconfigure

1  : Environment class
2  : SFP class
3  : Port class
4  : Fabric class
5  : E-Port class
6  : F/FL Port (Copper) class
7  : F/FL Port (Optical) class
8  : Alpa Performance Monitor class
9  : EE Performance Monitor class
10 : Filter Performance Monitor class
11 : Security class
12 : Switch Availability Monitor class
13 : Quit
Select a class => : (1..13) [13] 5

1  : Link loss
2  : Sync loss
3  : Signal loss
4  : Protocol error
5  : Invalid words
6  : Invalid CRCS
7  : RXPerformance
8  : TXPerformance
9  : State Changes
10 : return to previous page
Select an area => : (1..10) [10] 7

Index ThresholdName                        Status       CurVal
          LastEvent                LasteventTime      LastVal       LastState
================================================================================
    0 eportRXPerf000                       enabled          0 KB/s
              below    Fri Apr 16 16:49:38 2004          0 KB/s     Informative
   15 eportRXPerf015                       enabled          0 KB/s
              below    Fri Apr 16 16:49:38 2004          0 KB/s     Informative

1 : refresh
2 : disable a threshold
3 : enable a threshold
4 : advanced configuration
5 : return to previous page
Select choice => : (1..5) [5] 4

Index ThresholdName    BehaviorType    BehaviorInt
    0 eportRXPerf000       Triggered              1
   15 eportRXPerf015       Triggered              1

Threshold boundary level is set at : Default

                  Default        Custom
      Unit           KB/s          KB/s
 Time base
       Low         120000        120000
      High         220000        220000
   BufSize              0             0

Threshold alarm level is set at : Default

Errlog-1, SnmpTrap-2, PortLogLock-4
RapiTrap-8, EmailAlert-16

Valid alarm matrix is 31
```

```
                Default        Custom
   Changed            0             0
  Exceeded            0             0
     Below            0             0
     Above            0             0
 InBetween            0             0

1  : change behavior type             11 : change threshold alarm level
2  : change behavior interval         12 : change changed alarm
3  : change threshold boundary level  13 : change exceeded alarm
4  : change custom unit               14 : change below alarm
5  : change custom time base          15 : change above alarm
6  : change custom low                16 : change inBetween alarm
7  : change custom high               17 : apply threshold alarm changes
8  : change custom buffer             18 : cancel threshold alarm changes
9  : apply threshold boundary changes 19 : return to previous page
10 : cancel threshold boundary changes
Select choice => : (1..19) [19] 14

Errlog-1, SnmpTrap-2, PortLogLock-4
RapiTrap-8, EmailAlert-16

Valid alarm matrix is 31
Enter below alarm matrix => : (0..31) [0] 3

Index ThresholdName    BehaviorType    BehaviorInt
    0 eportRXPerf000       Triggered              1
   15 eportRXPerf015       Triggered              1

Threshold boundary level is set at : Default

                Default        Custom
      Unit         KB/s          KB/s
 Time base
       Low       120000        120000
      High       220000        220000
   BufSize            0             0

Threshold alarm level is set at : Default

Errlog-1, SnmpTrap-2, PortLogLock-4
RapiTrap-8, EmailAlert-16

Valid alarm matrix is 31

                Default        Custom
   Changed            0             0
  Exceeded            0             0
     Below            0             3
     Above            0             0
 InBetween            0             0

1  : change behavior type             11 : change threshold alarm level
2  : change behavior interval         12 : change changed alarm
3  : change threshold boundary level  13 : change exceeded alarm
4  : change custom unit               14 : change below alarm
5  : change custom time base          15 : change above alarm
6  : change custom low                16 : change inBetween alarm
7  : change custom high               17 : apply threshold alarm changes
8  : change custom buffer             18 : cancel threshold alarm changes
9  : apply threshold boundary changes 19 : return to previous page
10 : cancel threshold boundary changes
Select choice => : (1..19) [19] 15
Errlog-1, SnmpTrap-2, PortLogLock-4
RapiTrap-8, EmailAlert-16

Valid alarm matrix is 31
Enter above alarm matrix => : (0..31) [0] 3
```

```
Index ThresholdName    BehaviorType    BehaviorInt
    0 eportRXPerf000       Triggered              1
   15 eportRXPerf015       Triggered              1

Threshold boundary level is set at : Default

                 Default        Custom
      Unit          KB/s          KB/s
 Time base
       Low        120000        120000
      High        220000        220000
   BufSize             0             0

Threshold alarm level is set at : Default

Errlog-1, SnmpTrap-2, PortLogLock-4
RapiTrap-8, EmailAlert-16

Valid alarm matrix is 31

                 Default        Custom
   Changed             0             0
  Exceeded             0             0
     Below             0             3
     Above             0             3
 InBetween             0             0

1  : change behavior type             11 : change threshold alarm level
2  : change behavior interval         12 : change changed alarm
3  : change threshold boundary level  13 : change exceeded alarm
4  : change custom unit               14 : change below alarm
5  : change custom time base          15 : change above alarm
6  : change custom low                16 : change inBetween alarm
7  : change custom high               17 : apply threshold alarm changes
8  : change custom buffer             18 : cancel threshold alarm changes
9  : apply threshold boundary changes 19 : return to previous page
10 : cancel threshold boundary changes
Select choice => : (1..19) [19] 11
1  : Default
2  : custom
Enter alarm level type => : (1..2) [1] 2

Index ThresholdName    BehaviorType    BehaviorInt
    0 eportRXPerf000       Triggered              1
   15 eportRXPerf015       Triggered              1

Threshold boundary level is set at : Default

                 Default        Custom
      Unit          KB/s          KB/s
 Time base
       Low        120000        120000
      High        220000        220000
   BufSize             0             0

Threshold alarm level is set at : Custom

Errlog-1, SnmpTrap-2, PortLogLock-4
RapiTrap-8, EmailAlert-16

Valid alarm matrix is 31

                 Default        Custom
   Changed             0             0
  Exceeded             0             0
     Below             0             3
     Above             0             3
 InBetween             0             0
```

```
1  : change behavior type             11 : change threshold alarm level
2  : change behavior interval         12 : change changed alarm
3  : change threshold boundary level  13 : change exceeded alarm
4  : change custom unit               14 : change below alarm
5  : change custom time base          15 : change above alarm
6  : change custom low                16 : change inBetween alarm
7  : change custom high               17 : apply threshold alarm changes
8  : change custom buffer             18 : cancel threshold alarm changes
9  : apply threshold boundary changes 19 : return to previous page
10 : cancel threshold boundary changes
Select choice => : (1..19) [19] 17

Index ThresholdName    BehaviorType    BehaviorInt
    0 eportRXPerf000       Triggered              1
   15 eportRXPerf015       Triggered              1

Threshold boundary level is set at : Default

                 Default        Custom
      Unit          KB/s          KB/s
 Time base
       Low        120000        120000
      High        220000        220000
   BufSize             0             0

Threshold alarm level is set at : Custom

Errlog-1, SnmpTrap-2, PortLogLock-4
RapiTrap-8, EmailAlert-16

Valid alarm matrix is 31

                 Default        Custom
   Changed             0             0
  Exceeded             0             0
     Below             0             3
     Above             0             3
 InBetween             0             0

1  : change behavior type             11 : change threshold alarm level
2  : change behavior interval         12 : change changed alarm
3  : change threshold boundary level  13 : change exceeded alarm
4  : change custom unit               14 : change below alarm
5  : change custom time base          15 : change above alarm
6  : change custom low                16 : change inBetween alarm
7  : change custom high               17 : apply threshold alarm changes
8  : change custom buffer             18 : cancel threshold alarm changes
9  : apply threshold boundary changes 19 : return to previous page
10 : cancel threshold boundary changes
Select choice => : (1..19) [19]
```

These configuration settings are now saved to the switch and will be visible in the configshow output under the "thresh" variables, as shown in Figure 26-5.

Figure 26-5. Fabric Watch parameters in `configshow` output.

```
brocadessm:USERID> configshow "thresh"
thresh.cust.eport.RXPerf.above:3
thresh.cust.eport.RXPerf.below:3
thresh.eportRXPerf.actLevel:2
```

The WebTools, Fabric Watch interface can also accomplish the same set of tasks. Use the Fabric Watch button from the Main WebTools switch view to enter the Fabric Watch configuration applet. Figure 26-6.

Figure 26-6. Example using "fwConfigure" to change the Threshold Alarm Levels to "3".

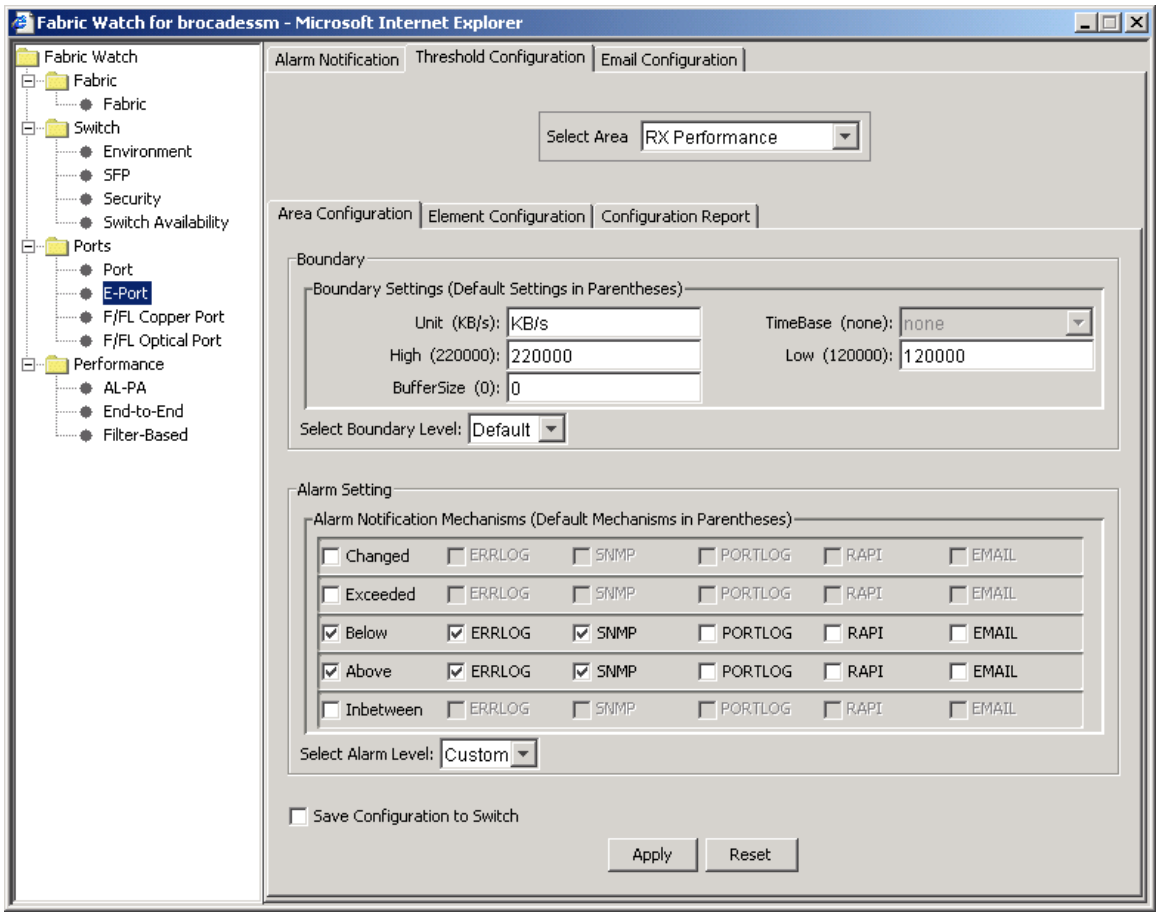

Guideline	Even though E-mail notification can be specified for any class element, it is recommended that it should be set to receive notification only for critical alerts that require immediate attention. For example, instances like fabric re-configurations, fabric segmentations, switch failures or security violations. The remaining informational and warning messages can be left directed to the system Error log.

Track Changes Setup

As an added security monitoring feature Brocade switches have the ability to track the occurrence of an administrator logging into the switch or logging out of a switch and whether a switches configuration has changed.

The commands used to configure the Track Changes feature are listed under the Track Changes help command "**trackChangesHelp**". To determine if this feature is turned on use the "**trackChangesShow**" command. The Track Changes feature can also send its event messages to SNMP. To configure Track Changes and send SNMP Traps when events occur use the "**trackChangesSet**" command. The "*swTrackChangesTrap*" in the file named "SW_v5_1.mib" described in the SNMP section above, receives Track Change notifications of an event.

Figure 26-7 shows an example of turning on the Track Changes feature. This feature can only be configured from the CLI.

Figure 26-7. Configuring Track Changes to send SNMP Traps.

```
brocadessm:USERID> trackchangesshow
Track changes status: OFF
Track changes generate SNMP-TRAP: NO
brocadessm:USERID>
brocadessm:USERID> trackchangesset 1,1
Committing configuration...done.
brocadessm:USERID>
brocadessm:USERID> trackchangesshow
Track changes status: ON
Track changes generate SNMP-TRAP: YES
brocadessm:USERID>
```

Switch Status Policy Setup

The Brocade Switch Status Policy feature monitors the overall status of the switch based on several contributing parameters. The policy parameter values determine how many failed or faulty units of each contributor are allowed before triggering a status change in the switch from HEALTHY to MARGINAL or DOWN. These status events are integrated into Brocade Web Tools and Fabric Manager so that if the overall status of the switch is HEALTHY the switch color is Green. If the overall switch status is MARGINAL the switch color is Yellow. Finally if the overall switch status is DOWN the switch color is RED. The overall status is calculated based on the most severe status of all contributors. For the Brocade SAN Switch Module the following are monitored by this function:

- Faulty Ports – triggers if the port goes faulty
- Missing GBICs/SFPs – triggers if there are missing GBICs/SFPs
- Temperatures – triggers if a temperature sensor goes out of its Fabric Watch range
- Port Status – triggers if a port changes its status (i.e. online, offline etc.)
- ISL Status – triggers if a certain number of ISLs go down
- Internal Switch Status – triggers if the switch is disabled or enabled

The overall status can be one of the following:

- Healthy/OK - every contributor is healthy
- Marginal/Warning - one or more components are causing a warning status
- Down/Failed - one or more contributors have failed

The commands used to configure the Switch Status Policy feature are "**switchStatusPolicyShow**" and "**switchStatusPolicySet**". The "*swEventTrap*" in the file named "SW_v5_1.mib" described in the SNMP

section above, receives Switch Status Policy notifications of an event because they are logged to the Error/Event Log. This feature can only be configured from the CLI.

Figure 26-8 shows the default parameter settings for the Switch Status Policy feature for the Brocade SAN Switch Module, other Brocade platforms may vary in function.

Note	Marginal indicates the switch will turn yellow. Down indicates the switch will turn red. If the value of a policy parameter is set to "0", it means that this factor is not used to determine the status of the switch. If both of the values are "0" that contributor is not used in the calculation of the overall switch status.

Figure 26-8. Switch Status Policy settings.

```
brocadessm:USERID> switchstatuspolicyshow
The current overall switch status policy parameters:
                   Down    Marginal
 ---------------------------------
       FaultyPorts  2         1
       MissingSFPs  0         0
      Temperatures  2         1
        PortStatus  0         0
         ISLStatus  0         0
brocadessm:USERID>
```

Figure 26-9 shows the configuration of the Switch Status Policy feature.

Figure 26-9. Switch Status Policy configuration.

```
brocadessm:USERID> switchstatuspolicyset

To change the overall switch status policy parameters

The current overall switch status policy parameters:
                    Down    Marginal
 ---------------------------------
       FaultyPorts  2         1
       MissingSFPs  0         0
      Temperatures  2         1
        PortStatus  0         0
         ISLStatus  0         0

Note that the value, 0, for a parameter, means that it is
NOT used in the calculation.
** In addition, if the range of settable values in the prompt is (0..0),
** the policy parameter is NOT applicable to the switch.
** Simply hit the Return key.


The minimum number of
  FaultyPorts contributing to
                      DOWN status: (0..16) [2]
  FaultyPorts contributing to
                      MARGINAL status: (0..16) [1]
  MissingSFPs contributing to
                      DOWN status: (0..16) [0]
  MissingSFPs contributing to
                      MARGINAL status: (0..16) [0]
  Bad Temperatures contributing to
                      DOWN status: (0..4) [2]
  Bad Temperatures contributing to
                      MARGINAL status: (0..4) [1]
  Down PortStatus contributing to
                      DOWN status: (0..16) [0]
  Down PortStatus contributing to
                      MARGINAL status: (0..16) [0]
  Down ISLStatus contributing to
                      DOWN status: (0..16) [0]
  Down ISLStatus contributing to
                      MARGINAL status: (0..16) [0]
No change
brocadessm:USERID>
```

Figure 26-10 shows two sample outputs of "switchstatusshow". The first shows a Healthy condition the second reveals error conditions have triggered for several contributors.

Figure 26-10. Switch Status Policy messages.

```
brocadessm:USERID> switchstatusshow
The overall switch status is HEALTHY/OK
brocadessm:USERID>


brocadessm:USERID> switchstatusshow
The overall switch status is Marginal/Warning
  Contributing factors:
   *  Switch Offline triggered the Marginal/Warning status

brocadessm:USERID>
```

Chapter 27

Advanced Performance Monitoring

Brocade Advanced Performance Monitoring is a tool which helps SAN managers measure the efficiency of their SAN resources. Trends and patterns can be read using charts and graphs.

The overall throughput performance of an application is entirely dependent on the configuration of the host and storage hardware being utilized. A SAN configuration providing connectivity to host and storage devices must not be a performance factor. Ideally, a storage port providing connectivity to one or more hosts must be able to satisfy the aggregated bandwidth requirement of all the hosts. However if not carefully planned, a congestion condition may exist when a port is over-subscribed. Advanced Performance Monitoring helps identify and correct this over-provisioning condition. Other SAN management areas which can benefit from the monitoring system

- Capacity Planning
- SAN Performance tuning
- End-to-end visibility into fabric
- Increasing productivity via Pre-formatted Reports

The following sections discuss:

- WebTools Performance Graphs
- Configuring Advanced Performance Monitoring

WebTools Performance Graphs

WebTools Performance Monitoring allows you to set up a canvas of performance graphs from the **Action** menu.

- An existing graph can be selected from a list of graphs that are predefined or in some cases can be customized to monitor the specific objects of the fabric.
- The canvas window displaying the graph can hold up to a maximum of eight graphs simultaneously. Any graph can be magnified, added or removed from the main canvas.
- Up to 20 individual canvases can be saved for later retrieval.
- Graphs can be printed.

Configuring Advanced Performance Monitoring

The Advanced Performance Monitoring license must be loaded on your Brocade SAN Switch Module. This can be verified with the "**licenseshow**" command. See Figure 27-1.

Figure 27-1. Advanced Performance Monitoring License.

```
brocadessm:USERID> licenseshow
SzdyQeSzRh0ezRj:
    Web license
    Zoning license
    Fabric license
    Fabric Watch license
    Performance Monitor license
```

Although there are CLI commands to configure Advanced Performance Monitoring, the WebTools interface provides some easy to use pre-defined and customizable graphs to get you up and running quickly.
Start by clicking on the Performance Monitoring icon "Perf" from the Main Switch Explorer window of Web Tools for the Brocade SAN Switch Module. Figure 27-2.

Figure 27-2. Advanced Performance Monitoring icon on Switch Explorer.

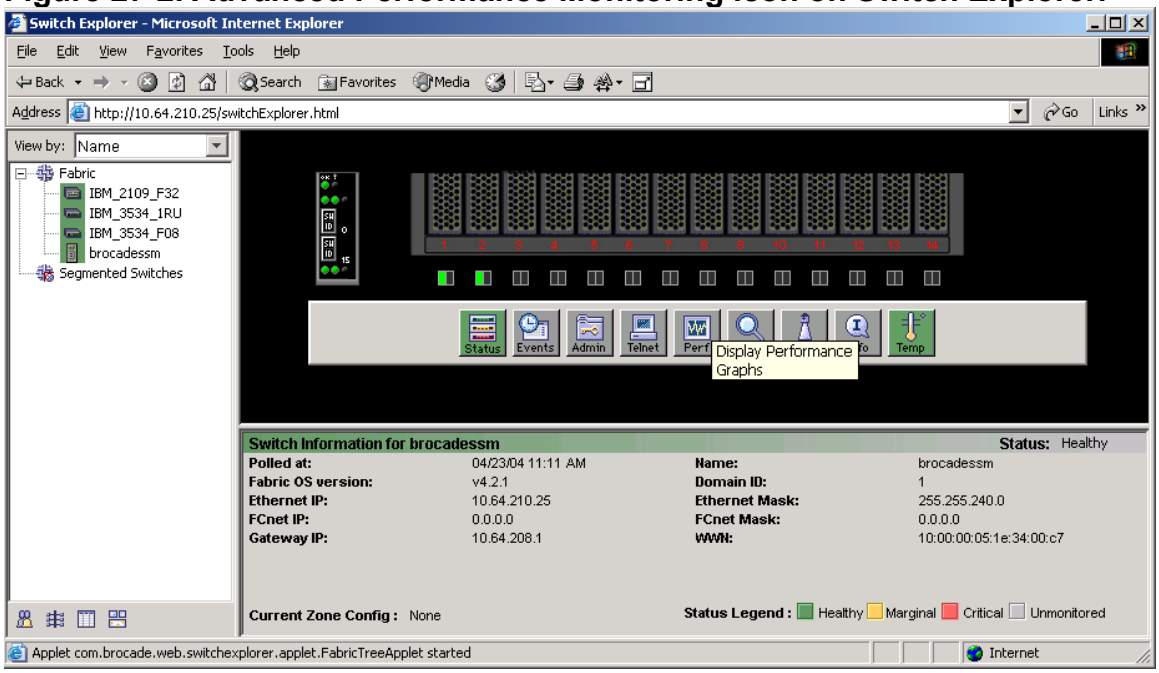

The following default switch performance Monitoring Graph appears. Figure 27-3.

Figure 27-3. Default Performance Monitoring Graph.

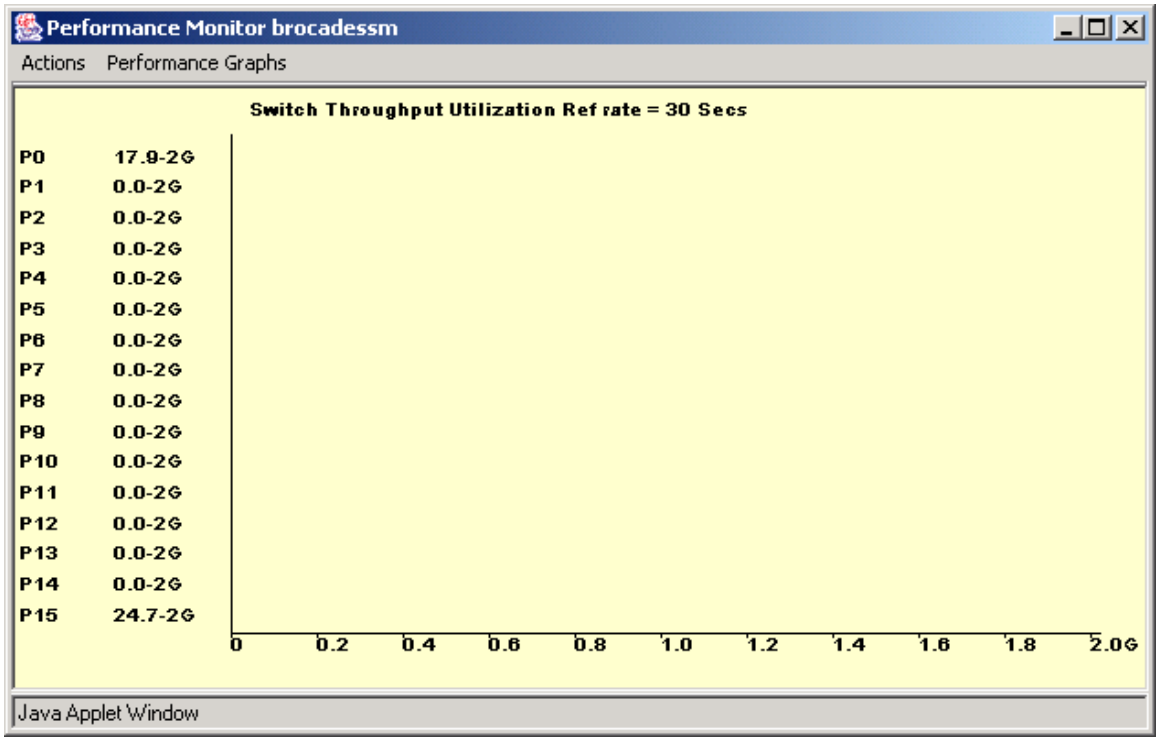

The following example will add two new Advanced Performance Monitoring Graphs that will display the Transmit and Receive Performance from each internal Server Blade in a BladeCenter (2 HS20 Server Blades) that are each communicating with one Fibre Channel Disk Drive from an FC JBOD.

Step 1. Use the Name Server icon on the bottom left half of the Switch Explorer WebTools window to access a list of devices in the fabric. Figure 27-4.

Figure 27-4. Name Server icon in WebTools.

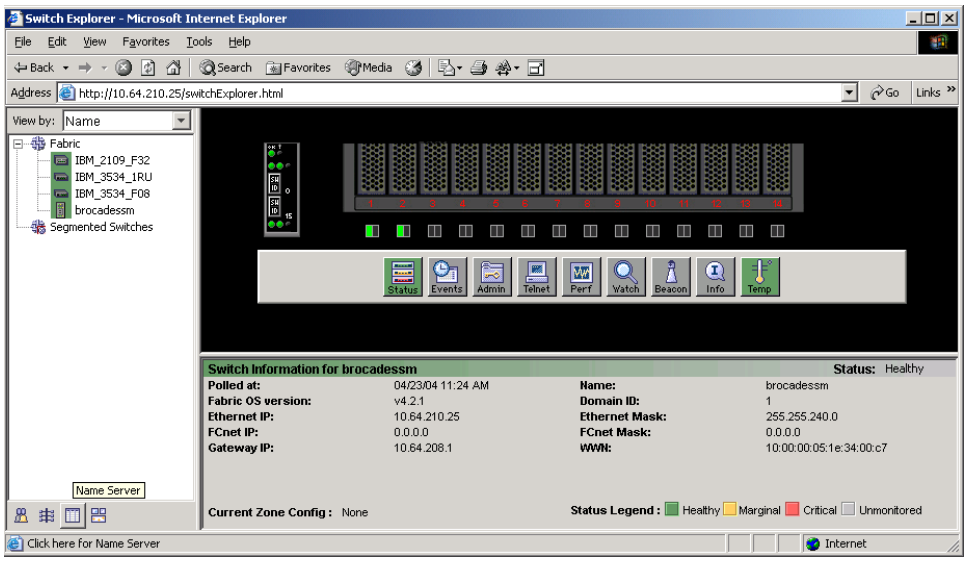

Step 2. Sort the entries by Port and then Domain. Then look to the "Port ID" column to obtain the Source ID's (SID) and Destination ID's (DID) of the devices you wish to monitor. In this example the two Server Blades are SID_1=010100 (Port 1 on Domain 1) and SID_2=010200 (Port 2 on Domain 1). The Fibre Channel JBOD drives are located on the same port, Domain 3 Port 7. DID_1=0307ba (the Port ID of the disk drive that SID_1 is communicating with) and DID_2=0307bc (the Port ID of the disk drive that SID_2 is communicating with). Figure 27-5.

Figure 27-5. Name Server icon in WebTools.

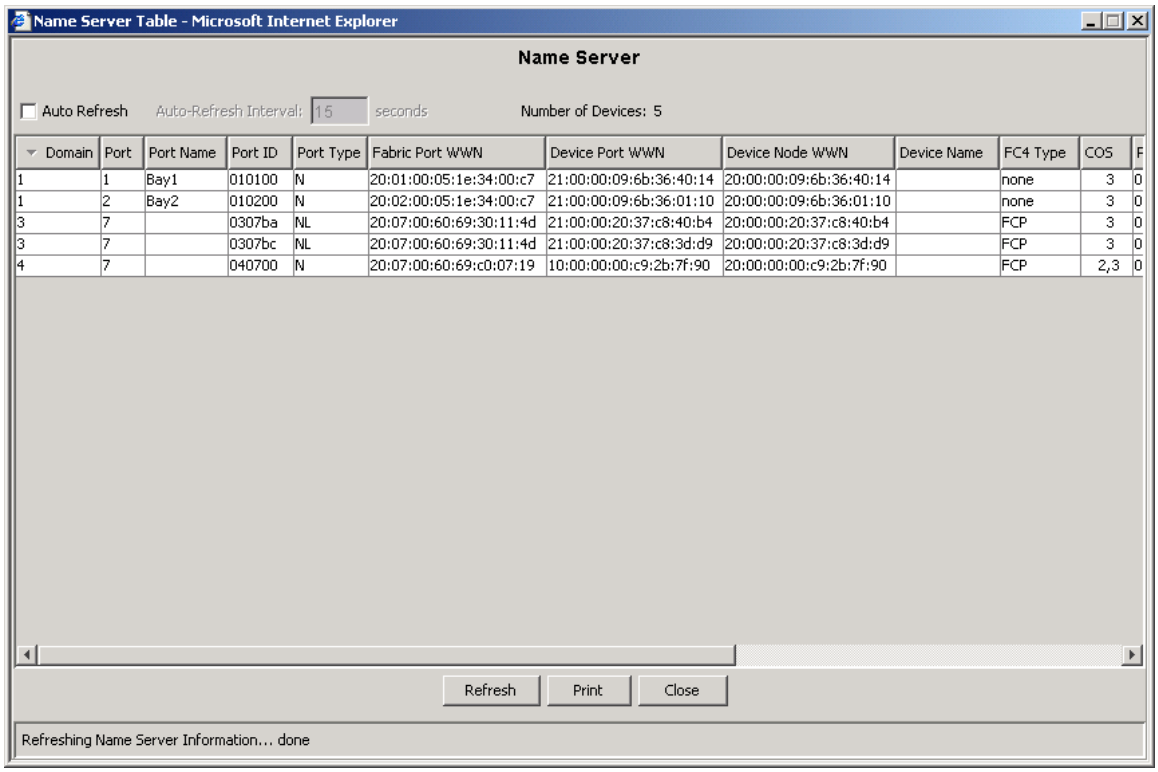

Domain	Port	Port Name	Port ID	Port Type	Fabric Port WWN	Device Port WWN	Device Node WWN	Device Name	FC4 Type	COS	F
1	1	Bay1	010100	N	20:01:00:05:1e:34:00:c7	21:00:00:09:6b:36:40:14	20:00:00:09:6b:36:40:14		none	3	0
1	2	Bay2	010200	N	20:02:00:05:1e:34:00:c7	21:00:00:09:6b:36:01:10	20:00:00:09:6b:36:01:10		none	3	0
3	7		0307ba	NL	20:07:00:60:69:30:11:4d	21:00:00:20:37:c8:40:b4	20:00:00:20:37:c8:40:b4		FCP	3	0
3	7		0307bc	NL	20:07:00:60:69:30:11:4d	21:00:00:20:37:c8:3d:d9	20:00:00:20:37:c8:3d:d9		FCP	3	0
4	7		040700	N	20:07:00:60:69:c0:07:19	10:00:00:00:c9:2b:7f:90	20:00:00:00:c9:2b:7f:90		FCP	2,3	0

This same data can be obtained from the “Device Ports” table in Fabric Manager. Figure 27-6.

Figure 27-6. Device Ports Table in Fabric Manager.

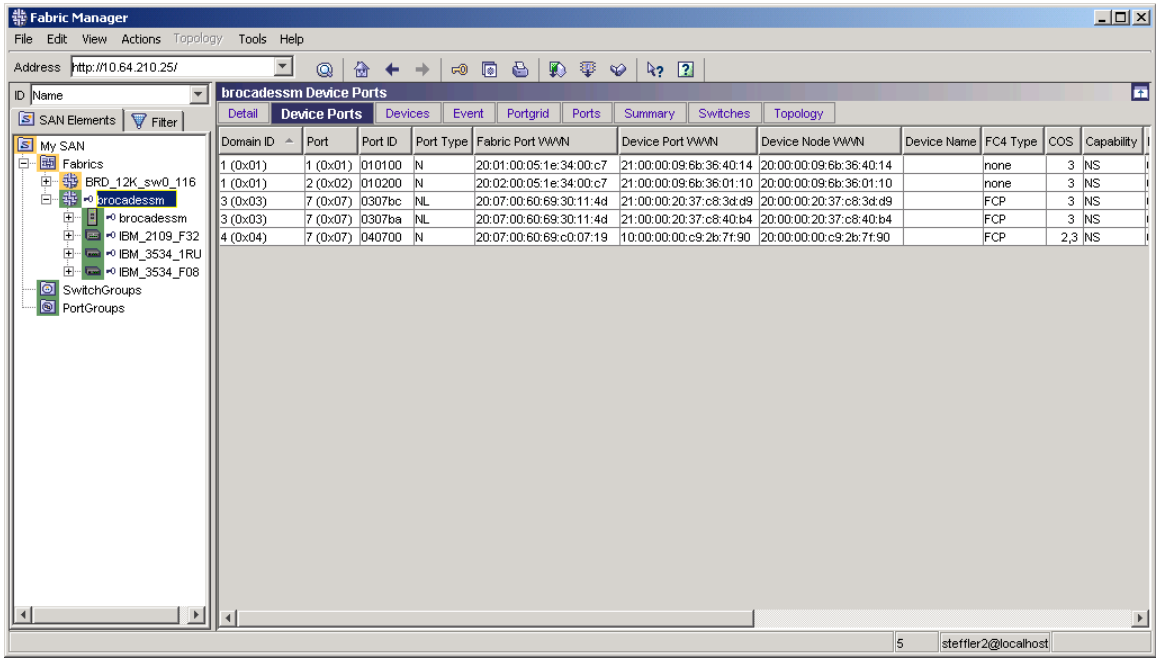

Step 4. From the WebTools Performance Monitor window locate the SID/DID performance option in the Advanced Monitoring drop down list of the Performance Graphs Tab. Figure 27-7.

Figure 27-7. SID/DID Performance Graphs in WebTools.

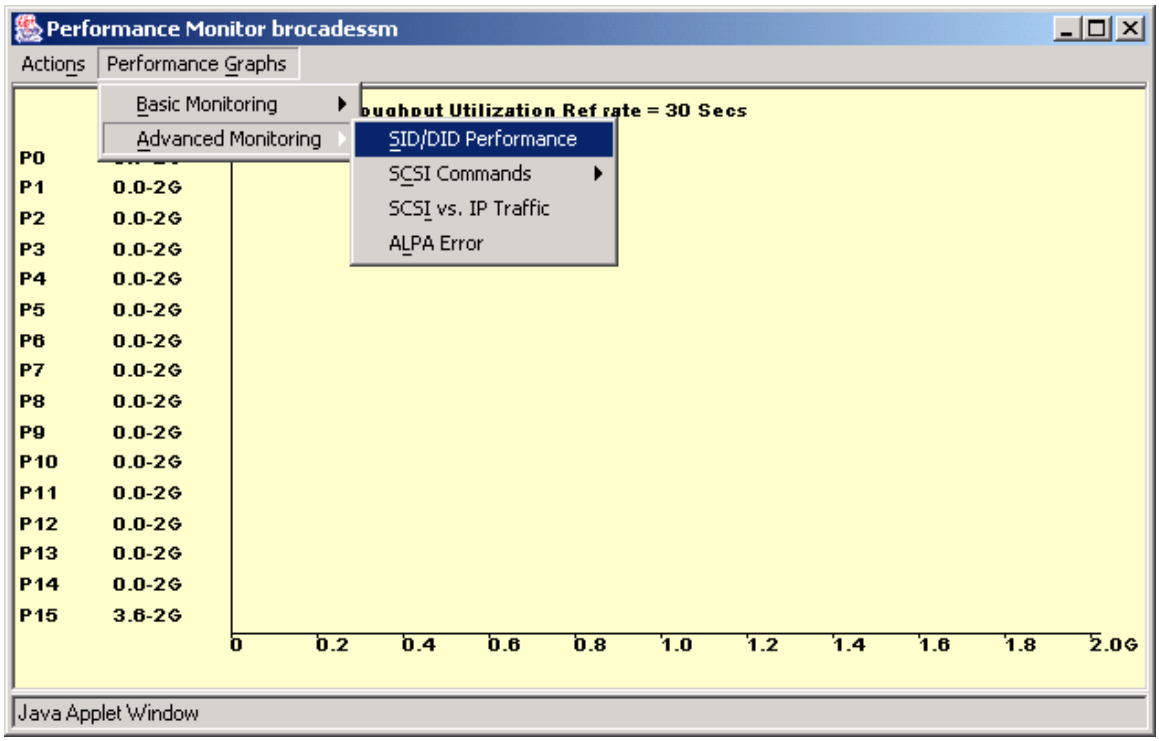

Step 5. The following SID/DID Performance Setup window appears. Figure 27-8.

Figure 27-8. SID/DID Performance Graphs in WebTools.

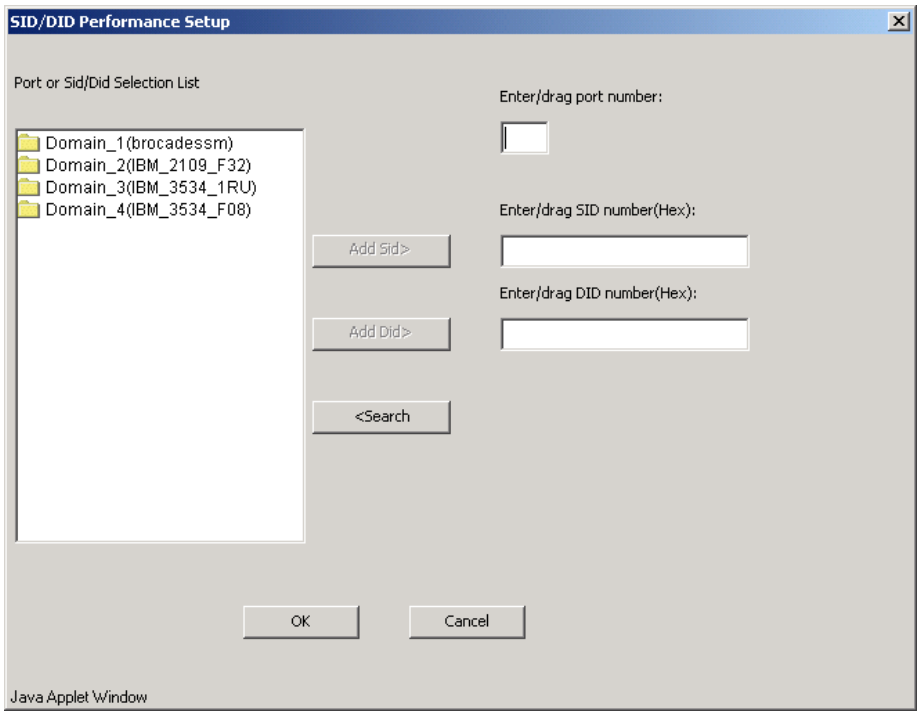

Step 6. Maneuver through the Port Selection List on the left to Port 1 (IBM eServer BladeCenter Server Blade in bay 1) on Domain 1 (the Brocade SAN Switch Module) and find the SID for that device. Figure 27-9.

Figure 27-9. SID "010100" in Port Selection List.

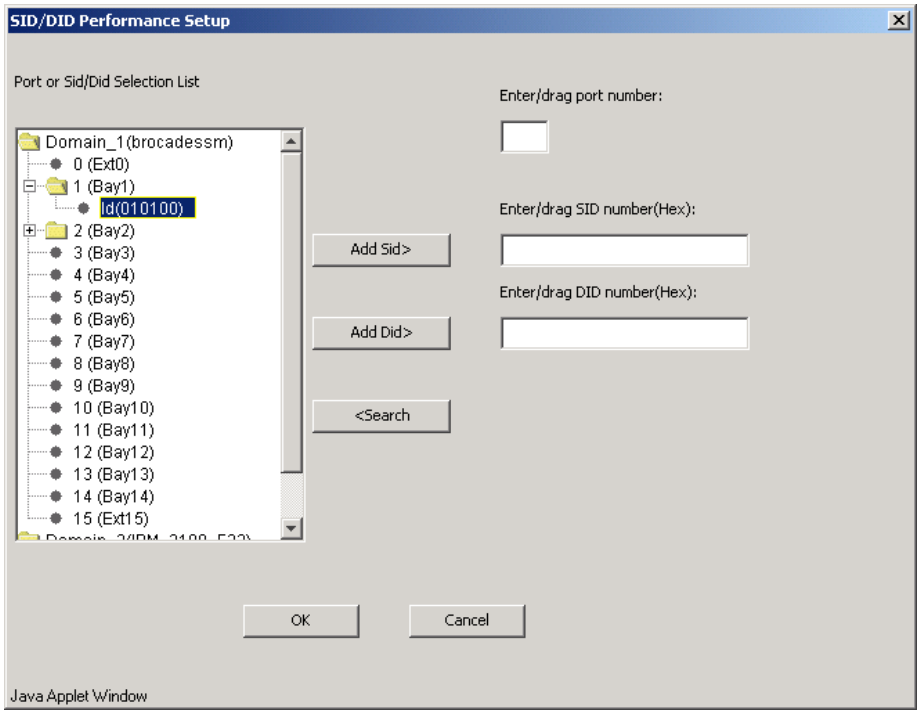

Step 7. Enter or drag and drop Port Number 1 "*1 (Bay1)*" into the "Enter/drag port number box. The box should have a "1" in it. Enter or drag and drop the SID number by clicking on the "*Id(010100)*" bullet to highlight it and then click on the "*Add Sid>*" button in the middle to place the value of "010100" into the SID box. Figure 27-10.

Figure 27-10. Place SID "010100" into the SID box.

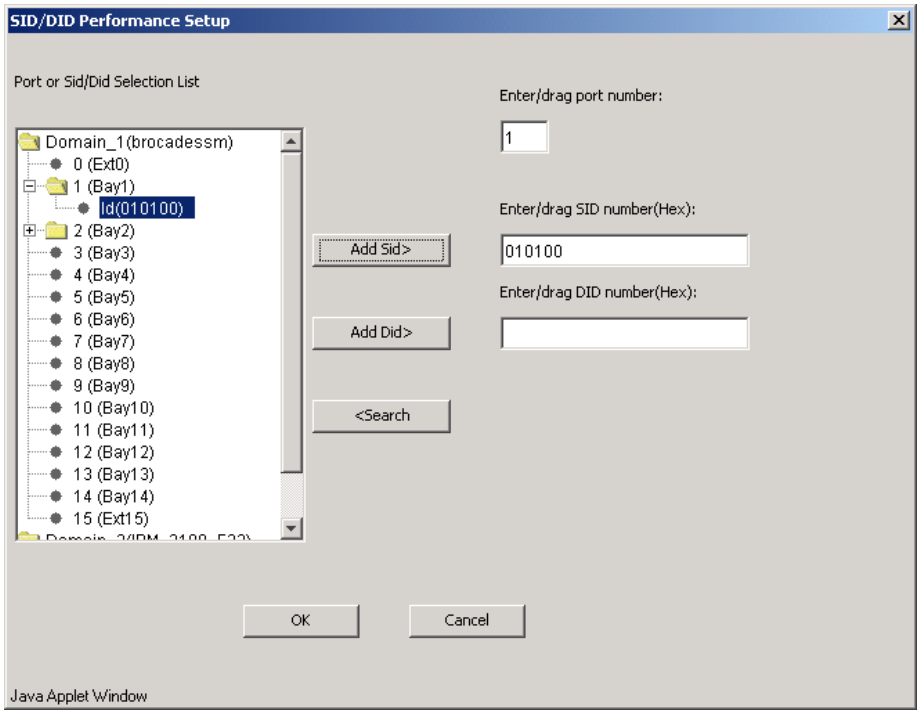

Step 8. Repeat 7 by finding the Destination ID the Disk Drive or by placing the value of "0307ba" into the DID box. Figure 27-11.

Figure 27-11. Place DID "0307ba" into the DID box.

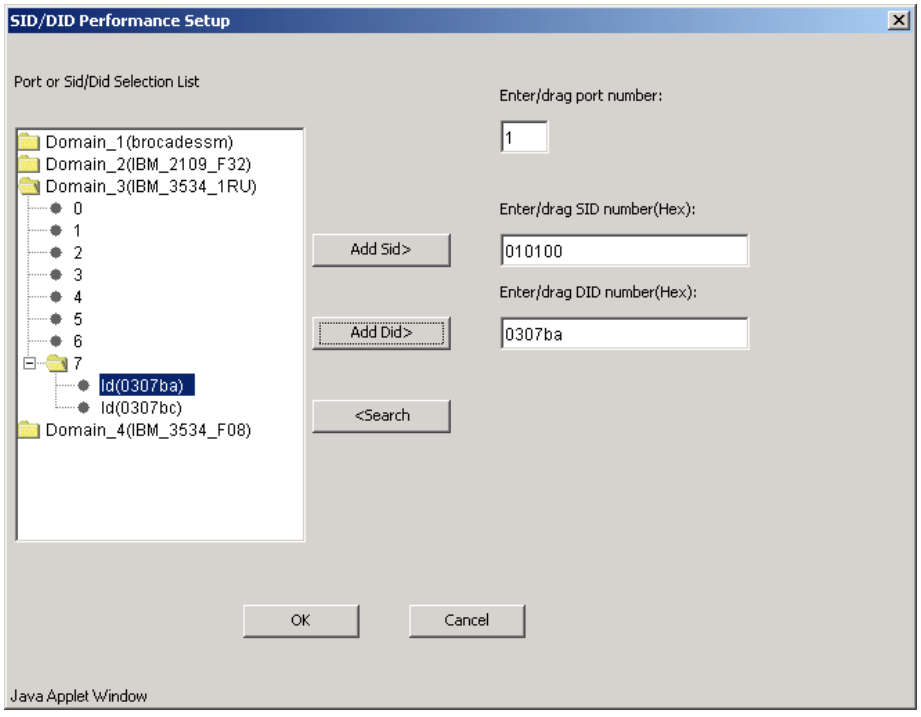

Step 9. Click the “OK” button and the new graph will get added to the existing canvas. Figure 27-12.

Figure 27-12. New Advanced Performance Monitoring graph is added.

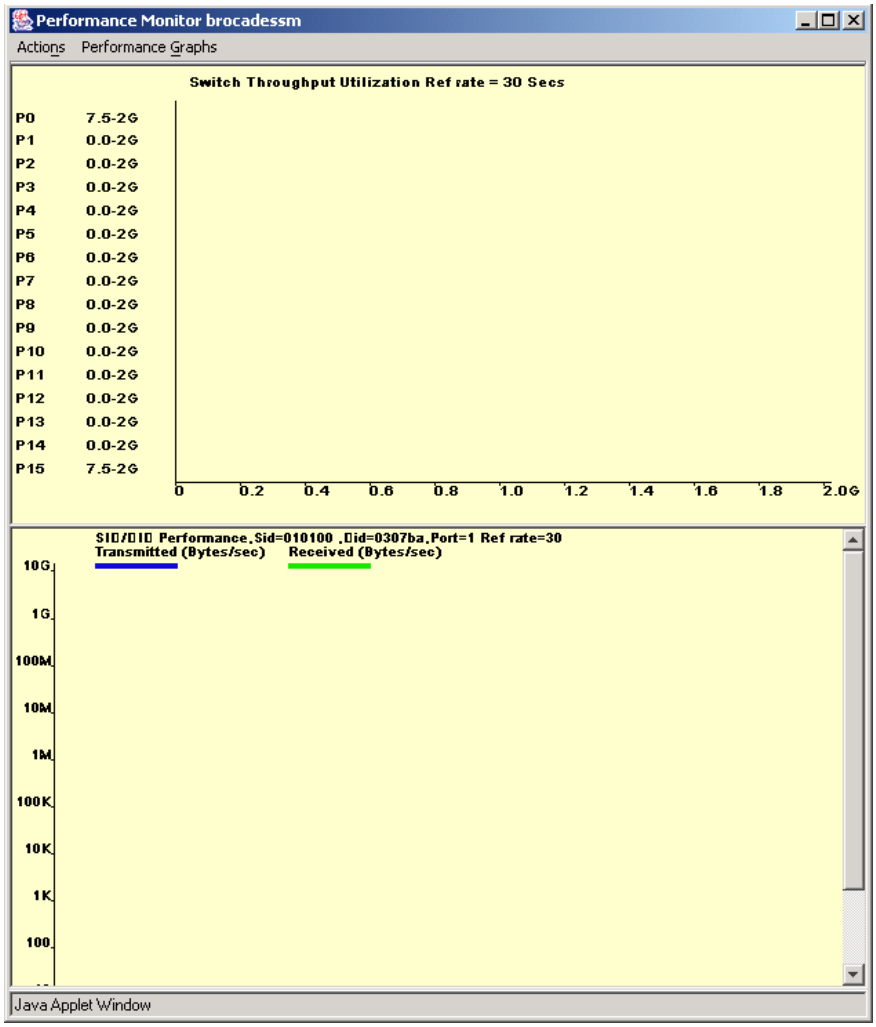

Step 10. Repeat Steps 4-8 to add a second Advanced Performance Monitoring Filter but this time add it to Port 2 and use the SID of 010200 and DID of 0307bc. Figure 27-13.

Figure 27-13. Second Advanced Performance Monitoring Filter applied to Port 2.

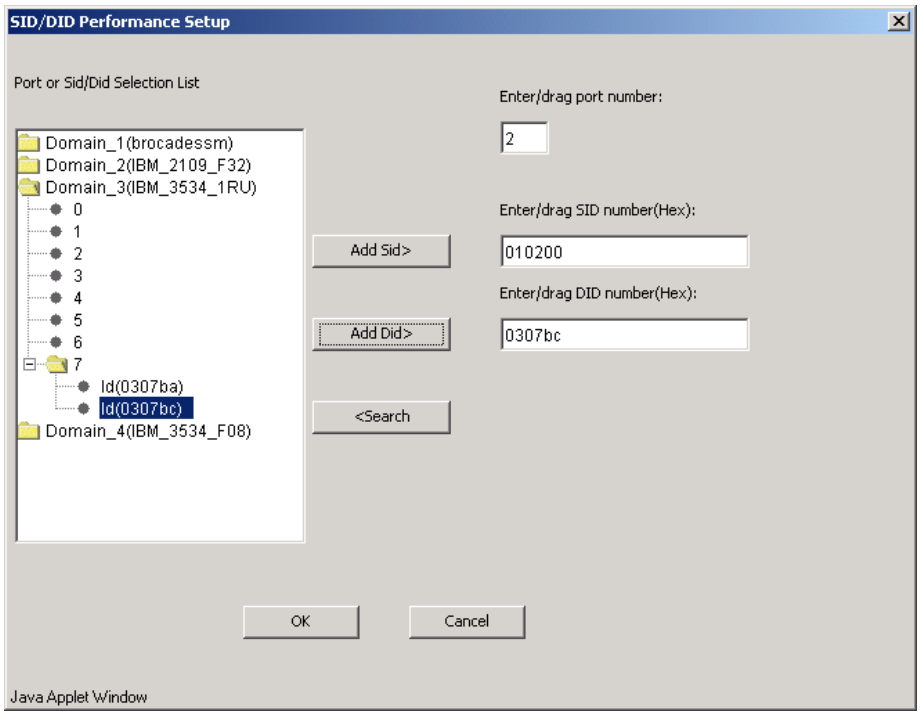

Step 11. Finally start I/O on the Server Blades and view the throughput of the devices. In this example Ports 0 and 15 are ISL Trunked and sending a shared load seen in the top graph. The two Internal Server Blades are in the middle and bottom graphs. Figure 27-14.

Figure 27-14. View Switch Throughput on a per port basis and the two specific SID/DID Advanced Performance Monitoring Filters for each of the 2 Server Blades.

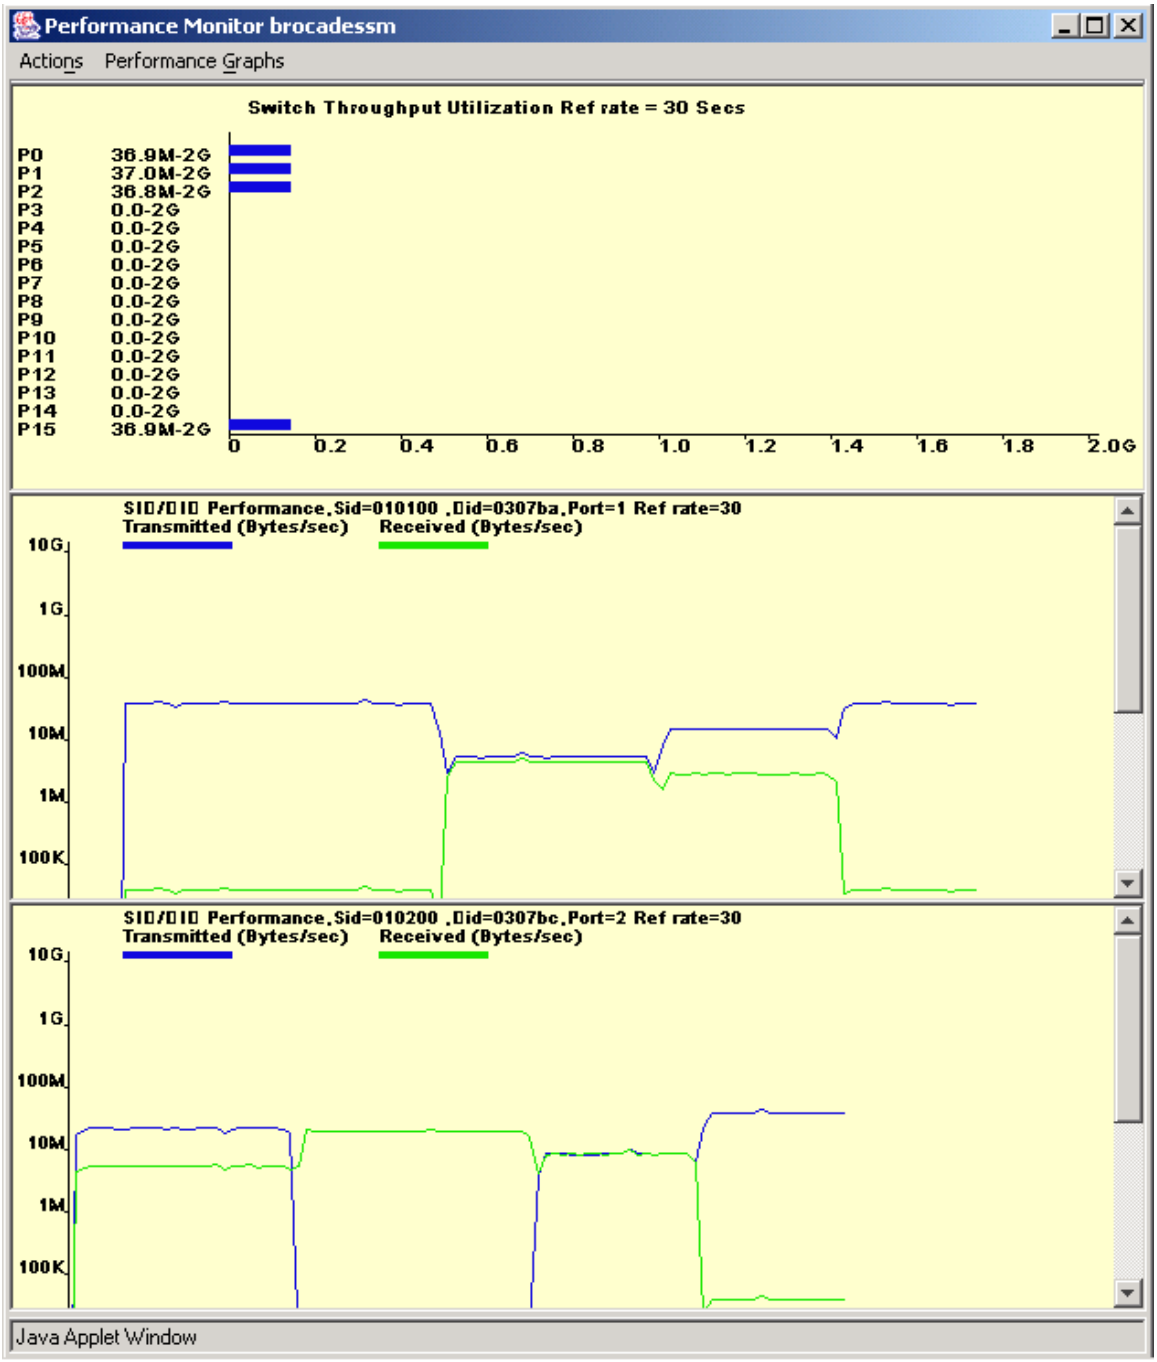

Step 12. To save the Graphs configurations for use at a later time. Click on the Actions Menu and choose the "Save Current Canvas Configuration" from the drop down menu list. Figure 27-15.

Figure 27-15. Saving the Performance Monitoring Canvas.

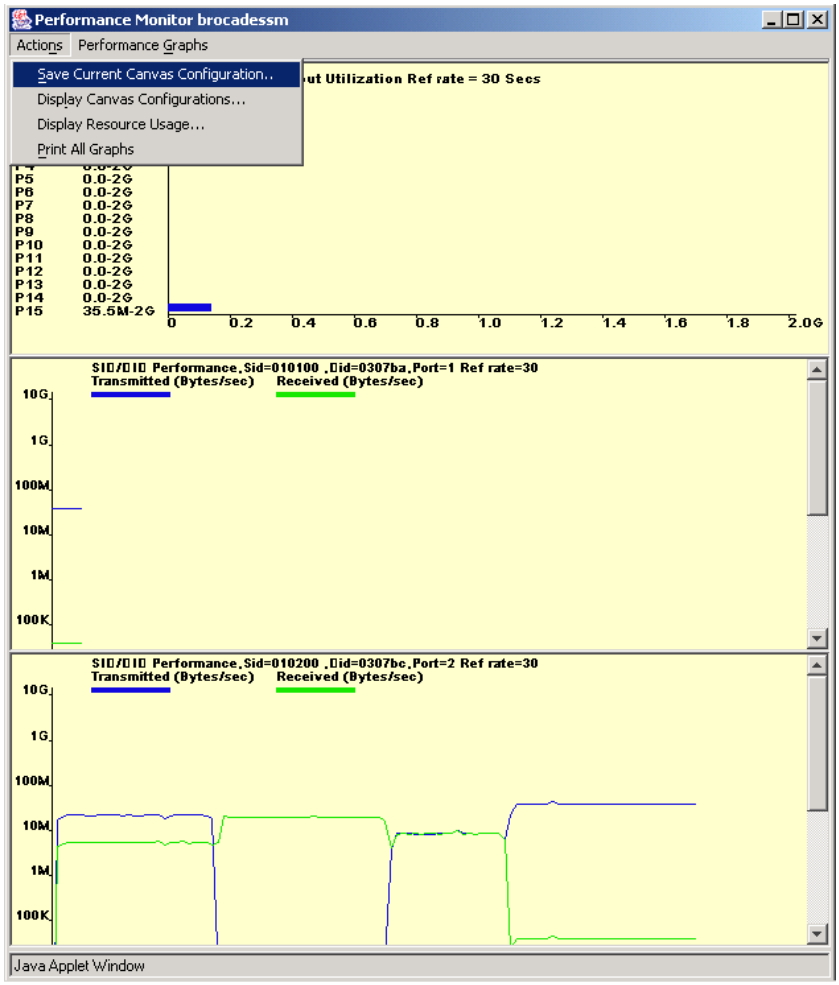

Step 13. The "Save Canvas Configuration" dialog box appears. Choose an Appropriate Name and an appropriate description for this set of graphs. Click on the "Save Canvas" button and then click on "OK" on the confirmation message to save the canvas. Figure 27-16.

Figure 27-16. Save Canvas Configuration.

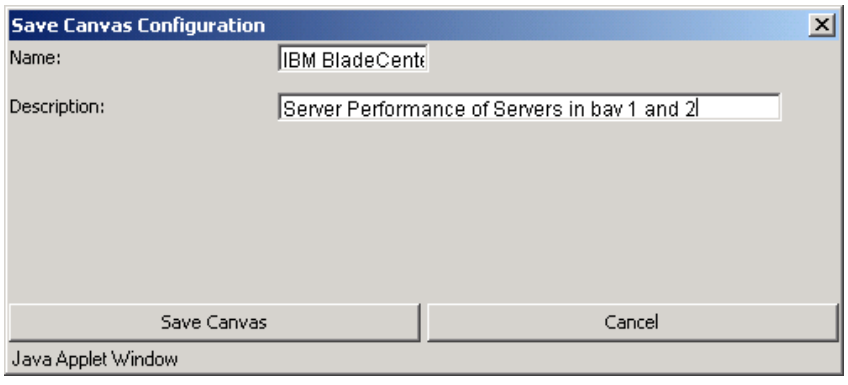

Step 14. To display or load the Saved Canvas Configurations at a later time (i.e. when starting a new session), choose the "Display Canvas Configurations" from the Actions menu. A list of the Saved Canvas Configurations will appear. You can Load, Edit, Copy, or Remove any of the Canvas Configurations from the list. Figure 27-17.

Figure 27-17. Save Canvas Configuration.

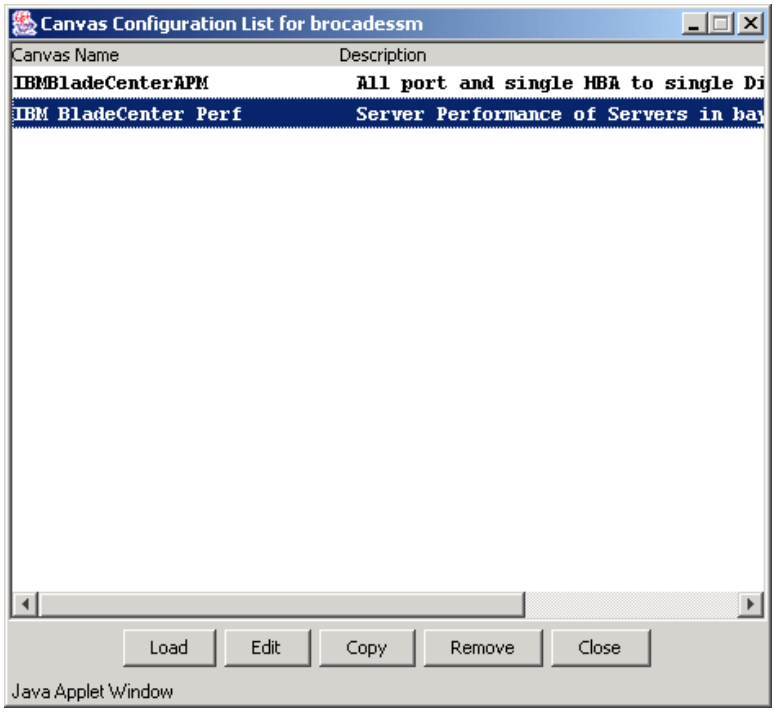

This is a simplified example to show how easy it is to create an Advanced Performance Monitoring Filter. In more complex environments when many Server devices are talking to many storage ports, the ability to analyze the "Hot Spots" of data throughput in the SAN will assist in being able to scale the SAN infrastructure as your bandwidth needs increase.

Appendix A

The Brocade SAN Switch Module can only be used inside the IBM eServer BladeCenter chassis. The Brocade SAN Switch Module can be directly connected to external Fibre Channel storage devices or to an existing or new Fibre Channel SAN fabric. To extend the SAN outside of the IBM eServer BladeCenter chassis refer to the IBM TotalStorage SAN Switch family of products. For reference the entire IBM TotalStorage SAN Switch family names and their associated Brocade names are listed below.

Naming Conventions for the IBM TotalStorage SAN Switch family

IBM Name	**Brocade Name**
1 Gbit Family	
3534-1RU SAN Managed Hub	Brocade SilkWorm 2010, (20X0)
2109-S08 SAN Switch	Brocade SilkWorm 2400
2109-S16 SAN Switch	Brocade SilkWorm 2800
2 Gbit Family	
3534-F08 SAN Switch	Brocade SilkWorm 3200
2109-F16 SAN Switch	Brocade SilkWorm 3800
2109-F32 SAN Switch	Brocade SilkWorm 3900
2109-M12 SAN Switch	Brocade SilkWorm 12000
**new* 2 Gbit Family*	
2005-H08 SAN Switch	Brocade SilkWorm 3250
2005-H16 SAN Switch	Brocade SilkWorm 3850
2109-M14 SAN Switch	Brocade SilkWorm 24000
Brocade SAN Switch Module for IBM eServer BladeCenter	Brocade SilkWorm 3016