Web Tools

Administrator's Guide

Supporting Fabric OS v6.4.0

BROCADE

Brocade Communications Systems, Incorporated

Document History

Document Title	Publication Number	Summary of Changes	Publication Date
Web Tools User's Guide v2.0	*53-0001536-01*	N/A	September 1999
Web Tools User's Guide v2.2	*53-0001558-02*	N/A	May 2000
Web Tools User's Guide v2.3	*53-0000067-02*	N/A	December 2000
Web Tools User's Guide v3.0	*53-0000130-03*	N/A	July 2001
Web Tools User's Guide v2.6	*53-0000197-02*	N/A	December 2001
Advanced Web Tools User's Guide v3.0 / v4.0	*53-0000185-02*	N/A	March 2002
Advanced Web Tools User's Guide v4.0.2	*53-0000185-03*	N/A	September 2002
Advanced Web Tools User's Guide v3.1.0	*53-0000503-02*	N/A	April 2003
Advanced Web Tools User's Guide v4.1.0	*53-0000522-02*	N/A	April 2003
Advanced Web Tools User's Guide v4.1.2	*53-0000522-04*	Insistent Domain ID Mode. Port Swapping information. Minor editorial changes.	October 2003
Advanced Web Tools Administrator's Guide, v4.2.0	*53-0000522-05*	Updates to support new switch types: Brocade 3250, 3850, 24000. Structural changes, support changes, installation changes.	December 2003
Advanced Web Tools User's Guide	*53-0000522-06*	Clarifications on software and hardware support, minor enhancements in procedure text, minor rearranging of content.	March 2004
Advanced Web Tools Administrator's Guide	*53-0000522-07*	Updates to support new switch types (3016, 4100) and Fabric OS v4.4.0, including Ports on Demand, user administration, and zoning wizards.	September 2004
Web Tools Administrator's Guide	*53-0000522-08*	Updates to support new switch types (200E, 48000) and Fabric OS v5.0.1, including switchAdmin role, upfront login, and Web Tools EZ.	April 2005
Web Tools Administrator's Guide	*53-0000522-09*	Updates to add additional information about refresh and polling rates.	July 2005
Web Tools Administrator's Guide	*53-1000049-01*	Updates to support new switch types (4900, 7500) and Fabric OS v5.1.0, including FCR, FCIP, and the FR4-18i port blade. Web Tools EZ information is moved to a separate book.	January 2006
Web Tools Administrator's Guide	*53-1000049-02*	Updates to the FCIP chapter to clarify how to configure tunnels.	April 2006
Web Tools Administrator's Guide	*53-1000194-01*	Updates for Fabric OS v5.2.0 and the FC4-16IP blade. Also new security for Web Tools, including Role-Based Access Control and administrative domains.	September 2006
Web Tools Administrator's Guide	*53-1000435-01*	Updates to reflect interface enhancements, support for new switch types, IPv6 support, and other enhancements.	June 2007

Document Title	Publication Number	Summary of Changes	Publication Date
Web Tools Administrator's Guide	*53-1000606-01*	Updates to reflect updates to enhanced Access Gateway support, changes to FCIP tunneling wizard, and other enhancements.	October 2007
Web Tools Administrator's Guide	*53-1000606-02*	Updates for support for new switches, traffic isolation zoning, F_Port trunking, removal of enhanced Access Gateway support, and other enhancements.	March 2008
Web Tools Administrator's Guide	*53-1001080-01*	Updates to add features that require the Enhanced Group Management license, removal of features that are now available from the Brocade Data Center Fabric Manager (DCFM).	July 2008
Web Tools Administrator's Guide	*53-1001133-01*	Updates to add support for Brocade Encryption Switch and FS8-18 Encryption Blade.	August 2008
Web Tools Administrator's Guide	*53-1001194-01*	Updates to add support for Virtual Fabrics, IPsec, and consistency with DCFM.	November 2008
Web Tools Administrator's Guide	*53-1001343-01*	Updates to add support for Brocade 7800 Extension Switch, Brocade 8000, FCOE10-24 DCX Blade, and FX8-24 DCX Extension Blade.	July 2009
Web Tools Administrator's Guide	*53-1001772-01*	Updates to add support for Brocade Fabric OS 6.4.0.	March 2010

Contents

Chapter 2   **Using the Web Tools Interface**

Chapter 3   **Managing Fabrics and Switches**

Chapter 5   **Managing Administrative Domains**

Chapter 6   **Managing Ports**

Figures

Tables

About This Document

In this chapter

How this document is organized

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

The document contains the following components:

- Chapter 1, "Introducing Web Tools" provides some basic information about the Web Tools interface, including system requirements and installation instructions.
- Chapter 2, "Using the Web Tools Interface" describes the components of the Web Tools interface.
- Chapter 3, "Managing Fabrics and Switches" provides information on how to manage your fabric and switches using the Web Tools interface.
- Chapter 4, "Maintaining Configurations and Firmware" provides information about uploading and downloading configuration files and downloading firmware.
- Chapter 5, "Managing Administrative Domains" provides information on managing Admin Domains.
- Chapter 6, "Managing Ports" provides information about managing FC and GbE ports.
- Chapter 7, "Enabling ISL Trunking" provides information on managing the licensed ISL Trunking feature.
- Chapter 8, "Monitoring Performance" provides information on how to use the Brocade Advanced Performance Monitoring feature to monitor your fabric performance.
- Chapter 9, "Administering Zoning" provides information on how to use the Brocade Advanced Zoning feature to partition your storage area network (SAN) into logical groups of devices that can access each other.

- Chapter 10, "Working with Diagnostic Features" provides information about trace dumps, viewing switch health, and interpreting the LEDs.
- Chapter 11, "Using the FC-FC Routing Service" provides information on using the FC-FC Routing Service to share devices between fabrics without merging those fabrics.
- Chapter 12, "Using the Access Gateway" provides information on how to configure and manage the Brocade Access Gateway.
- Chapter 13, "Administering Fabric Watch" provides information on how to use the Fabric Watch feature to monitor the performance and status of switches and alert you when problems arise.
- Chapter 14, "Administering Extended Fabrics" provides information on how to configure a port for long distance.
- Chapter 15, "Administering the iSCSI Target Gateway" provides information on how to configure and manage the iSCSI Target Gateway.
- Chapter 16, "Routing Traffic" provides information on how to configure routes.
- Chapter 17, "Configuring Standard Security Features" provides information on managing user accounts, SNMP, and the RADIUS server.
- Chapter 18, "Administering FICON CUP Fabrics" provides information on how to administer and manage FICON CUP fabrics. You can enable FMS mode, edit and create configurations, and edit FMS parameters.
- Chapter 19, "Configuring FCoE with Web Tools" provides information on how to configure FCoE features.
- Chapter 20, "Limitations" discusses limitations of and provides workarounds for using Web Tools.

Supported hardware and software

In those instances in which procedures or parts of procedures documented here apply to some switches but not to others, this guide identifies exactly which switches are supported and which are not.

Although many different software and hardware configurations are tested and supported by Brocade Communications Systems, Inc. for Fabric OS v6.4.0, documenting all possible configurations and scenarios is beyond the scope of this document.

The following hardware platforms are supported by this release of Web Tools:

- Brocade 300
- Brocade 4100
- Brocade 4900
- Brocade 5000
- Brocade 5100
- Brocade 5300
- Brocade 5410
- Brocade 5480
- Brocade 5424
- Brocade 5450

- Brocade 7500

- Brocade 7500E

- Brocade 7600

- Brocade 7800 Extension Switch

- Brocade 8000 FCoE Switch

- Brocade VA-40FC

- Brocade Encryption Switch

- Brocade DCX Backbone and Brocade DCX-4S Backbone

  - FA4-18 Fibre Channel application blade

  - FCOE10-24 DCX Blade

  - FS8-18 Encryption Blade

  - FC8-16 port blade

  - FC8-32 port blade

  - FC8-48 port blade

  - FC8-64 port blade

  - FC10-6 port blade

  - FR4-18i router blade

  - FX8-24 DCX Extension blade

- Brocade 48000 director

  - FA4-18 Fibre Channel application blade

  - FC4-16 port blade

  - FC4-16IP blade

  - FC4-32 port blade

  - FC4-48 port blade

  - FC10-6 port blade

  - FR4-18i router blade

What's new in this document

The following major additions have been made since this document was last released:

- Port beaconing

- Enhanced TI zones

- WWN-based Persistent PID

- WWN–N Port mapping

- Switch Status Policy

- Lossless DLS support

- GigE ports—Media type configuration

- FICON—CUP Logical Path and Link Incident Registered Recipient support

- Fabric Watch—Low Above alarms and System Monitor

For further information, refer to the release notes.

Document conventions

This section describes text formatting conventions and important notice formats used in this document.

TEXT FORMATTING

The narrative-text formatting conventions that are used are:

bold text	Identifies command names
	Identifies the names of user-manipulated GUI elements
	Identifies keywords and 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 command syntax examples

For readability, command names in the narrative portions of this guide are presented in mixed lettercase: for example, **switchShow**. In actual examples, command lettercase is often all lowercase. Otherwise, this manual specifically notes those cases in which a command is case sensitive.

NOTES, CAUTIONS, AND WARNINGS

The following notices and statements are used in this manual. They are listed below in order of increasing severity of potential hazards.

NOTE
A note provides a tip, guidance, or advice, emphasizes important information, or provides a reference to related information.

ATTENTION
An Attention statement indicates potential damage to hardware or data.

CAUTION

A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.

DANGER

A Danger statement indicates conditions or situations that can be potentially lethal or extremely hazardous to you. Safety labels are also attached directly to products to warn of these conditions or situations.

KEY TERMS

For definitions specific to Brocade and Fibre Channel, see the *Brocade Glossary*.

For definitions of SAN-specific terms, visit the Storage Networking Industry Association online dictionary at:

http://www.snia.org/education/dictionary

Notice to the reader

This document may contain references to the trademarks of the following corporations. These trademarks are the properties of their respective companies and corporations.

These references are made for informational purposes only.

Corporation	Referenced Trademarks and Products
Microsoft Corporation	Windows, Windows NT, Internet Explorer
Oracle Corporation	Oracle, Solaris
Netscape Communications Corporation	Netscape
Red Hat, Inc.	Red Hat, Red Hat Network, Maximum RPM, Linux Undercover
Mozilla	Firefox

Additional information

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

BROCADE RESOURCES

To get up-to-the-minute information, go to *http://my.brocade.com* and register at no cost for a user ID and password.

White papers, online demonstrations, and data sheets are available through the Brocade website at:

   http://www.brocade.com/products-solutions/products/index.page

For additional Brocade documentation, visit the Brocade website:

   http://www.brocade.com

Release notes are available on the MyBrocade website and are also bundled with the Fabric OS firmware.

OTHER INDUSTRY RESOURCES

For additional resource information, visit the Technical Committee T11 website. This website provides interface standards for high-performance and mass storage applications for Fibre Channel, storage management, and other applications:

   http://www.t11.org

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

   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

   - Switch model

   - Switch operating system version

   - Software name and software version, if applicable

   - Error numbers and messages received

   - **supportSave** command output

   - Detailed description of the problem, including the switch or fabric behavior immediately following the problem, and specific questions

   - Description of any troubleshooting steps already performed and the results

   - Serial console and Telnet session logs

   - syslog message logs

2. Switch Serial Number

   The switch serial number and corresponding bar code are provided on the serial number label, as illustrated below:

   |‖‖‖‖‖‖‖‖‖‖‖‖‖‖‖‖‖‖‖‖‖|

   FT00X0054E9

   The serial number label is located as follows:

- *Brocade 300, 4100, 4900, 5100, 5300, 7500, 7800, 8000*, *VA-40FC*, and *Brocade Encryption Switch*—On the switch ID pull-out tab located inside the chassis on the port side on the left

- *Brocade 5000*—On the switch ID pull-out tab located on the bottom of the port side of the switch

- *Brocade 7600*—On the bottom of the chassis

- *Brocade 48000*—Inside the chassis next to the power supply bays

- *Brocade DCX*—On the bottom right on the port side of the chassis

- *Brocade DCX-4S*—On the bottom right on the port side of the chassis, directly above the cable management comb

3. World Wide Name (WWN)

   Use the **licenseIdShow** command to display the WWN of the chassis.

   If you cannot use the **licenseIdShow** command because the switch is inoperable, you can get the WWN from the same place as the serial number, except for the Brocade DCX. For the Brocade DCX, access the numbers on the WWN cards by removing the Brocade logo plate at the top of the nonport side of the chassis.

Document feedback

Quality is our first concern at Brocade and 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 of the document and as much detail as possible about your comment, including the topic heading and page number and your suggestions for improvement.

Introducing Web Tools

In this chapter

Web Tools overview

Brocade Web Tools is an embedded graphical user interface (GUI) that enables administrators to monitor and manage single or small fabrics, switches, and ports. Web Tools is launched directly from a web browser, or from the DCFM client.

A limited set of features is accessible using Web Tools without a license, and is available free of charge. Additional switch management features are accessible using Web Tools with the Enhanced Group Management (EGM) license. Refer to "Web Tools, the EGM license, and DCFM" for more information.

Web Tools, the EGM license, and DCFM

Beginning with Fabric OS version 6.1.1, Web Tools functionality is tiered and integrated with DCFM. If you are migrating from a Web Tools release prior to Fabric OS version 6.1.1, this may impact how you use Web Tools.

A Web Tools license is not required, and a basic version of Web Tools is available for free. Additional functionality may be added by obtaining the Enhanced Group Management (EGM) license. Table 1 compares Basic Web Tools features to Web Tools with the EGM license. The EGM license is only for 8 Gbps platforms, such as the Brocade DCX and DCX-4S enterprise-class platforms, the Encryption Switch, and the 300, 5100, and 5300 switches. For non-8 Gbps platforms, all functionalities are available without the EGM license.

Beginning with Fabric OS version 6.1.1, some Web Tools capabilities are moved from Web Tools to DCFM. Table 2 summarizes these changes.

Web Tools features enabled by the EGM license

Table 1 describes those Web Tools features that require the EGM license.

TABLE 1 Basic Web Tools features and EGM licensed features

Feature	Basic Web Tools	Web Tools with EGM License
Active Directory support	yes	yes
AD Context Switching	no	yes
AD filtered views	yes	yes
Admin Domain Management	no	yes
AG Management	yes	yes
Analyze zone config	no	no
Basic Zoning and TI Zoning	yes	yes
Blade Management	yes	yes
Cloning a zone	no	yes
Configuration upload/download	yes	yes
Convenience function from Tools menu	no	no
Device Accessibility Matrix	no	no
Easy to configure iSCSI wizard	yes	yes
Extended Fabric Management	no	yes
F_Port Trunk Management	no	yes
Fabric Events	no	no
Fabric Summary	no	no
Fabric Tree	yes	yes
FCIP Tunnel configuration	no	no
FCIP Tunnel Display	yes	yes
FCR Management	yes	yes
FCR Port Config	yes	yes
FICON CUP Tab	no	yes
FRU Monitoring	yes	yes
High Availability	yes	yes
IP Sec Policies	yes	yes
ISL Trunk Management	no	yes
ISL Trunking information	yes	yes
License Management	yes	yes
Long Distance	no	yes
Logical Switch Context Switching	no	yes

TABLE 1   Basic Web Tools features and EGM licensed features (Continued)

Feature	Basic Web Tools	Web Tools with EGM License
PDCM Matrix	no	yes
Performance Monitoring Dialog	no	yes
Port Administration	yes	yes
Print zone database summary	no	no
RBAC	yes	yes
Routing and DLS Configuration	no	yes
Security Policies Tab (like ACL)	yes	yes
Switch Info tab	yes	yes
Switch Status	yes	yes
Switch View right-click options	yes	yes
Trace dump	yes	yes
USB Management	yes	yes
User Management	yes	yes
Verify and troubleshoot accessibility between devices	yes	yes

Web Tools functionality moved to DCFM

The functionality that was moved from Web Tools into DCFM is applicable to both DCFM Professional and DCFM Enterprise. Table 2 details these changes.

TABLE 2   Web Tools functionality moved to DCFM

Function	Web Tools 6.1.0	DCFM	Comments
Add Un-Zoned Devices	Zone Admin	**Configure** > **Zoning** Reverse Find in the Zoning dialog box provides the view of the zoned and unzoned devices in the fabric if all zone members are selected for Find.	
Analyze Zone Config	Zone Admin	1 **Configure** > **Zoning**: Reverse Find in the Zoning dialog box provides the view of the zoned and unzoned devices in the fabric if all zone members are selected for Find. 2 Device Tree and Topology: Connected End Devices — Custom Display from the top level in the main frame provides the device tree and topology view for all the zoned devices if all zones are selected in the active zone configuration.	
Define Device Alias	Zone Admin	**Configure** > **Zoning**	

TABLE 2   Web Tools functionality moved to DCFM (Continued)

Function	Web Tools 6.1.0	DCFM	Comments
Device Accessibility Matrix	Zone Admin	**Configure** > **Zoning** the Compare dialog box provides the Storage-Host and Host-Storage view in a tree representation that is comparable to the Device Accessibility Matrix when all devices are selected.	
Fabric Events	**Monitor** > **Fabric Events**	**Monitor** > **Logs** > **Events**	
Fabric Summary	**Reports** > **Fabric Summary**	**Monitor** > **Reports** > **Fabric Summary Report**	
FCIP Tunnel Configuration	**Port Admin Module** > **GigE** tab	**Configure** > **FCIP Tunnel**	Viewing FCIP tunnels is still supported in Web Tools 6.1.1, but New, Edit Config, and delete are only available in DCFM.
GigE Ports Interface	**Port Admin Module** > **GigE** tab	**Configure** > **FCIP Tunnel**	
GigE Ports Route	**Port Admin Module** > **GigE** tab	**Configure** > **FCIP Tunnel**	
Non-local switch ports display in zoning tree	**Zone Admin** **Admin Domain** **Switch Admin** > **DCC policies** **Performance Monitoring**	**Configure** > **Zoning**	In Web Tools, non-local switch port id/WWN can be added using text box.
Remove Offline or Inaccessible Devices	**Zone Admin**	**Configure** > **Zoning** Replace/Replace All zone members by selecting the offline devices from the zone tree. Offline devices have an unknown overlay badge with good visibility.	
Zone database summary print	**Zone Admin**	**Configure** > **Zoning** Zoning report for both online and offline database.	

System requirements

Before you install Web Tools on your workstation, verify that your switches and workstation meet the Web Tools requirements listed in this chapter.

Web Tools requires any browser that conforms to HTML version 4.0, JavaScript version 1.0, and Java Plug-in 1.6.0_16 or later.

Brocade has certified and tested Web Tools on the platforms shown in Table 3.

TABLE 3   Certified and tested platforms

Operating System	Browser
Windows 7	Internet Explorer 7.0/8.0
Windows Server 2008 Standard	Internet Explorer 7.0/8.0
Windows Vista Business	Internet Explorer 7.0/8.0
Windows XP, SP3	Firefox 2.0/3.0 and Internet Explorer 7.0/8.0
RedHat Enterprise Server 5 Advanced Platform	Firefox 2.0
SUSE Linux Enterprise Server 10	Firefox 2.0/3.0
Solaris 10 (SPARC only)	Firefox 2.0
Linux RedHat AS 3.0	Firefox 2.0
Linux RedHat AS 4.0	Firefox 2.0/3.0
Windows 2000	Firefox 2.0, Internet Explorer 6.0
Windows 2003 Server, SP2	Firefox 2.0/3.0 and Internet Explorer 7.0/8.0
Windows XP, SP2	Firefox 2.0/3.0 and Internet Explorer 7.0/8.0
Solaris 9 (SPARC only)	Firefox 2.0

For Windows systems, a minimum of 256 MB of RAM for fabrics comprising up to 15 switches, 512 MB of RAM for fabrics comprising more than 15 switches, and a minimum of 8 MB of video RAM are recommended. Additionally, a DCX with a fully populated FC8-64 blade requires a minimum of 512 MB of RAM.

Setting refresh frequency for Internet Explorer

Correct operation of Web Tools with Internet Explorer requires specifying the appropriate settings for browser refresh frequency and process model. Browser pages should be refreshed frequently to ensure the correct operation of Web Tools.

User the following procedure to set the Internet Explorer options.

1. Open your web browser and select **Tools** > **Internet Options**.

2. Select **General** > **Temporary Internet Files > Settings**.

3. Choose **Every visit to the page** under "Check for newer versions of stored pages:" as shown in
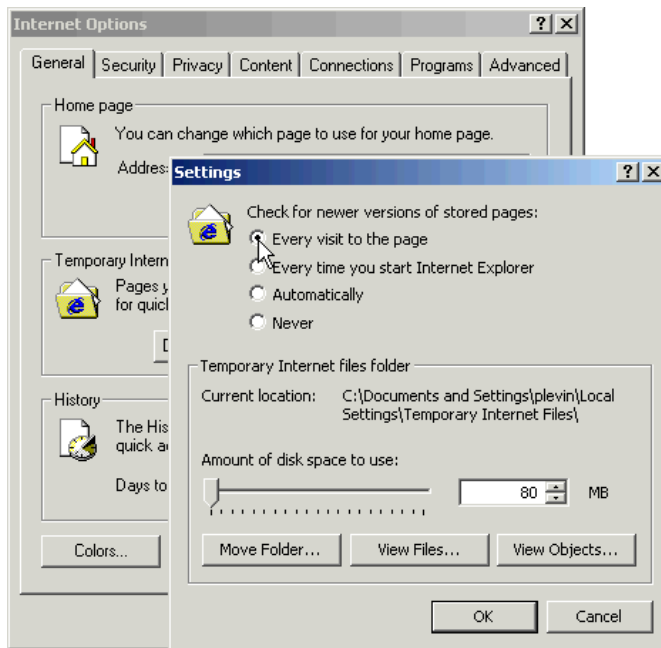
FIGURE 1   Configuring Internet Explorer

Deleting temporary internet files used by Java applications

For Web Tools to operate correctly, you must delete the temporary internet files used by Java applications.

Use the following procedure to delete these files.

1. From the Control Panel, open Java.

2. Select the **General** tab and click **Settings**.

3. Click **Delete Files** to remove the temporary files used by Java applications.

4. Click **OK** on the confirmation dialog box.

   You can clear the **Trace and Log files** check box if you want to keep those files.

5. Click **OK**.

6. On the Java Control Panel, click **View** to review the files that are in the Java cache.

   If you have deleted all the temporary files, the list is empty.

Java installation on the workstation

Java Plug-in must be installed on the workstation. If you attempt to open Web Tools without any Java Plug-in installed:

- Internet Explorer automatically prompts and downloads the proper Java Plug-in.
- Firefox downloads the most recently released Java Plug-in.

If you attempt to open Web Tools with a later version of Java Plug-in installed:

- Internet Explorer might prompt for an upgrade, depending on the existing Java Plug-in version.
- Firefox uses the existing Java Plug-in.

Installing the JRE on your Solaris or Linux client workstation

Use the following procedure to do the JRE installation.

1. Locate the JRE on the Internet, at the following URL:

    http://java.sun.com/products/archive/j2se/5.0_06/index.html

    NOTE
    This URL points to a non-Brocade website and is subject to change without notice.

2. Download **JRE 5.0 Update 13**.

3. Follow the instructions to install the JRE.

4. Create a symbolic link from this location:

    $FIREFOX/plugins/libjavaplugin_oji.so

    To this location:

    $JRE/plugin/$ARCH/ns600/libjavaplugin_oji.so

Installing patches on Solaris

Use the following procedure to install patches on Solaris.

1. Search for any required patches for your current version of the JRE at the following website:

    http://sunsolve.sun.com/pub-cgi/show.pl?target=patchpage

    NOTE
    This URL points to a non-Brocade website and is subject to change without notice.

2. Follow the link to download the patch

3. Exit the browser when you have downloaded the patch.

4. Install the patch and restart the system.

Installing the Java plug-in on Windows

Use the following procedure to Install the Java plug-in on Windows.

1. Select **Start Menu** > **Settings** > **Control Panel** and select the Java Plug-in Control Panel.

2. Select the **About** tab.

3. Determine whether the correct Java Plug-in version is installed:

    - If the correct version is installed, Web Tools is ready to use.

- If no Java Plug-in is installed, point the browser to a switch running Fabric OS 5.2.0 or later to install JRE 1.6.0. For Fabric OS 6.3.0 install JRE 1.6.0 update 13. Web Tools guides you through the steps to download the proper Java Plug-in.

- If an outdated version is currently installed, uninstall it, restart your computer, reopen the browser, and enter the address of a switch running Fabric OS 5.2.0 or later to install JRE 1.6.0. For Fabric OS 6.3.0 install JRE 1.6.0 update 13. Web Tools guides you through the steps to download the proper Java Plug-in.

Java plug-in configuration

If you are managing fabrics with more than 10 switches or 1000 ports, or if you are using the iSCSI Gateway module extensively, you should increase the default heap size to 256 MB to avoid out-of-memory errors.

If you are using a Mozilla family browser (Firefox, Netscape), you should set the default browser in the Java control panel.

The following procedures instruct you in increasing the default heap size in the **Java Control Panel** and in setting the default browser.

Configuring the Java plug-in for Windows

Use the following procedure to configure Java plug-in for Windows.

1. From the **Start** menu, select **Settings** > **Control Panel** > **Java**.

2. Click the **Java** tab.

3. In the section **Java Applet Runtime Settings**, click **View.**

   The **Java Runtime Settings** dialog box displays.

4. Double-click the **Java Runtime Parameters** field and enter the following information to set the minimum and maximum heap size:

   ```
   -Xms256m  -Xmx256m
   ```

   In this example, the minimum and maximum sizes are both 256 MB.

5. Click **Apply** to apply your settings and close the Java Control Panel.

Configuring the Java plug-in for Mozilla family browsers

Use the following procedure to configure Java plug-in for Mozilla family browsers.

1. From the **Start** menu, select **Settings** > **Control Panel.**

2. Click the **Advanced** tab and expand the **Default Java for browsers** option, as shown in Figure 2 on page 9.
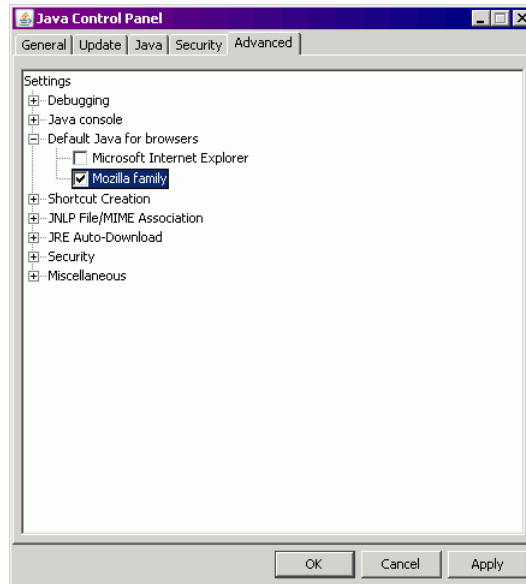
FIGURE 2   Default Java for browsers option

3. Select **Mozilla family** and click **OK**.

4. Click **OK** to apply your settings and close the Java Control Panel.

Value line licenses

If you open Web Tools on a switch with a limited license, and if the fabric exceeds the switch limit indicated in the license, then Web Tools displays a warning message. Web Tools allows a 30-day grace period, during which you can still monitor the switch while continuing to display warning messages periodically.

These messages warn you that your fabric size exceeds the supported switch configuration limit and tells you how long you have before Web Tools is disabled. After the 30-day grace period, you are no longer able to open Web Tools from the switch with the limited switch license.

Web Tools is part of the Fabric OS of a switch. When you open Web Tools on a switch, you can manage other switches in the fabric that have lower or higher firmware versions. It is important to note that when accessing these switches you are opening the remote switch's version of Web Tools, and the functionality available for those switches might vary.

Opening Web Tools

You can open Web Tools on any workstation with a compatible Web browser installed. For a list of Web browsers compatible with Fabric OS v6.4.0, refer to Table 3. Web Tools supports both HTTP and HTTPS protocol.

Use the following procedure to open Web Tools.

1. Open the Web browser and enter the IP address of the device in the **Address** field:

   `http://10.77.77.77`

   or

   `https://10.77.77.77`

2. Press **Enter**. The Web Tools login dialog box displays. Refer to "Logging in" on page 11 for more information.

   > **NOTE**
   > If you are using Firefox, the browser window is left open. You can close it anytime after the login dialog box displays. If you are using Internet Explorer, the browser window automatically closes when the login dialog box displays.

   > **NOTE**
   > If you have installed EZSwitchSetup on your workstation, the EZSwitchSetup Switch Manager displays the first time you access the device. EZSwitchSetup provides an easy to use wizard interface that may be used to simplify the initial setup procedure for smaller switches. Refer to *EZSwitchSetup Administrator's Guide* for information about the EZSwitchSetup interface.
   >
   > If you want to use Web Tools instead of EZSwitchSetup, click **Advanced Management** in the lower-left corner of the window to open the Web Tools interface. This book describes only the Web Tools interface.
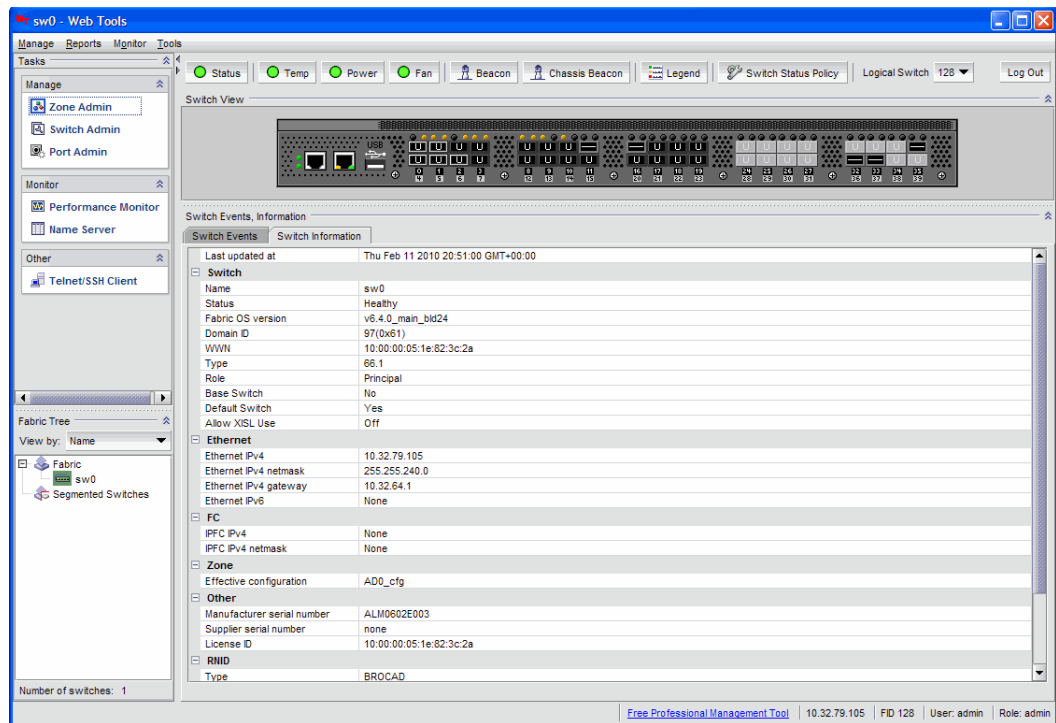
FIGURE 3   Web Tools interface

Logging in

When you use Web Tools, you must log in before you can view or modify any switch information. This section describes the login process.

Prior to displaying the login window, Web Tools displays a security banner (if one is configured for your switch), that you must accept before logging in. The security banner displays every time you access the switch.

When you are presented with the login screen you must provide a user name and a password. Your home Admin Domain is automatically selected. You can select to log in to an Admin Domain other than your home domain.

NOTE
You must login before you can view Switch Explorer (shown in Figure 3 on page 10).

Use this procedure to log in to the Admin Domain.

1. Click **Run** on the signed certificate applet.

   A warning dialog box may display. If you select the check box **Always trust content from this publisher**, the warning dialog box is not displayed when you open Web Tools again.

2. Click **OK** in the security banner window, if one displays.

3. In the login dialog box, enter your user name and password.

   If your current password has expired, you must also provide a new password and confirm the new password.

Logging in to a Virtual Fabric

If you are logging in to a platform that is capable of supporting Virtual Fabrics, the login dialog box provides the option of logging in to a virtual fabric. The following platforms support virtual fabrics:

- Brocade DCX and DCX-4S
- Brocade 5300
- Brocade 5100

Use the following procedure to log in to a Virtual Fabric.

1. Select **Options** to display the Virtual Fabric options.

   You are given a choice between **Home Logical Fabric** and **User Specified Virtual Fabric** (Figure 4). **Home Logical Fabric** is the default.
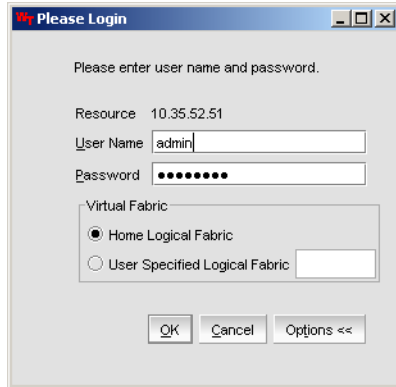
FIGURE 4   Virtual Fabric login option

2. Log in to a logical fabric.

   - To log in to the home logical fabric, select **Home Logical Fabric** and click **OK**.

   - To log in to a logical fabric other than the home logical fabric, select **User Specified Logical Fabric**, enter the fabric ID number, and click **OK**.

Logging in to an Admin Domain

If you are logging in to a platform that is capable of supporting Admin Domains, the login dialog box displays. You do not have an Admin Domain option if the Access Gateway or Interoperability mode is enabled. Admin Domains and Virtual Fabrics are mutually exclusive.

1. Select **Options** to select an Admin Domain other than your default home domain.

   You are given a choice of **Home Domain** (the default), or **User Specified Domain**.

2. Log in to an Admin Domain.

- To log in to the home domain, select **Home Domain** and click **OK**.
- To log in to an Admin Domain other than the home domain, select **User Specified Domain**, enter the Admin Domain name or number, and click **OK** (Figure 5).

FIGURE 5   Login dialog box with Admin Domain options

If the user name or password is incorrect, a dialog box displays indicating an authentication failure.

If you entered valid credentials, but specified an invalid Admin Domain, a dialog box displays from which you can select a valid Admin Domain or click **Cancel** to log in to your home domain (Figure 6).
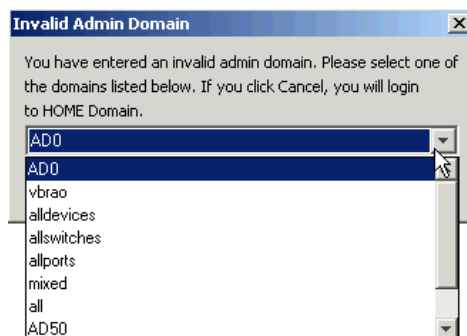
FIGURE 6   Invalid Admin Domain dialog box

Logging out

You can end a Web Tools session either by logging out or by closing the **Switch Explorer** window.

You might be logged out of a session involuntarily, without explicitly clicking the **Logout** button, under the following conditions:

- A physical fabric administrator changes the contents of your currently selected Admin Domain.
- Your currently selected Admin Domain is removed or invalidated.
- Your currently selected Admin Domain is removed from your Admin Domain list.
- You initiate a firmware download from Web Tool's **Switch Administration** window. In this case, you are logged out a few minutes later when the switch restarts.

- Your session times out.

Role-Based Access Control

Role-Based Access Control (RBAC) defines the capabilities that a user account has based on the assigned role. For each role, there is a set of predefined permissions on the jobs and tasks that can be performed on a fabric and its associated fabric elements.

When you log in to a switch, your user account is associated with a predefined role. The role determines the level of access you have on that switch and in the fabric. Table 4 describes these roles.

TABLE 4   Predefined Web Tools roles

Role	Description
admin	You have full access to all of the Web Tools features.
operator	You can perform any actions on the switch that do not affect the stored configuration.
securityadmin	You can perform actions that do not affect the stored configuration.
switchadmin	You can perform all actions on the switch, except the following: • You cannot modify zoning configurations. • You cannot create new accounts. • You cannot view or change account information for any accounts. You can only view your own account and change your account password.
zoneadmin	You can only create and modify zones.
fabricadmin	You can do everything the Admin role can do except create new users.
basicswitchadmin	You have a subset of Admin level access.
user	You have nonadministrative access and can perform tasks such as monitoring system activity.

For information about changing user account roles, refer to "User-defined accounts" on page 207.

Session management

A Web Tools session is the connection between the Web Tools client and its managed switch. A session is established when you log in to a switch through Web Tools. When you close Switch Explorer, Web Tools ends the session.

A session remains in effect until one of the following happens:

- You log out
- You close the **Switch Explorer** window
- The session ends due to inactivity (time out)

A session automatically ends if no information was sent to the switch for more than two hours. Because user key strokes are not sent to the switch until you apply or save the information, it is possible for your session to end while you are entering information in the interface. For example, entering a zoning scheme in the Zoning module does not require you to send information to the switch until you save the scheme.

Web Tools does not display a warning when the session is about to time out. If your session ends due to inactivity, all Web Tools windows become invalid and you must restart Web Tools and log in again.

Web Tools enables sessions to both secure and nonsecure switches.

Access rights for your session are determined by your role-based access rights and by the contents of your selected Admin Domain. After you log in, you can change to a different Admin Domain at any time. However, you cannot change your role-based permissions.

Ending a Web Tools session

To end a Web Tools session, perform one of the following actions:

- Click **Logout** in the **Switch Explorer** window.
- Click the X in the upper-right corner of the **Switch Explorer** window to close it.
- Close all open Web Tools windows.

NOTE
If you click **Logout** in the **Switch Explorer** window, and Web Tools leaves the **Temperature**, **Fan**, **Power,** and the **Switch status** windows open, you must manually close them.

Web Tools system logs

Web Tools uses the log4j framework to write the logs into a file

When you launch Web Tools for the first time, it automatically creates the following directories. These directories are created under Web Tools directory if they are not available:

- A <Web Tools> directory under the user home directory.
- The Web Tools Switch Support Save directory with the name format <Core Switch Name-Switch IP Address-Switch WWN>.

The Web Tools Switch Support Save directory contains the following files:

- Log4j.xml
- WebTools.log
- SwitchInfo.txt

The SwitchInfo.txt file contains the following basic switch information:

- Switch Name
- Fabric OS version
- Switch Type
- Ethernet Ipv4
- Ethernet IPv4 subnet mask
- Ethernet IPv4 gateway

The maximum size of the webtools.log file is 2MB. It is rolled into new file when the 5mb file size limit is exceeded. A backup file named webtools1.log is automatically created. Web Tools maintains only one webtools.log backup file at a time.

The Web Tools debug dialog box can be used to enable the debug state and level for a module at runtime.

If you are familiar with XML scripting, you can edit the configuration file (log4j.xml) to collect the data at startup. If you edit the configuration file, Web Tools need to be restarted. Contact your switch support supplier for assistance.

Requirements for IPv6 support

The following list provides requirements for Web Tools IPv6 support:

- In a pure IPv6 environment, you must configure your DNS maps to the IPv6 address of the switch.
- The switch name is required to match the DNS name that is mapped to the IPv6 address.
- If both IPv4 and IPv6 addresses are configured, Web Tools can be launched using any configured IP address.
- Use a switch with v5.3.0 or later firmware to manage a mixed fabric of IPv4 and IPv6 switches.
- Switches running on version 5.2.0 do not discover IPv6 address-only switches in the same fabric until the IPv4 address is configured.

Using the Web Tools Interface

In this chapter

Viewing Switch Explorer

The first thing you see when you log in to a switch with Web Tools is Switch Explorer, shown in Figure 7 on page 19. Switch Explorer is divided into areas that provide access to, and information about, the switch and fabric. The Switch Explorer areas are:

- The left pane, displaying the **Tasks** and **Fabric Tree** areas.

  The **Tasks** area lets you perform management, monitoring, and other tasks. The Fabric Tree displays a list of all the switches in the fabric.

- The menu bar, at the top of the window, providing access to commands and actions. The menu bar displays the same commands as the left pane of Switch Explorer.

  If you choose to collapse the left pane, you still have access to:

  - Management tasks, such as zone administration, switch administration, port administration, admin domain administration, and Fabric Watch administration.

    > **NOTE**
    > You can manage basic zoning and Traffic Isolation zoning using Web Tools and Web Tools with the Enhanced Group Management (EGM) license. To perform clone operations for zoning, the EGM license must be installed on the switch; otherwise, access to this feature is denied and an error message displays.
    > You must use DCFM Professional or Enterprise Edition to print the zone database summary configuration and to analyze zone configurations. For more information on zoning management, refer to "Zone configuration and zoning database management" on page 131.

  - Reporting tasks, such as viewing the status of a switch.

  - Monitoring tasks, such as performance monitoring, and viewing the temperature or power status.

> **NOTE**
> To perform monitoring tasks such as performance monitoring the EGM license must be installed on the switch; otherwise, access to this feature is denied and an error message displays.

- Tools tasks, such as opening the Telnet window.

• The Switch View buttons above the Switch View provides access to switch information: status, temperature, power, and fan data, beaconing, and the legend for the Switch View.

Although clicking a button can open a separate dialog box or window displays the management tasks, all access control is established when you first log in to the switch.

Buttons in Switch Explorer are unavailable for two reasons: your account does not have sufficient privileges to access this feature, or your currently selected Admin Domain does not meet some condition to access the feature.

• The Admin Domain context field indicates the administrative domain you are viewing and allows you to change it.

• The Switch View displays an interactive graphic of the switch.

• The Switch Events and Switch Information tabs allow you to view event information and switch information, including connectivity, port, zone and other information.

• An indicator bar in the lower-right corner of every module window contains the Admin Domain you are currently viewing, the current user name logged in to the switch, and the role associated with that user account.

• The small right arrow near the **Switch Event** tab displays the switch. When you log out of Web Tools, it remembers the last window settings the next time you log in to the application. If you display the switch, the next time you log in to Web Tools, by default Switch View displays.

The EGM license is required only for 8 Gbps platforms, such as the following:

• Brocade DCX and DCX-4S enterprise-class platforms

• Encryption Switch

• 300, 5300, and 5100 switches

For non-8 Gbps platforms, all functionalities are available without EGM license.
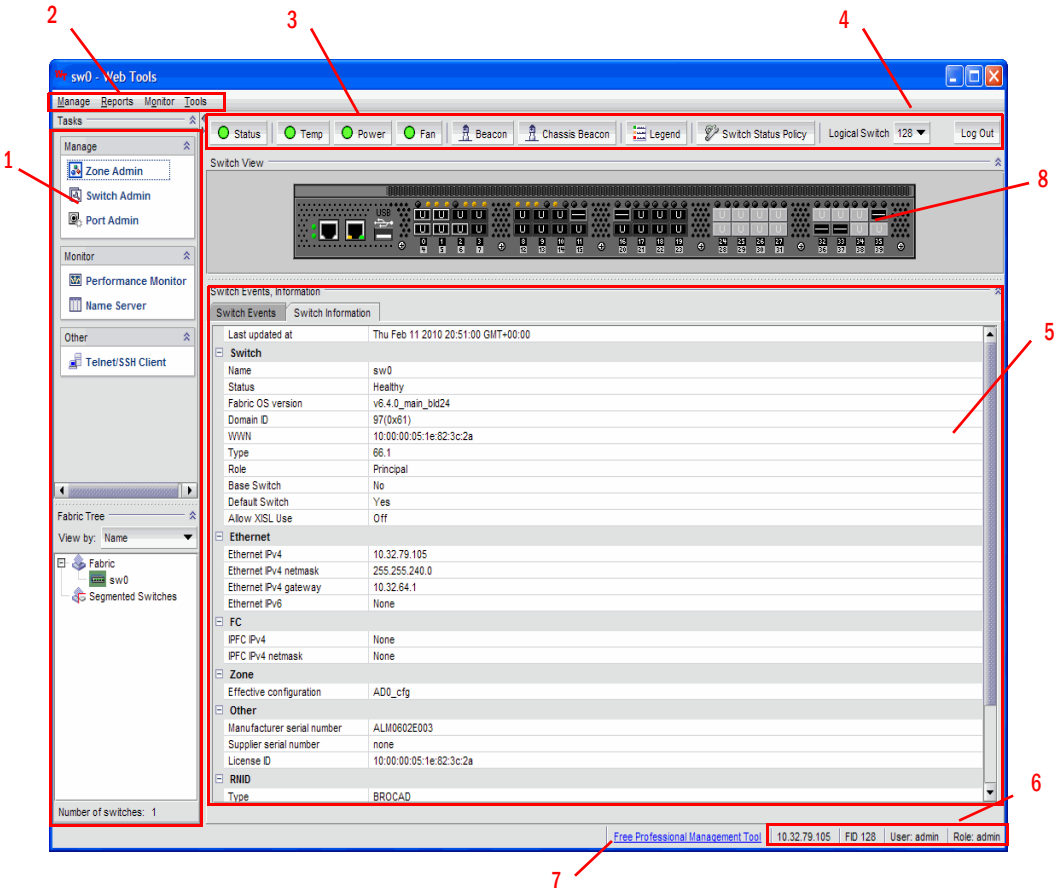
1	Tasks and Fabric Tree
2	Menu bar
3	Switch View buttons
4	Changing the Virtual Fabric ID, or Changing the Admin Domain
5	Switch Events and Switch Information
6	Indicator bar
7	Professional Management Tool offering
8	Switch View

FIGURE 7   Switch Explorer

Changes for consistency with DCFM

Beginning with Fabric OS version 6.2.0, Web Tools icons are changed to be consistent with DCFM. Table 5 summarizes these changes.

TABLE 5 Icon image changes

Image name	Old image	New image
Switch		
Director or DCX		
Fabric		
Fabric segment		
Admin port		
Generic port		
Zone		
Zone Admin		
Zone config		
Zone folder		
Status button - Healthy		
Status button - Marginal		
Status button - Critical		
Status button - Unmonitored		
Switch event - Fatal		
Switch event - Informational		

TABLE 5　　Icon image changes　(Continued)

Image name	Old image	New image
Switch event - Warning		
Refresh		
Enable		
Disable		
Prohibit		

The **Search**, **Copy**, and **Export** buttons are removed from the Web Tools tree and table headers, and are replaced by right-click operations, as shown in Figure 8.
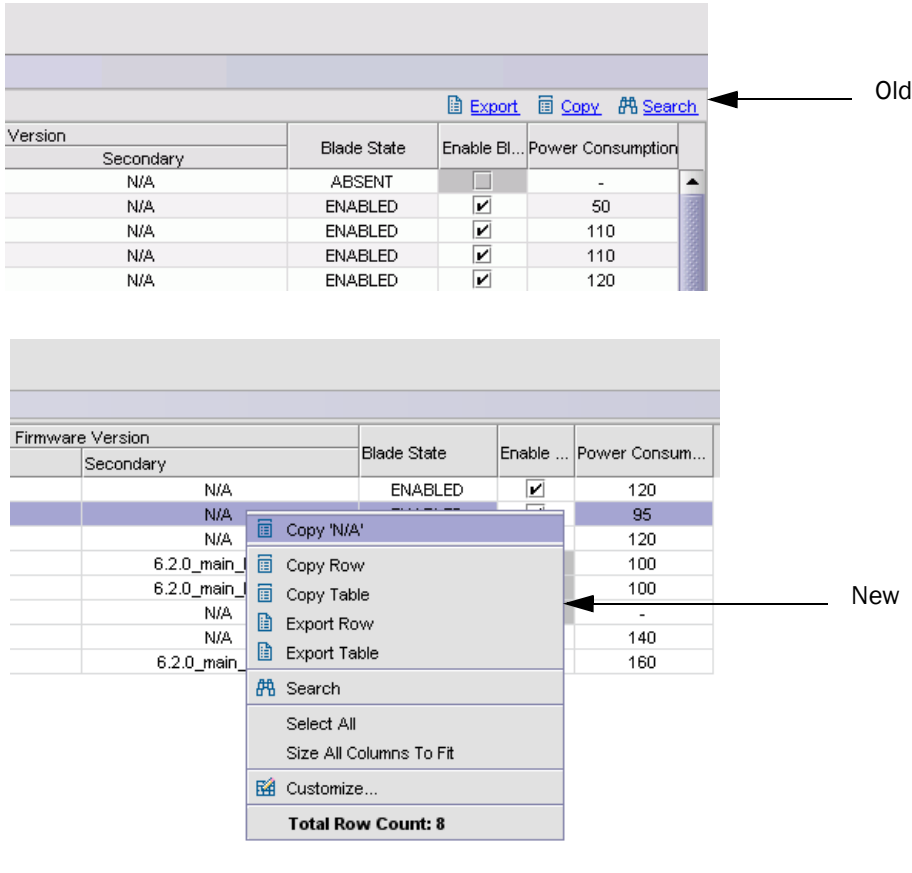
FIGURE 8　　Right-click for Copy, Export, and Search

Persisting GUI preferences

Web Tools persists your GUI preferences across sessions for the Port Admin, Switch Admin, and Name Server dialog boxes on all web-browser platforms. Persistence is performed on a per host basis.

The Port Admin GUI preferences that persist are:

- Basic or Advance mode
- Last selected tab by the user
- Table column sorting
- Table column positions

The Switch Admin GUI preferences that persist are:

- Basic or Advance modes

The Name Server GUI preferences that persist are:

- Table column sorting
- Table column positions
- Auto refresh interval selection check box
- Auto refresh interval value

Tasks

The **Tasks** menu lets you manage, monitor, and perform other tasks.

The **Management** section of the **Tasks** menu provides access to the following options:

- Zone administration

  Zone information is collected from the selected switch. If an ACL-based FCS policy is in effect, zoning can be administered only from the primary fabric configuration server (FCS) switch. Refer to "Zoning management" on page 121 for more information.

- Switch administration
- Port administration
- Admin Domain administration
- FCR (present only on the base switch when the Virtual Fabrics capability is enabled.)
- iSCSI administration
- Fabric Watch

  NOTE
  Some of these functions require a license key to activate them.

The **Monitor** section of the **Tasks** menu provides access to the following options:

- Performance monitoring—You must use Web Tools with the EGM license to perform performance monitoring operations; otherwise, access to this feature is denied and an error message displays.
- Name Server information—This feature is available with Web Tools and Web Tools with the EGM license. Name Server information is collected from the selected switch. Refer to "Displaying the Name Server entries" on page 56 for more information.

The **Other** section of the **Tasks** menu provides access to Telnet tools.

NOTE
There are certain Fabric OS features that are available only on particular switch types, and the system displays only the icons that are appropriate for the switch type.

Fabric Tree

Fabric Tree displays all switches in the fabric, even those that do not have a Web Tools license and that are not owned by your selected Admin Domain. Switches that are not owned by the Admin Domain are shown in the Fabric Tree with switch status. Fabric Tree does not display switches segmented before you opened Web Tools.

Use the drop-down menu at the top of the Fabric Tree area to view switches in the Fabric Tree by switch name, IP address, or WWN. The background color of the switch icon indicates the current status of the switch. You can rest on the cursor over a switch to display the IP address and current status. To manually refresh the status of a switch within the fabric, right-click the switch in the Fabric Tree and select **Refresh**.

Although Fabric Tree displays all the switches in the fabric, you can manage switches that support 6.1 and later versions (does not requires Web Tools license), the versions earlier to 6.1 (requires a Web Tools license), and, if applicable, an EGM license installed. Other switches must be managed through the Fabric OS command line interface (CLI) or another management application or by using DCFM Professional Edition or Enterprise Edition.

Changing the Admin Domain context

The **Admin Domain** field displays the currently selected Admin Domain and allows you to change to a different one. The ability to change Admin Domain context requires that the EGM license is enabled on the switch. Otherwise, the following error message displays (Table 9).

FIGURE 9   **Missing EGM license**

If you are logged in to Web Tools without the EGM license, you must log in again using a specific Admin Domain.

Figure 10 displays the login dialog box. After you log in, all Admin Domains assigned to you are available in the drop-down menu (Figure 11). For most administrative tasks, you must be in either AD0 or the physical fabric.
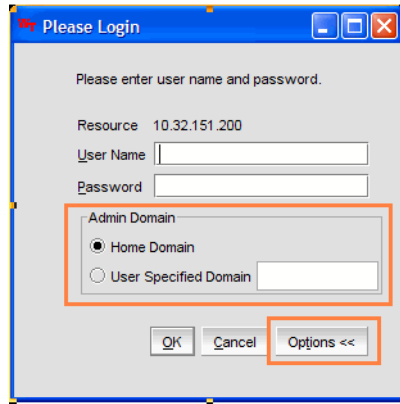
FIGURE 10   Login dialog box with Admin Domain

Figure 11 displays the **Admin Domain** context drop-down menu highlighted for changing the Admin Domain context.
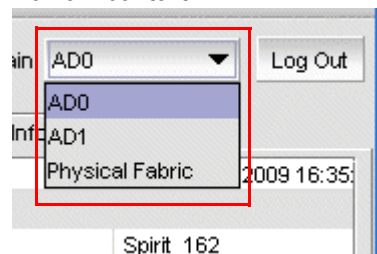
FIGURE 11   Changing the Admin Domain context

When changing the Admin Domain context, the option for selecting AD from the drop-down list is not available if the EGM license is not present.

Use the following procedure to change the Admin Domain context.

1. Select an **Admin Domain** from the **Admin Domain** menu.

2. Click **OK** in the confirmation window.

   Switch Explorer refreshes to display the new Admin Domain context. You can monitor the progress using the progress bar.

   The system displays a list of all open windows. You can choose to change the Admin Domain, which closes all the open windows, or cancel the action and return to Switch Explorer.

   NOTE
   The Telnet window and the Fabric Details are not AD-filtered and do not need to be closed.

Switch View buttons

The Switch View buttons let you access the following switch information:

- **Status**—Click the button to view the status of the switch.
- **Temperature**—Click the button to view temperature monitors.
- **Power**—Click the button to view power supply information.
- **Fan**—Click the button to view the status of the switch fans.
- **Beaconing**—Click this button to enable or disable beaconing and to view the status of beaconing from the button's icon.
- **Legend**—Click the button to view the legend for the Switch View.
- **Switch Status Policy**—Click this button to define the Switch Status Policy.

NOTE
For all status displays based on errors per time interval, any errors cause the status to show faulty until the entire sample interval has passed.

Switch View

You can click the small right arrow towards the left of the **Switch Event** tab to display the Switch View. The Switch View displays a graphical representation of the switch, including a real-time view of switch and port status. Refer to area 8 in Figure 7 on page 19.

NOTES: With the upgrade license installed:

- For 7500E, all the FC ports are enabled
- For 7800, all FC ports and 6 GbE ports are enabled

  Without the upgrade license installed:

- For 7500E, FC ports 2-15 are disabled
- For 7800, 4 FC ports and 2 GbE ports are enabled

NOTE
Blades are graphically represented in the Web Tools GUI. They are vertical in the DCX, and horizontal in the DCX-4S.

The Switch View display is refreshed at 15 second intervals. However, the initial display of Switch Explorer might take from 30 to 60 seconds after the switch is booted. Refresh rates are fabric-size dependent. The larger the fabric, the longer it takes to poll the fabric and refresh the view. F_Port and L_Port connection changes refresh immediately.

Auto refresh intervals may not be exactly 15 seconds. The refresh rate varies depending on the activity in the fabric and on the host system you are using.

Port representations

The ports in the Switch View show the port type. Borders around the accessible ports indicate that SFP modules are present. A colored border indicates the status of the port; for example, a green border indicates that the port is connected and traffic is flowing. Ports that are not accessible do not display the port type and do not have borders.

The port LEDs in the Switch View match the LEDs on the physical switch. However, the blink rate of the LEDs in the Switch View does not necessarily match the blink rate of the LEDs on the physical switch. Refer to "Port LED interpretation" on page 146 for more information.

Right-click a port in Switch View to get a menu that opens the **Port Administration** window, allowing you to view detailed information about the port. From **Port Administration**, you can access information on all other ports. Refer to Chapter 6, "Managing Ports" for more information.

NOTE
For detailed information on ISL Trunking, F_Port Trunking, and long distance, you must install the EGM license on the switch; otherwise access to these features is denied and an error message is displayed.

If the selected Admin Domain does not include ownership of some ports that are physically present on the switch, these ports are represented as black rectangles with horizontal gray bars indicating they are not accessible.

E_Ports are visible in all domains. You cannot open the **Port Administration** window by clicking these ports. For the Brocade DCX, the Port Admin view is launched for ICL ports.

USB port representation

For switches with USB ports, the **USB Storage Management** view is launched for USB ports (Figure 12).
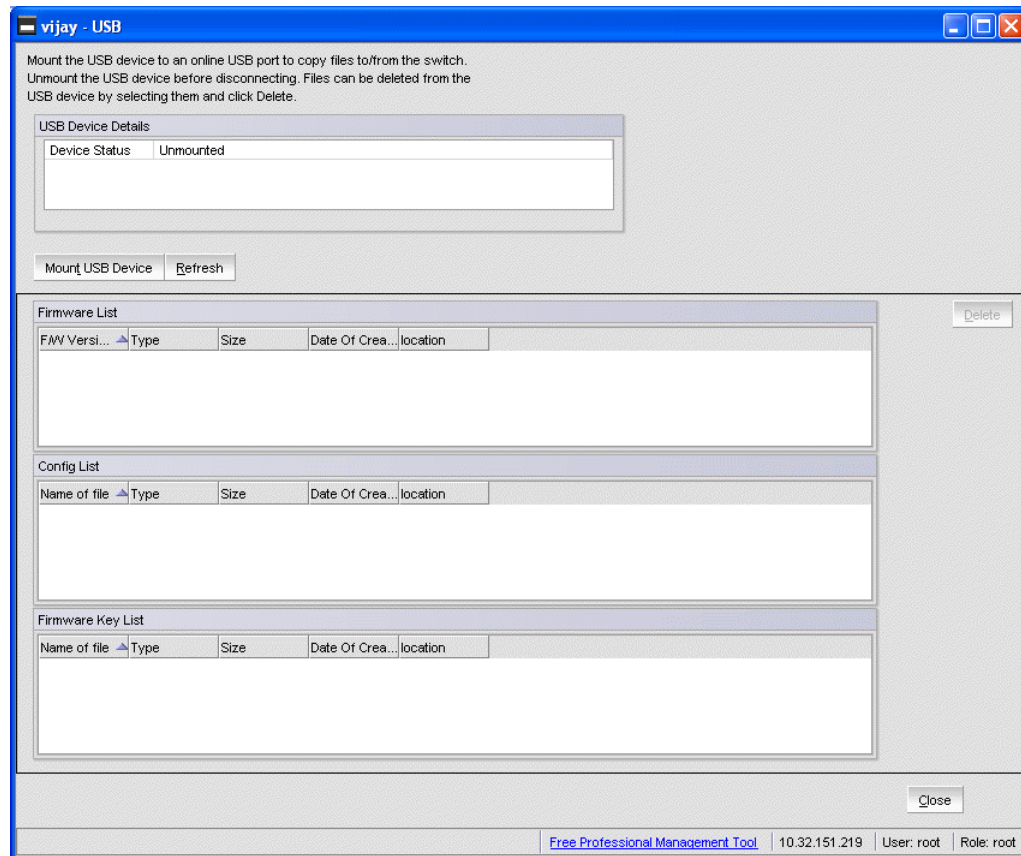
FIGURE 12   USB port storage management

NOTE
Click the USB port on the switch to launch the **USB Storage Management** window.

Switch Events and Switch Information

Switch Events and Switch Information display as tab forms under Switch View. The information in the Switch Information View is polled every 60 seconds.

The **Switch Information** tab displays information about the following items:

- **Switch**
    - Name   Name of the switch.
    - Status   Status of the switch.
    - Fabric OS Version   Fabric OS version of the switch.
    - Domain ID   Domain ID of the switch.
    - WWN   World Wide Name of the switch.
    - Type   Type of the switch.
    - Role   Role of the switch.

    The following information is specific to Virtual Fabrics:
    - Base Switch   Indicates whether or not the logical switch can act as a base switch.
    - Default Switch   Indicates whether or not the logical switch is the default logical switch.
    - Allow XISL Use   Indicates whether or not the logical switch is allowed to connect to other logical switches using an extended inter-switch link (XISL).

- **Ethernet**
    - Ethernet IPv4   Ethernet IPv4 address of the switch.
    - Ethernet IPv4 subnet mask   Ethernet IPv4 subnet mask address of the switch.
    - Ethernet IPv4 gateway   Ethernet IPv4 gateway address of the switch.
    - Ethernet IPv6   Ethernet IPv6 address of the switch.

- **FC**
    - IPFC IPv4   Fiber Channel IPv4 address.
    - IPFC IPv4 subnet mask   Fiber Channel IPv4 subnet mask address.

- **Zone**
    - Effective Configuration   Indicates whether zone configuration is enabled or not.

- **Other**
    - Manufacturer serial number   Displays the serial number of the manufacturer.
    - Supplier serial number   Displays the serial number of the supplier.
    - License ID   Displays the license ID.

- **RNID**

    - Type   Type of the switch.
    - Model   Model of the switch.
    - Tag   Tag of the switch.
    - Sequence number   Sequence number of the switch.
    - Insistent Domain ID   Current status of the Insistent Domain ID mode of the switch.
      Mode
    - Manufacturer   Manufacturer of the switch.
    - Manufacturer Plant   Plant where the switch was manufactured.

- **Name Server**

    - Capability   Capabilities of the switch.
    - Device Type   Displays the device type.
    - Unit Type   Displays the unit type.
    - Model   Model of the switch.
    - Tag   Tag of the switch.
    - Sequence number   Sequence number of the switch.
    - Manufacturer   Manufacturer of the switch.
    - Manufacturer Plant   Plant where the switch was manufactured.

For more information, refer to "Displaying switch information" on page 141.

Free Professional Management tool

You can use the Professional Management tool with Web Tools to view connectivity for each fabric, to back up and restore last-known configurations, and more. You can also use it with the Enhanced Group Management license to manage groups of switches, download firmware, manage security settings, and deploy configurations across groups of switches. Contact your preferred storage supplier to get a complimentary copy of the Professional Management tool.

Launch the install wizard for the free Professional Management tool through the link located at the bottom of the Switch Explorer.

Displaying tool tips

When you rest the cursor over a Web Tools button, the system displays a brief description of the button. If you rest the cursor over most components, the system displays *tool tip* information about the component.

In the Fabric Tree you can rest the cursor over a switch to view its type, Ethernet IP, IPFC, and status of the switch.

In Switch View, you can rest the cursor over a blade to view the blade ID and its status. It is easier to use the top of the blade to display the tool tip so that you do not inadvertently display the port tool tips. Firmware versions and IP addressing are displayed for CP blades.

When you rest the cursor over a port, you can view the port number, port index, port type (E, F, L, or U_Port), port status (online or offline), and port state (in-sync, no_sync, no light, or no module). If you right-click the port, the system displays the tool tip information and the port world-wide name.

Right-click options

You can right-click a port to quickly perform some basic port administration tasks, as shown in Figure 13.
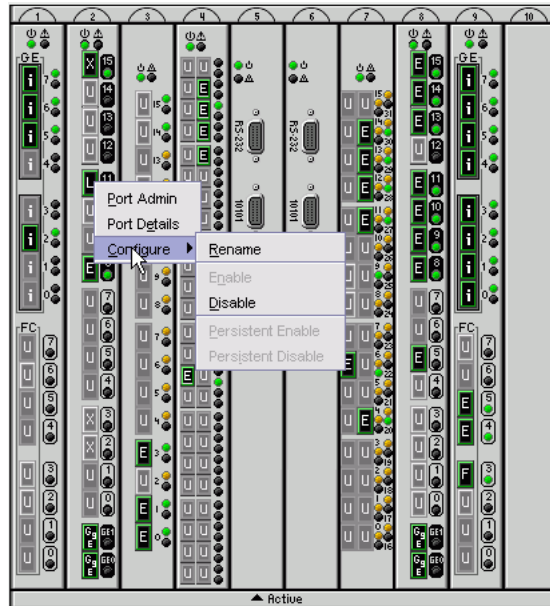
FIGURE 13   Right-click menu for ports (from Switch Explorer)

The tasks are:

- The **Port Admin** option displays the **Port Administration** window.

- The **Port Details** option displays read-only information about a port, without opening the **Port Administration** window. You can right-click on the table content to export or copy the information from the **Port Details** window.

- The **Configure** option provides another menu of options to allow you to rename, enable, and disable ports, and to set persistent enable or disable without opening the **Port Administration** window.

Refresh rates

Different panels of Web Tools refresh at different rates.

The refresh, or polling, rates listed in this section and throughout the book indicate the time between the end of one polling period and the start of the next, and not how often the screen is refreshed. A refresh rate of 15 seconds does not ensure that a refresh occurs every 15 seconds. It ensures that the time between each refresh activity is no more than 15 seconds.

Autorefresh intervals might be not be exactly 15 seconds. The refresh rate varies depending on the activity in the fabric and on the host system you are using. Following are some variables you should consider when refreshing the fabric:

- Retrieval time increases when you are in a large fabric because there is more data to retrieve from the switches.
- Processor speed of the system you are using may slow down the refresh rate.
- OS-Job Scheduling if you are using a host-system in the data center impacts the refresh rate.
- JVM-Performance can contribute to causing interval differences between what is on-screen and how long it is actually taking.

For these reasons, the time displayed in the port statistics tab might not be refreshed as expected. The counter time indicates only that "this statistics data is retrieved from the switch in this time period." To ensure the correct information, the time field is updated along with the port statistics data after every refresh.

The refresh rates are different for each module. Table 6 lists polling rates by module. Though these rates are sample rates, they correctly illustrate variance in the refresh rates throughout Web Tools.

TABLE 6   Polling rates

Module	Polling rate
Name Server	User-defined; 15 sec minimum
Zoning Database	60 sec
Fabric Watch	45 sec
Performance Monitor (This feature requires the EGM license.)	30 sec
Port Management	60 sec
FC Routing	45–90 sec, depending on network traffic

Displaying switches in the fabric

If your fabric has more than one switch, you can open Web Tools from one switch and then access other switches. You can also launch Web Tools from the DCFM client as Element Manager. This lets you manage Web Tool requests in the case where the fabric is in a private network.

Launch Web Tools from DCFM if you need to access the fabric from a host that is not in the same network and does not have direct access to the fabric.

NOTE
If you open switches running Fabric OS v4.4.x or later from a Fabric Tree displayed for a version earlier than a v4.4.x switch, some features might be disabled.

Use the following procedure to display switches in the fabric.

1. Open Web Tools as described in "Opening Web Tools" on page 9 and log in to the switch.

   Switch Explorer for your current switch.

2. If the Fabric Tree is not expanded, click the plus sign (+) in the Fabric Tree to view all the switches in the fabric.

3. Click a switch in the Fabric Tree.

   A separate browser dialog box displays the selected switch. (If the launch switch is running a Fabric OS version earlier than v5.0.1, the selected switch displays in the same browser window.)

   The graphic of the selected switch displays in Switch View. Additional switch information displays in the Switch Events and Switch Information dialog box.

Working with Web Tools: recommendations

Brocades makes the following recommendations for working with Web Tools:

* If you receive an error when saving changes in the **Switch Administration** window, note the error messages, refresh the window, and make your changes again. Do not continue making changes without refreshing the window and determining which changes were saved correctly.

* In a fabric containing switches and directors running different versions of firmware, use the switches or directors with the latest firmware versions to control the fabric.

* If switches are accessed simultaneously from different connections (for example, Web Tools, CLI, and API), changes from one connection might not be updated to the other, and some modifications might be lost. Make sure that, when you connect with simultaneous multiple connections, you do not overwrite the work of another connection.

* Several tasks in Web Tools make fabric-level changes, such as the tasks in Zone Administration. When executing fabric-level configuration tasks, wait until you have received confirmation that the changes are implemented before executing any subsequent tasks. For a large fabric, this can take several minutes.

* Some data collection and processing operations in the iSCSI Gateway module might take a long time to complete, especially in large fabrics or fabrics with large numbers of defined Discovery Domains and Discovery Domain Sets. In most cases, progress bars are provided. Allow the application a sufficient amount of time (30-40 seconds) to collect and display data before taking any action or assuming the application is "hanging."

* A maximum of five simultaneous HTTP sessions to any one switch is recommended. An HTTP session is considered a Fabric Manager or Web Tools connection to the switch.

Opening a Telnet or SSH client window

When you open a Telnet or SSH client window, it connects to the IP interface of the switch. You cannot connect to a CP blade on a director switch through a Telnet or SSH client window opened from Web Tools, even when the blade has an IP address and supports Telnet sessions. Refer to the *Fabric OS Command Reference* for information about the Telnet commands.

NOTE
Internet Explorer 7.0 default settings disable Telnet functionality. If you are using Internet Explorer 7.0, you must make the appropriate changes in the registry to open the Telnet window.

Use the following procedure to open a Telnet or SSH client window.

1. Select a switch in Fabric Tree.

   You are prompted to log in if the OS is version 5.3.0 or later; otherwise, the selected switch displays in the Switch View.

2. Select **Telnet/SSH Client** in the **Other** section of the **Tasks** panel. The **Preference** dialog box displays.

3. Select the client by clicking **Telnet** or **SSH**.

4. Enter the Telnet or SSH path, as defined for your implementation.

   To avoid the need to remember and key in the path, you can store the path on your PC and browse to the location. Clicking the button to the right of the field initiates the browse capability.

5. Click **OK**.

   The Telnet or SSH window displays.

6. Enter your user credentials at the login prompt.

7. To close the session, enter **exit** at the prompt and press the **Enter** key.

Collecting logs for troubleshooting

If you encounter problems using the Web Tools interface, collect Java logs for use in troubleshooting. From Microsoft Windows, perform this procedure.

1. Open **Control Panel** and select **Java**.

2. Click on the **Advanced** tab.

3. Expand **Java console**.

4. Select **Show console**.

5. Restart Web Tools.

   The Java console displays, along with the Web Tools opening page.

6. Perform the Web Tools operation that caused the problem.

7. Collect the logs shown on the Java console.

8. If you no longer want to see the Java console when you start Web Tools, go back to the **Control Panel**, repeat steps 1 and 2, and then deselect **Show console**.

Managing Fabrics and Switches

In this chapter

Fabric and switch management overview

Most of the management tasks described in this chapter are accessed through the **Switch Administration** window. Information in the **Switch Administration** window is retrieved from the selected switch, as shown in

If the switch is not a member of the selected Admin Domain, most tabs in the **Switch Administration** window display in read-only mode, regardless of your permission level. The **User** tab is editable because most of its information does not require switch membership in the current Admin Domain.
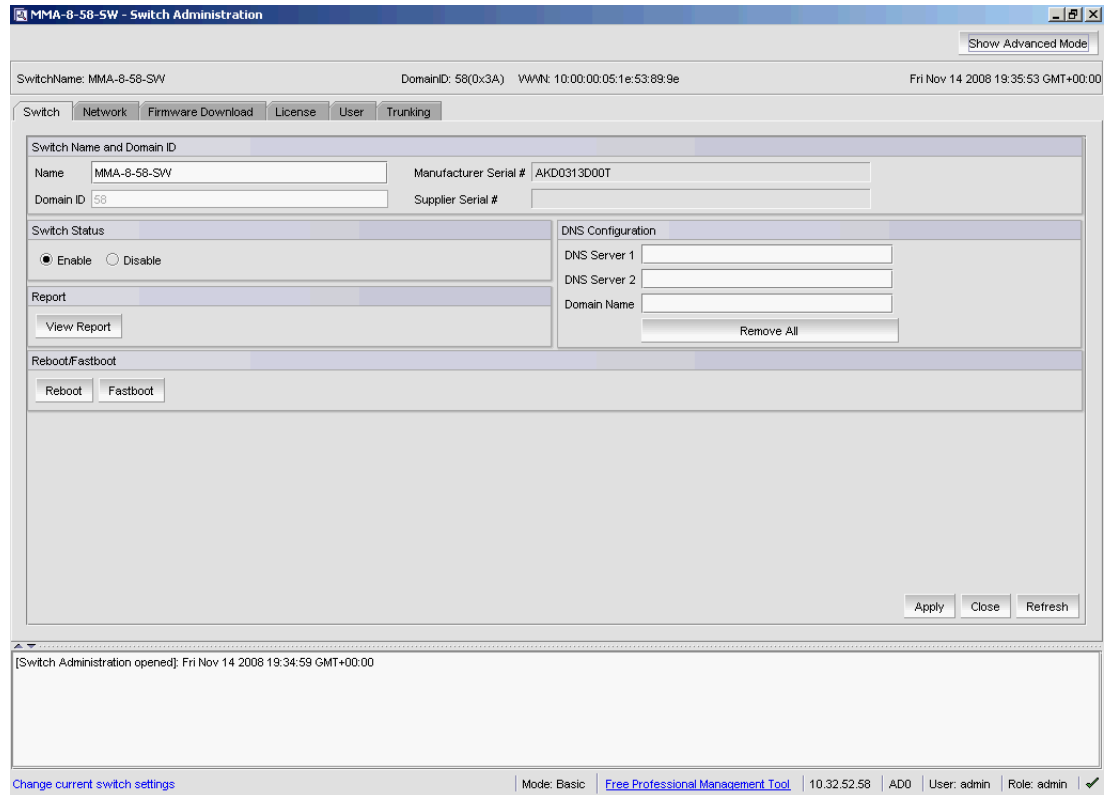
FIGURE 14   Switch Administration window, Switch tab

With the exception of switch time, information displayed in the **Switch Administration** window is not updated automatically by Web Tools. To update the information displayed in the **Switch Administration** window, click the **Refresh** button.

ATTENTION
Most changes you make in the **Switch Administration** window are buffered, and are *not* applied to the switch until you save the changes. If you close the **Switch Administration** window without saving your changes, your changes are lost. To save the buffered changes you make in the **Switch Administration** window to the switch, click **Apply** before closing the module or before switching to another tab.

The **License** tab, **Firmware Download** tab, and the **Security Policies** tab are exceptions. The changes you make on these tabs take effect immediately and there is no **Apply** button. There is an **Apply** button in all the subtabs of security policies except ACL.

You can also use Telnet commands to perform management tasks. Refer to *"Opening a Telnet or SSH client window"* on page 31 for information on how to launch a Telnet window using Web Tools.

Opening the Switch Administration window

Most of the management procedures in this chapter are performed from the **Switch Administration** window.

To open the **Switch Administration** window, perform the following steps.

1. Select **Tasks** > **Manage** > **Switch Admin**.

   The **Switch Administration** dialog box displays in basic mode, as shown in on page 34. The basic mode displays the "basic" tabs and options.

2. Click **Show Advanced Mode** to see all the available tabs and options.

Configuring IP and subnet mask information

Before proceeding, collect all the information you need to configure the Ethernet IP interface. This includes the subnet mask, gateway IP address, or IPFC, and subnet mask for your system. When you configure or change the Ethernet IP, subnet mask, gateway IP, or IPFC, and subnet mask from Web Tools, there is a normal loss of network connection to the switch. Close all current windows and restart Web Tools with the new IP address.

NOTE
The IPFC address is specific for each logical switch. The IPFC address is set to FC0 for switches that do not support Virtual Fabrics.

Use the following procedure to configure the IP and subnet mask information.

1. Select the **Network** tab.

2. In the appropriate IP address section, enter the IP address you want to use for the IP interface.

   Use the **IPv4 Address** section or the **IPv6 Address** section to specify IP addresses.

3. In the **IPv4 Address** section:

   - In the **Ethernet IP** field, enter the Ethernet IP address.
   - In the **Ethernet Mask** field, enter the Ethernet mask address.
   - In the **GateWay IP address** field, enter the gateway IP address.

4. In the **IPv6 Address** section, in the **Ethernet IPv6** field, enter the Ethernet IP address.

5. You can also enable automatic configuration of IPv6 addresses by selecting **Enable IPV6 Auto Configuration**.

   The automatically generated IPv6 addresses are displayed under **Auto Configured IPV6 Addresses**. Eight auto-configured addresses are created per switch, and up to 24 for a 48000, DCX, or DCX-4S chassis (eight per chassis, and eight per each installed CP).

Configuring Auto Refresh

The Netstat Performance window displays the details about Ethernet management port statistics like the Interface, MTU, Met, RX-OK, RX-ERR, RX-DRP, RX-OVR, TX-OK, TX-ERR, TX-DRP, TX-OVR, and Flag.

Use the following procedure to configure Auto Refresh.

1. Open the **Switch Administration** window as described in .

2. Select the **Network** tab.

3. Click **Netstat Performance**.

4. Select the **Auto Refresh** check box to automatically refresh the statistics details.

    Clear the check box to disable auto refresh.

5. When enabled, enter the interval time in seconds in the **Auto-Refresh Interval** field.

    The statistics details are automatically refreshed, based on the configured time interval.

Configuring a syslog IP address

The syslog IP represents the IP address of the server that is running the syslog process. The syslog daemon reads and forwards system messages to the appropriate log files or users, depending on the system configuration. When one or more IP addresses are configured, the switch forwards all error log entries to the syslog on the specified servers. Up to six servers are supported. Refer to *Fabric OS Administrator's Guide* for more information on configuring the syslog daemon.

Use the following procedure to configure a syslog IP address.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Network** tab.

3. In the **Syslog IP's Configuration** section, in the **New IP** field, enter an IP address in either IPv4 or IPv6 format.

4. Click **Add**.

    The new IP address displays in the Syslog IP area.

5. Click **Apply**.

Removing a syslog IP address

Use the following procedure to remove a syslog IP address.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Network** tab.

3. Select a syslog IP in the table and click **Remove**.

    You can click **Clear All** to remove all of the syslog IP addresses from the table.

4. Click **Apply**.

Configuring IP Filtering

Web Tools provides the ability to control what client IP addresses may connect to a switch or fabric.

Use the following procedure to set up IP Filtering.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Security Policies** tab.

3. Select **IPFilter** on the **Security Policies** menu.

4. Click **Create Policy**.

   The **Create IP Filter Policy** dialog box displays.

5. Enter a policy name, select a policy type, and then click the **Add Rule** button.

6. Enter the rule order and source IP address, and modify the service or destination port, protocol, and action as necessary.

7. Click **OK**.

   After you create a policy, you can use the following controls on this tab to manage the policies:

   - The **Edit Policy** button lets you select an existing policy and make changes to it.
   - The **Show Policy** button lets you view the details of the policy in a read-only window.
   - The **Delete Policy** button lets you delete a policy.
   - The **Clone Policy** button lets you copy a policy. Use this feature when you want to create similar policies. After you create a clone, you can edit the policy to make the appropriate changes.
   - The **Activate Policy** button lets you make an existing policy active.
   - The **Distribute Policy** button lets you distribute a policy to various switches.
   - The **Accepts Distribution** check box lets you set the policy to accept or reject distributions.

Blade management

Web Tools provides the ability to enable and disable blades, and to set slot-level IP addresses for blades. The procedure in this section applies only to the Brocade 48000 and Brocade DCX and DCX-4S enterprise-class platforms.

Enabling or disabling a blade

Use the following procedure to enable or disable a blade.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Blade** tab (Figure 15).

The **Firmware Version** columns display the firmware loaded onto each blade. A blade can have more than one firmware image loaded onto it. The **Enable Blade** column in the **Blade** tab pane indicates whether the blade is enabled.
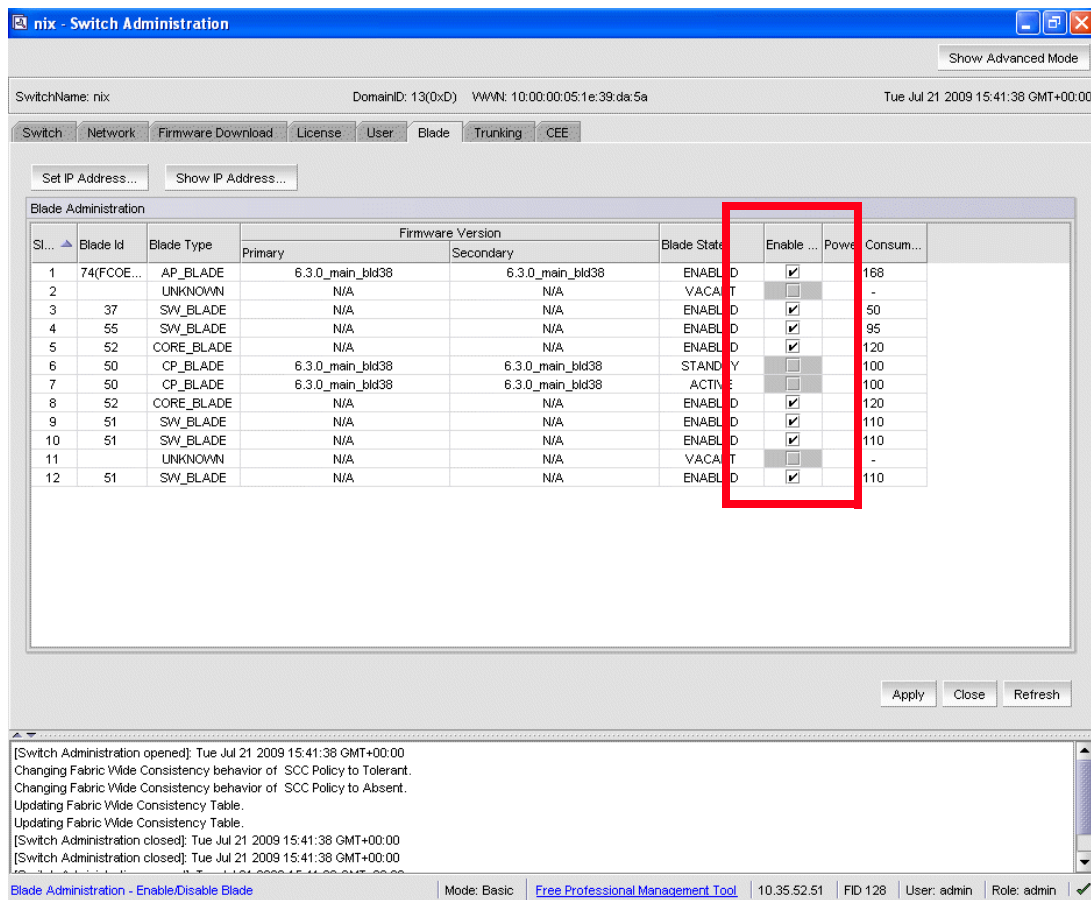
FIGURE 15   Blade tab

3.   Select the **Enable Blade** check box for each blade you want to enable.

Clear the check box to disable the blade. You cannot enable or disable the CP blades.

4.   Click **Apply**.

Setting a slot-level IP address

Use the following procedure to set an IP address.

1.   Open the **Switch Administration** window as described in *"Opening the Switch Administration window"* on page 35.

2.   Select the **Blade** tab.

3.   Click **Set IP address**.

4.   Select a slot number from the list.

5.   Enter the IP address, subnet mask, and Gateway IP address.

6.   Select a type from the list.

7. Click **Add** to add the new entry to the table.

   When you click **Add**, the values remain in the fields. The **Clear Gateway** and **Clear IP** buttons are available for clearing fields in the table.

   NOTE
   To remove a configuration, select a row in the table and click **Delete**.

8. Click **Apply** to save the values currently shown in the table or click **Cancel** to close the dialog box without saving any of your changes.

9. To update the switch with your changes, update the table using the **Add** and **Delete** buttons, and then click **Apply**.

Viewing IP addresses

If you want to view the IP addresses configured on the switch for the currently populated slots, use the **Show IP Address** button.

Use this procedure to display the IP addresses.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Blade** tab.

3. Click **Show IP Address**.

4. Scroll through the list to view all the information.

5. When you are finished, click **Close**.

Switch configuration

Use the **Switch** tab of the **Switch Administration** window to perform basic switch configuration. Figure 14 on page 34 displays an example of the **Switch** tab.

Enabling and disabling a switch

You can identify whether a switch is enabled or disabled in the **Switch Administration** window by looking at the lower-right corner. If you rest the cursor over the icon, the system displays text that indicates the status of the switch.

Use this procedure to enable or disable a switch.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Switch** tab.

3. In the **Switch Status** section, click **Enable** to enable the switch or **Disable** to disable the switch.

4. Click **Apply**.

   The system displays a confirmation window that asks if you want to save the changes to the switch. You must click **Yes** to save the changes.

Changing the switch name

Switches can be identified by IP address, domain ID, World Wide Name (WWN), or switch names. Names must begin with an alphabetic character, but otherwise can consist of alphanumeric, hyphen, and underscore characters. The maximum number of characters is 30, unless FICON mode is enabled. When FICON mode is enabled, the maximum number of characters is 24.

NOTE
Some system messages identify a switch service by the chassis name. If you assign meaningful chassis names and switch names, system logs are easier to use.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Switch** tab.

3. Enter a new name in the **Name** field and click **Apply**.

Changing the switch domain ID

Although domain IDs are assigned dynamically when a switch is enabled, you can request a specific ID to resolve a domain ID conflict when you merge fabrics.

Use the following procedure to change the switch domain ID.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Disable the switch, as described in "Enabling and disabling a switch" on page 40.

3. Select the **Switch** tab.

4. Enter a new domain ID in the **Domain ID** field.

    The domain ID range depends on the switch interop mode:

    - For IM0, the range is between 1 and 239.
    - For IM2, it depends on the selected offset value.
    - For IM3, it depends on the selected offset value.

5. Click **Apply**.

6. Enable the switch, as described in "Enabling and disabling a switch" on page 40.

Viewing and printing a switch report

The switch report includes the following information:

- A list of switches in the fabric
- Switch configuration parameters
- A list of ISLs and ports
- Name Server information
- Zoning information
- SFP serial ID information

Perform the following steps to view or print a report.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Switch** tab.

3. Click **View Report**.

4. In the new window that displays the report, view or print the report using your browser.

Switch restart

When you restart the switch, the restart takes effect immediately. Ensure that there is no traffic or other management on the switch, because traffic is interrupted during the restart; however, frames are not dropped. Be sure to save your changes before the restart, because any changes not saved are lost.

Performing a fast boot

A fast boot reduces boot time significantly by bypassing the power-on self test (POST).

Use the following procedure to perform a fast boot.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Click **Fastboot**.

3. On the **Fastboot Confirmation** window, click **Yes** to continue.

4. Click **Apply**.

Performing a reboot

Use the following procedure to reboot the CP and execute the normal power-on booting sequence.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Click **Reboot**.

3. On the **Reboot Confirmation** window, click **Yes** to continue.

4. Click **Apply**.

System configuration parameters

You must disable the switch before you can configure fabric parameters.

You can change the following system configuration parameters:

- Switch fabric settings
- Virtual channel settings
- Arbitrated loop parameters
- System services

- Signed firmware

WWN-based Persistent PID assignment

WWN-based PID assignment allows you to configure a PID persistently using a device's WWN. When the device logs into the switch, the PID is bound to the device WWN. If the device is moved to another port in the same switch, or a new blade is hot plugged, the device receives the same PID (area) at its next login. For information on configuring WWN-based PID assignment, refer to "Configuring fabric settings" on page 43.

This feature is deactivated by default. When the feature is enabled, bindings are created dynamically; as new devices log in, they automatically enter the WWN-based PID database. The bindings exist until you explicitly unbind the mappings through the CLI or change to a different addressing mode. If there are any existing devices when you enable the feature, you must manually enter the WWN-based PID assignments through the CLI.

Once WWN-based PID assignment is enabled you must manually enter the WWN-based PID assignments through the CLI for any existing devices. Any new devices logging in are automatically entered in the WWN-based PID database. Current WWN-based PID bindings are cleared when you change to a different addressing mode.

PID assignments are supported for a maximum of 4096 devices; this includes both point-to-point and NPIV devices. The number of point-to-point devices supported depends directly on the areas available. For example, 448 are available on an enterprise-class platform and 256 are available on switches. When the number of entries in the WWN-based PID database reaches the number 4096 or areas are used up, the oldest unused entry is purged from the database to free up the reserved area for the new FLOGI. Refer to Table 7 for complete information.

TABLE 7   Switches that support WWN-based Persistent PID on Web Tools

Platform	VF	Default switch	Logical switch	Area mode	FICON mode
DCX/DCX-4S	Enabled	Yes, if dynamic area addressing is enabled in the default switch.	Yes	0	FMS is not supported
				1	Can be set
				2	Not supported
Brocade 5100 Brocade 5300	Enabled	Yes	Yes	Default-8 bit dynamic	Configurable
Brocade 300 Brocade 5100 Brocade 5300	Disabled	N/A	N/A	Default-8bit dynamic	Configurable

Configuring fabric settings

Perform the following steps to configure the fabric settings.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Configure** tab.

3. Select the **Fabric** subtab.

4. Make the fabric parameter configuration changes.

5. Click **Apply**.

6. Enable the switch as described in "Enabling and disabling a switch" on page 40.

Fabric settings

Configure the following fabric settings on the **Fabric** subtab of the **Configure** tab:

BB Credit	The buffer-to-buffer credit is the number of buffers available to attached devices for frame receipt. The default BB Credit is 16. The range is 1–27.
R_A_TOV	Resource allocation timeout value (in milliseconds). This variable works with the E_D_TOV to determine switch actions when presented with an error condition. The default is 10000. The possible range is (2*E_D_TOV) –120000. Values must be multiples of 1000.
E_D_TOV	Error detect timeout value (in milliseconds). This timer is used to flag a potential error condition when an expected response is not received within the set time. The valid range is 1000 - (R_A_TOV/2)
Addressing mode	Displays the addressing mode present in the switch.
Datafield size	The largest possible data field size (in bytes). The valid range is 256–2112.
Sequence Level Switching	Select this box to enable frames of the same sequence from a particular group to be transmitted together. When this option is not selected, frames are transmitted interleaved among multiple sequences. Under normal circumstances, sequence-level switching should be disabled for better performance. However, some host adapters have issues when receiving interleaved frames among multiple sequences.
Disable Device Probing	Set this mode only if the switch N_Port discovery process (PLOGI, PRLI, INQUIRY) causes an attached device to fail. When set, devices that do not register with the Name Server are not present in the Name Server database.
Per-Frame Routing Priority	Select whether to select per-frame routing priority. When enabled, the virtual channel ID is used in conjunction with a frame header to form the final virtual channel ID.
Suppress Class F Traffic	Applies only if VC-encoded address mode is also set. When selected, translative addressing (which allows private devices to communicate with public devices) is disabled.
Insistent Domain ID Mode	Set this mode to make the current domain ID insistent across reboots, power cycles, and failovers. This mode is required fabric wide to transmit FICON data.
WWN-based Persistent PID	Set this mode to configure a PID persistently using a device's WWN. When the device logs into the switch, the PID is bound to the device WWN. Refer to "WWN-based Persistent PID assignment" on page 43.

Enabling insistent domain ID mode

Use the following procedure to enable insistent domain ID mode.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Disable the switch as described in "Enabling and disabling a switch" on page 40.

3. Select the **Configure** tab.

4. Select the **Fabric** subtab.

5. Select the **Insistent Domain ID Mode** check box.

6. Click **Apply**.

7. Enable the switch as described in "Enabling and disabling a switch" on page 40.

Configuring virtual channel settings

You can configure parameters for eight virtual channels (VC) to enable fine-tuning for a specific application. You cannot modify the first two virtual channels because these are reserved for switch internal functions.

ATTENTION
The default virtual channel settings are already optimized for switch performance. Changing the default values can improve switch performance, but can also degrade performance. Do not change these settings without fully understanding the effects of the changes.

VC Priority specifies the class of frame traffic given priority for a virtual channel.

Perform the following procedure to configure the virtual channel settings.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Disable the switch as described in "Enabling and disabling a switch" on page 40.

3. Select the **Configure** tab.

4. Select the **Virtual Channel** subtab.

5. Enter a value in the **VC Priority** field you want to change.

   The only valid numeric values for all fields are either "2" or "3".

6. Click **Apply**.

7. Enable the switch as described in "Enabling and disabling a switch" on page 40.

Configuring arbitrated loop parameters

Use the following procedure to configure arbitrated loop parameters.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Disable the switch as described in "Enabling and disabling a switch" on page 40.

3. Select the **Configure** tab.

4. Select the **Arbitrated Loop** subtab.

5. Select or clear the check boxes to enable or disable the corresponding arbitrated loop parameters.

6. Click **Apply**.

7. Enable the switch as described in "Enabling and disabling a switch" on page 40.

Arbitrated loop parameters

Configure the following arbitrated loop parameters on the **Arbitrated Loop** subtab of the **Configure** tab:

Send Fan Frames	Select this check box to specify that fabric address notification (FAN) frames are sent to public loop devices to notify them of their node ID and address.
Always Send RSCN	Following the completion of loop initialization, a remote state change notification (RSCN) is issued when FL_Ports detect the presence of new devices or the absence of pre-existing devices. Select this check box to issue an RSCN upon completion of loop initialization, regardless of the presence or absence of new or pre-existing devices.

Configuring system services

You can enable or disable FCP read link status (RLS) probing for F_Ports and FL_Ports. It is disabled by default.

Perform the following procedure configure system services.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Disable the switch as described in "Enabling and disabling a switch" on page 40.

3. Select the **Configure** tab.

4. Select the **System** subtab.

5. Select the **Disable RLS Probing** check box to disable RLS probing.

   or

   Clear the check box to enable RLS probing.

6. Click **Apply**.

7. Enable the switch as described in "Enabling and disabling a switch" on page 40.

Configuring signed firmware

When the firmware is downloaded to a device, the system can validate the firmware based on a configuration setting. By default, the signed firmware download is not validated.

NOTE
During the first download, the system ignores the signed firmware. After the first download, the public key is downloaded and then, in subsequent downloads, you can turn on the feature. You can view the public key on the **Firmware Download** tab in the **Switch Administration** window.

Perform the following procedure to configure the signed firmware.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Configure** tab.

3. Select the **Firmware** subtab.

4.   Select the **Enable Signed Firmware Download** check box.

5.   Click **Apply**.

Licensed feature management

The licensed features currently installed on the switch are listed in the **License** tab of the **Switch Administration** window, as shown in Figure 16. If the feature is listed, such as the EGM license, it is installed and immediately available. When you enable some licenses, such as ISL Trunking, you might need to change the state of the port to enable the feature on the link. For time-based licenses, the expiry date is included. Right-click a license key to export data, copy data, or search the table.

FIGURE 16   License tab

Activating a license on a switch

Before you can unlock a licensed feature, you must obtain a license key. You can either use the license key provided in the paperpack document supplied with switch software or refer to the *Fabric OS Administrator's Guide* for instructions on how to obtain a license key at the Brocade website (*www.brocade.com*).

Use the following procedure to activate a license.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **License** tab and click **Add**.

   The **Add License** dialog box displays.

3. Paste or enter a license key in the field.

4. Click **Add License**.

5. Click **Refresh** to display the new licenses in the **License** tab.

   Some licenses, such as the Trunking or the 7500E and 7800 upgrade license, do not take effect until the switch is restarted.

Assigning slots for a license key

This feature allows to increase the capacity without disrupting the slots that already have licensed features running.

NOTE
You can enable slot based licenses only on 10 Gigabit Ethernet (FTR_10G), Advanced Extension (FTR_AE), and Advanced FICON Acceleration (FTR_AFA) features.

Use the following procedure to assign slots for a license key.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **License** tab.

3. Select the license key for which you want to assign slots from the **License Administration** table.

   The **Assign Slots** window displays.

4. Select the slots you want to assign.

5. Click **OK.**

Removing a license from a switch

Use the following procedure to remove a license from a switch in the **Switch Administration** window.

ATTENTION
Use care when removing licenses. If you remove a license for a feature, that feature no longer works.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **License** tab.

3. Click the license you want to remove.

4. Click **Remove**.

Universal time-based licensing

After v6.3.0, Web Tools supports universal time-based licensing. Each universal key is for a single feature, and can be used on any product that supports the feature, for a defined trial period. At the end of the trial period, the feature gets disabled. You can extend the universal key license. For time-based licenses, the **Expiry Date** displays in the **License Administration** table (refer to).

The following features are supported for universal time-based license:

- Fabric
- Extended Fabric
- Performance Monitor
- Trunking
- High-Performance Extension over FCIP/FC
- Fabric Watch
- Integrated Routing
- Adaptive Networking
- Server Application Optimization

High Availability overview

High-Availability (HA) features provide maximum reliability and nondisruptive replacement of key hardware and software modules. High Availability is available only on the Brocade 48000 and Brocade DCX and DCX-4S enterprise-class platforms. Refer to the *Fabric OS Administrator's Guide* for additional information about High Availability.

The High Availability window, as shown in , displays information about the status of the HA feature on each control processor (CP), and enables you to perform CP failover.

The background color of the HA button indicates the overall status of high availability on the switch. The colors and their meanings are:

- **Green**—Healthy: HA Status is **HA enabled, Heartbeat Up, HA State synchronized**.
- **Yellow**—Disruptive mode: HA Status is **HA enabled, Heartbeat Up, HA State not in sync**.
- **Red**—HA is unavailable: HA Status is **Non-Redundant**.

Admin Domain considerations

HA is possible if the switch is a member of the current Admin Domain. If switch is not a member of current Admin Domain, the **Synchronize Services** and **Initiate Failover** buttons are unavailable.

Launching the High Availability window

Use the following procedure to launch the High Availability window.

1. Select a Brocade 48000, Brocade DCX, or Brocade DCX-4S from the Fabric Tree.

   The Switch View displays.

2.   Click the **HA** button in the Switch View.

The **High Availability** dialog box displays.
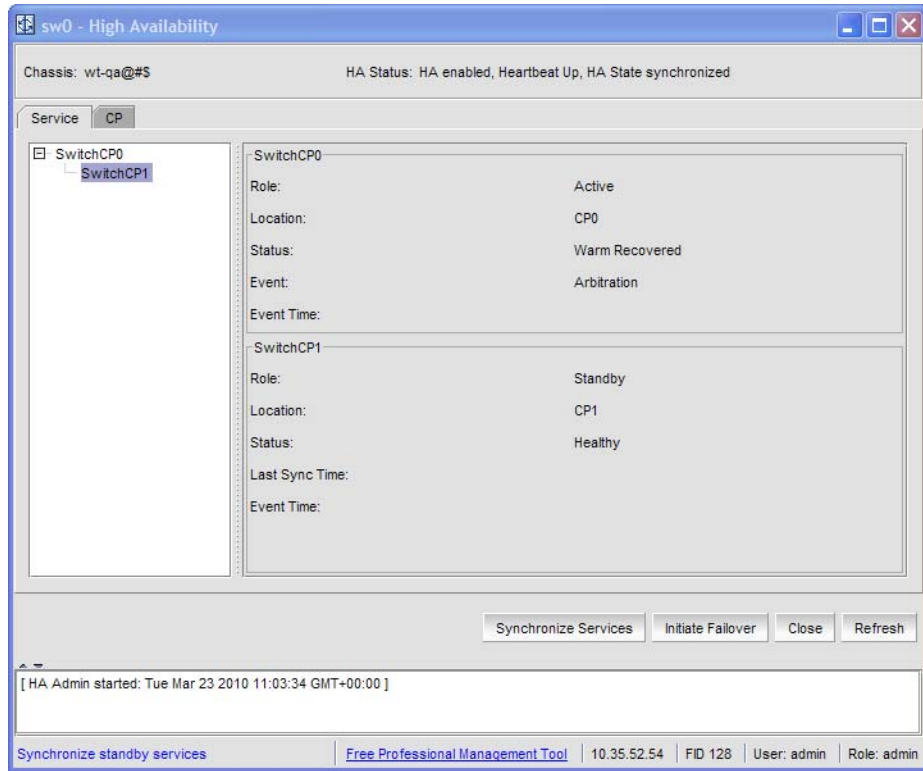
FIGURE 17   High Availability window, Service tab

FIGURE 18   High Availability window, CP tab

The **High Availability** window contains the following two tabs:

- The **Service** tab displays information about the switch. When the hardware is configured as a dual switch, the **Service** tab displays information about both switches.

- The **CP** tab displays information about slots. For Brocade 48000 and Brocade DCX-4S, CP blades are placed in slot 5 and slot 6, and for Brocade DCX, CP blades are placed in slot 6 and 7.

The **High Availability** window gets refreshed automatically. You can also click **Refresh** to update the information displayed in the **High Availability** window.

Admin Domain considerations

To open the **High Availability** window, the switch must be a member of your current Admin Domain. If the switch is *not* a member of the current Admin Domain, the **Synchronized Services** and **Initiate Failover** buttons are unavailable.

Synchronizing services on the CP

A nondisruptive CP failover is only possible when *all* the services are synchronized between both CPs.

Use the following procedure to synchronize services on the CP.

1. Open the **High Availability** window as described in "Launching the High Availability window" on page 49.

2. Verify that the **HA Status** field displays **HA enabled, Heartbeat Up, HA State synchronized**.

   If the **HA Status** field displays **HA enabled, Heartbeat Up, HA State synchronized** you are finished.

   If the **HA Status** field displays **HA enabled, Heartbeat Up, HA State not in sync**, continue with step 3.

3. Click **Synchronize Services**.

   The **Warning** dialog box displays.

4. Click **Yes** and wait for the CPs to complete a synchronization of services, so that a nondisruptive failover is ready.

5. Click **Refresh** to update the **HA Status** field.

   When the **HA Status** field displays **HA enabled, Heartbeat Up, HA State synchronized** a failover can be initiated without disrupting frame traffic on the fabric.

Initiating a CP failover

A nondisruptive failover might take about 30 seconds to complete. During the failover, all of the Web Tools windows and all associated child-windows are invalidated. You must close all Web Tools windows and open Web Tools again.

Use the following procedure to initiate a nondisruptive failover.

1. Open the **High Availability** window as described in "Launching the High Availability window" on page 49.

2. Verify that the **HA Status** field displays **HA enabled, Heartbeat Up, HA State synchronized** or **HA enabled, Heartbeat Up, HA State not in sync**.

3. Click **Initiate Failover**.

   The **Warning** dialog box displays.

4. Click **Yes** to initiate a nondisruptive failover.

5. When prompted, close the **Web Tools Switch Explorer** window and all associated windows, and re-open Web Tools.

Event monitoring

Web Tools displays fabric-wide and switch-wide events. Event information includes sortable fields for the following:

- Switch name
- Message number
- Time stamp
- Indication of whether the event is from a logical switch or a chassis
- The number of successive events of the same kind

- Severity level

- Unique message identifier (in the form *moduleID-messageType*)

- Detailed error message for root cause analysis

There are four message severity levels: Critical, Error, Warning, and Info. Table 8 lists the event message severity levels displayed on the **Switch Events** tab and explains what qualifies event messages to be certain levels.

On the **Switch Events** tab, you can click the **Filter** button to launch the **Filter Events** dialog box. The **Filter Events** dialog box allows you to define which events should be displayed on the **Switch Event**s tab. For more information on filtering events, refer to

TABLE 8   Event severity levels

Icon and level	Description
Critical	Critical-level messages indicate that the software has detected serious problems that will eventually cause a partial or complete failure of a subsystem if not corrected immediately. For example, a power supply failure or rise in temperature must receive immediate attention.
Error	Error-level messages represent an error condition that does not impact overall system functionality significantly. For example, error-level messages might indicate timeouts on certain operations, failures of certain operations after retries, invalid parameters, or failure to perform a requested operation.
Warning	Warning-level messages highlight a current operating condition that should be checked or it might lead to a failure in the future. For example, a power supply failure in a redundant system relays a warning that the system is no longer operating in redundant mode. The failed power supply must be replaced or fixed.
Info	Information-level messages report the current nonerror status of the system components, such as the online and offline status of a fabric port.

Displaying Switch Events

The **Switch Events** tab displays a running log of events for the selected switch. Switch events are polled and updated every 15 seconds; there is no refresh-on-demand option for switch events.

For two-switch configurations, all chassis-related events are displayed in the event list of each logical switch for convenience.

Perform the following procedure to display Switch Events.

1. Select the switch from the Fabric Tree.

   The Switch View displays.

2. Select the **Switch Events** tab, if necessary.

   NOTE
   You can click the column head to sort the events by a particular column, and drag the column divider to resize a column. You can also right-click a column heading to resize one or all columns, sort the information in ascending or descending order, or select which columns are displayed.

Filtering Switch Events

You can filter the fabric and switch events by time, severity, message ID, and service. You can apply either one type of filter at a time or multiple types of filters at the same time. When a filter is applied, the filter information displays at the bottom of the filtered information and the **Show All** link is available to allow you to view the information unfiltered.

Use the following the procedure to filter Switch Events.

1. Open the **Switch Events** tab as described in "Displaying Switch Events" on page 54.

2. Click **Filter**.

   The **Event Filter** dialog box displays.

3. To filter events within a certain time period:
   - Select the **From** check box and enter the start time and date in the fields.
   - Select the **To** check box and enter the finish time and date in the fields.
   - To filter events beginning at a certain date and time, select only the **From** check box and enter the start time and date.
   - To filter events up until a certain date and time, select only the **To** check box and enter the finish time and date.

4. Click **OK**.

   The filter is enabled and the window is refreshed to show the filtered information.

Filtering events by event severity levels

Use the following procedure to filter events by event severity levels.

1. Open the **Switch Events** tab as described in "Displaying Switch Events" on page 54.

2. Click **Filter**.

   The **Event Filter** dialog box displays.

3. Select **Level**.

4. Select the event levels you want to display.

5. Click **OK**.

   The filter is enabled and the window is refreshed to show the filtered information.

Filtering events by message ID

Use the following procedure to filter events by message ID.

1. Open the **Switch Events** tab as described in "Displaying Switch Events" on page 54.

2. Click **Filter**.

   The **Event Filter** dialog box displays.

3. Select **Message ID**.

4. Enter the message IDs in the associated field.

NOTE
You can enter multiple message IDs as long as you separate them by commas. You can enter
either the full message ID (moduleID-messageType) or a partial ID (moduleID only). The
message ID filtering is case-sensitive.

---

5. Click **OK**.

   The filter is enabled and the window is refreshed to show the filtered information.

Filtering events by service component

Use the following procedure to filter events by service component.

1. Open the **Switch Events** tab as described in "Displaying Switch Events" on page 54.

2. Click **Filter**.

   The **Event Filter** dialog box displays.

3. Select **Service**. The event service menu is enabled.

4. Select either **Switch** or **Chassis** from the menu to show only those messages from the logical
   switch or from the chassis.

5. Click **OK**.

   The filter is enabled and the window is refreshed to show the filtered information.

Displaying the Name Server entries

Web Tools displays Name Server entries listed in the Simple Name Server database. This includes
all Name Server entries for the fabric, not only those related to the local domain. Each row in the
table represents a different device. You can click the column head to sort the events by a particular
column, and drag the column divider to resize a column. You can also right-click a column heading
to resize one or all columns, sort the information in ascending or descending order, or select which
columns are displayed.

Admin Domain considerations: The **Name Server** table is filtered based on Admin Domain
membership of the fabric devices. The **Name Server** table lists only devices that are part of your
current Admin Domain. This includes devices that are direct members of the Admin Domain and
devices that are attached to ports that are direct members of the Admin Domain. All other fabric
devices are filtered out of the **Name Server** view for the current Admin Domain. Refer to "Admin
Domain membership" on page 71 for information about direct and indirect members.

For FICON devices: The **Name Server** table lists the request node identification (RNID) information.

Use the following procedure to display the Name Servers.

1. Select **Tasks > Monitor > Name Server**.

   The **Name Server** window displays.

2. To set an autorefresh rate for the **Name Server** entries, select the **Auto Refresh** check box in
   the **Name Server** window, and enter an auto-refresh interval (in seconds).

   The minimum (and default) interval is 15 seconds.

Printing the Name Server entries

Use the following procedure to set up printing preferences.

1. Select **Tasks > Monitor > Name Server**.

    The **Name Server** window displays.

2. Click **Print**.

3. On the **Page Setup** dialog box, set up your printing preferences and click **OK**.

    The **Print** dialog box displays.

4. Select a printer and click **OK**.

Displaying Name Server information for a particular device

Use the following procedure to display Name Server information for a particular device.

1. Select **Tasks > Monitor > Name Server**.

    The **Name Server** window displays.

2. Select a device from the **Domain** column.

3. Click **Detail View**.

    The **Name Server Information** dialog box displays the information specific to that device.

Displaying zone members for a particular device

Use the following procedure to display zone members for a particular device.

1. Select **Tasks > Monitor > Name Server**.

    The **Name Server** window displays.

2. Select a device from the **Domain** column.

3. Click **Accessible Devices**.

    The **Zone Accessible Devices** window displays accessible zone member information specific to that device.

Physically locating a switch using beaconing

Use the **Beacon** button to physically locate a switch in a fabric. The beaconing function helps to physically locate a switch by sending a signal to the specified switch, resulting in an LED light pattern that cycles through all ports for each switch (from left to right).

NOTE
You must have an RBAC role of admin to initiate switch beaconing. The LED light pattern is initiated on the actual switch or chassis. It is not mirrored in the **Switch View**.

Use the following procedure to use beaconing.

1. Select a logical switch using the drop-down list under **Fabric Tree** section in the Switch Explorer window.

   The selected switch displays in the Switch View.

2. Select **Beacon** for a switch, or **Chassis Beacon** for a chassis-based switch.

   The LED lights on the actual switch light up on the physical switch in a pattern running back and forth across the switch itself. In chassis-based switches, the LEDs glow across all the blades.

3. Look at the physical switches in your installation location to identify the switch.

Locating logical switches using chassis beaconing

Use the following procedure to locate all logical switches in a chassis.

1. Select a logical switch using the drop-down list under the **Fabric Tree** section in the **Switch Explorer** window.

   The selected switch displays in the Switch View.

2. Click **Chassis Beacon**.

   The LEDs on the logical switch light up on the blades associated with the logical switch.

Virtual Fabrics overview

Virtual Fabrics is an architecture that virtualizes hardware boundaries. Traditionally, SAN design and management is done at the granularity of a physical switch. Each switch and all the ports in the switch act as a single fabric element that participates in a single fabric. Virtual Fabrics allows SAN design and management to be done at the granularity of a port. This enables partitioning of a physical switch into multiple logical switches, which may be organized into logical fabrics.

The following platforms are Virtual Fabrics-capable:

- Brocade DCX and DCX-4S
- Brocade 5300
- Brocade 5100

Virtual Fabrics cannot be configured or managed from Web Tools. Configuration and management is done from either the Data Center Management Program (DCFM), or the Fabric OS command line interface. For information about configuring and managing Virtual Fabrics, refer to the *DCFM User Manual* if you are using DCFM, or *Fabric OS Administrator's Guide* if you are using the Fabric OS command line interface.

You can use Web Tools to view Virtual Fabrics and logical switch configurations.

Selecting a logical switch from the Switch View

You can log in to a specific logical switch, as described in Chapter 1, or you can select a logical switch from the Switch View. If you do not log in to a specific logical switch, you are presented with the default logical switch.

Under the **Switch Information** tab, **Base Switch**, **Default Switch**, and **Allow XISL Use** are specific to VIrtual Fabrics. These options perform these functions:

- **Base Switch** indicates whether or not the logical switch can act as a base switch. A base switch is a special logical switch that can be used for chassis interconnection. Each chassis may only designate only one logical switch as a base switch.

- **Default Switch** indicates whether or not the logical switch is the default logical switch. The default logical switch is equivalent to the normal, discovered physical switch topology. It is automatically assigned fabric ID 128. If you do not log in to a specific logical switch using **Options** on the login dialog box, the default logical switch displays in the **Switch View**.

- **Allow XISL Use** indicates whether or not the logical switch is allowed to connect to other logical switches using an extended inter-switch link (XISL). Base switches may use XISLs. Dynamically created logical switches can use the XISL for traffic, only if **Allow XISL Use** is enabled through CLI using the configure command.

Use the following procedure to select a logical switch.

1. Use the **Logical Switch** selector to select the fabric ID, as shown in .

   You must have the EGM license installed to view the Logical Switch selection. A dialog box displays asking you to confirm your selection.
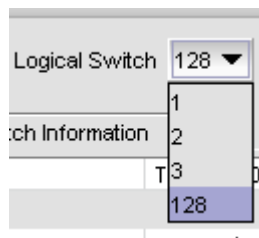
FIGURE 19   Logical switch selector

2. Click **Yes** to confirm.

   The selected logical switch displays. Note that the **Logical Switch** selector is relocated above the **Switch View**.

Viewing logical ports

When base switches are connected through XISLs, a base fabric is formed that includes logical switches in different chassis. A logical link is formed dynamically among logical switches that have the same FID to carry frames between the logical switches. Logical ports are created in the respective switches to support the logical link.

Logical ports are software constructs, and have no corresponding hardware to represent them on the Switch View. Logical port information is available in the **Port Administration** window.

1. Select **Port Administration**.

   The **Port Administration** window displays. Logical ports are displayed in the **FC Ports Explorer** tree structure.

2. To view logical port properties, expand the **Logical Ports** folder, and select a port.

   The **General** properties are displayed.

Maintaining Configurations and Firmware

In this chapter
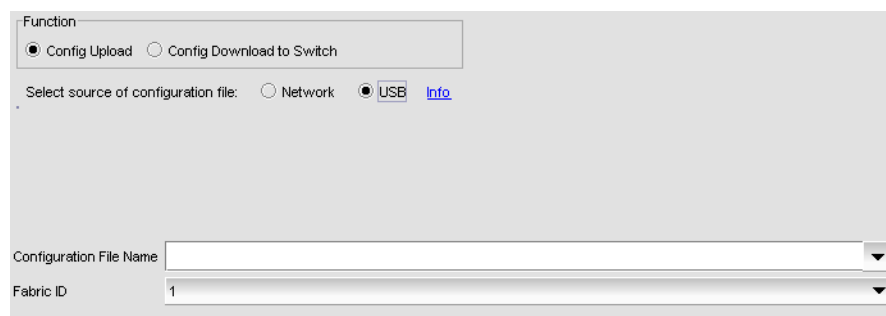
Creating a configuration backup file

Keep a backup copy of the configuration file in case the configuration is lost or unintentional changes are made. You should keep individual backup files for all switches in the fabric. You should avoid copying configurations from one switch to another.

If you upload from a network, enter the host name or IP address in the **Host Name or IP** field, the user ID and password required for access to the host in the **User Name** and **Password** fields, and select the **Protocol Type** used for the upload. The default is FTP. If you select "Secure Copy Protocol (SCP)," you cannot specify "anonymous" in the **User Name** field.

If you select USB as the configuration file source, the network parameters are not needed and are not displayed (Figure 20). You can skip to step 6.

FIGURE 20   USB configuration upload

An <u>info</u> link is enabled when USB is chosen as the source of the configuration file. If you click on <u>info</u>, the following information message displays (Figure 21).
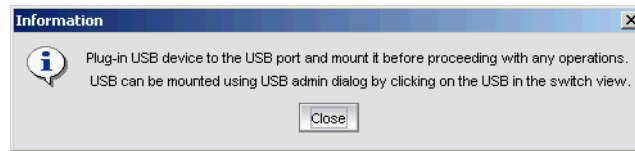
FIGURE 21   Information dialog box

To create a configuration backup file, perform the following task.

1. Open the **Switch Administration** window.

2. Select **Show Advanced Mode**.

3. Select the **Configure** tab.

   The **Configure** screen displays.

4. Select the **Upload/Download** tab.

   The Upload/Download configuration screen displays (Figure 20).

   NOTE
   By default, **Config Upload** is chosen under **Function**, and **Network** is chosen as the source of the configuration file.

5. Enter the configuration file with a fully-qualified path, or select the configuration file name in the **Configuration File Name** field.

6. Use the **Fabric ID** selector to select the fabric ID of the logical switch from which the configuration file is to uploaded.

   The selector displays all the virtual fabric IDs that have been defined, the default of 128 for the physical switch, chassis level configuration, and all chassis and switches.
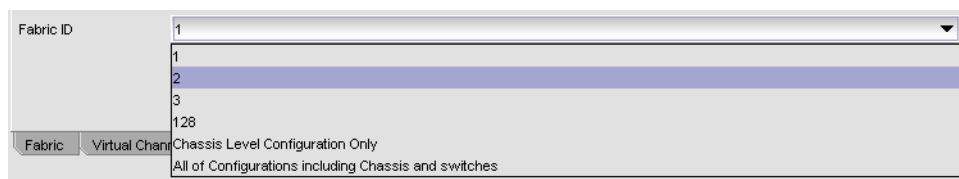
FIGURE 22   Fabric ID selector

NOTE
If you are using a USB device, it must be connected and mounted before you upload or download. Refer to "Uploading and downloading from USB storage" on page 65 for more information.

7. Click **Apply**.

   You can monitor the progress by watching the **Upload/Download Progress** bar.

Restoring a configuration

Restoring a configuration involves overwriting the configuration on the switch by downloading a previously saved backup configuration file. Perform this procedure during a planned down time.

Make sure that the configuration file you are downloading is compatible with your switch model. Configuration files from other model switches might cause your switch to fail.

If you download from a network, enter the host name or IP address in the **Host Name or IP** field, the user ID and password required for access to the host in the **User Name** and **Password** fields, and select the **Protocol Type** used for the upload. The default is FTP. If you select "Secure Copy Protocol (SCP)," you cannot specify "anonymous" in the **User Name** field.

To restore a configuration, perform the following task.

1. Open the **Switch Administration** window.

2. Select **Show Advanced Mode**.

3. Select the **Configure** tab.

   The **Configure** screen displays.

4. Select the **Upload/Download** tab.

   The Upload/Download configuration screen displays. By default, **Config Upload** is chosen under **Function**, and **Network** is chosen as the source of the configuration file.

5. Under **Function**, select **Config Download to Switch**.

6. If you select USB as the configuration file source, the network parameters are not needed and are not displayed (Figure 23), and you can skip to step 7.

FIGURE 23   **USB configuration download**

An **info** link is enabled when USB is chosen as the source of the configuration file. If you click **info**, the following information message displays (Figure 24).
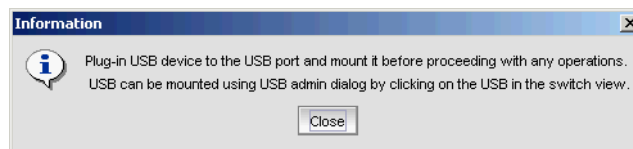
FIGURE 24   **USB information dialog box**

7. Enter the configuration file with a fully-qualified path, or select the configuration file in the **Configuration File Name** field.

8. Use the **Fabric ID** selector to select the fabric ID of the logical switch to which the configuration file is to downloaded.

   The selector displays all the virtual fabric IDs that have been defined, the default of 128 for the physical switch, chassis level configuration, and all chassis and switches.

9. Enter the fabric ID of the logical switch in **Template Fabric ID**.

NOTE
If you are using a USB device, it must be connected and mounted before you upload or download. Refer to "Uploading and downloading from USB storage" on page 65 for more information.

---

10. Click **Apply**.

    You can monitor the progress by watching the **Upload/Download Progress** bar.

Admin Domain configuration maintenance

When you log in to the switch as a physical fabric administrator and back up a configuration, all local switch configuration parameters are saved, as well as all Admin Domain membership information and Admin Domain zone databases.

To perform a configuration upload or download, you should have the Admin Domain of AD255 or AD0, if no other user-defined Admin Domains exist. A configuration upload or download gathers all the configuration files for the fabric, including Admin Domains and iSCSI Target Gateway information. For more information on Admin Domains, refer to "Requirements for Admin Domains" on page 69.

When the configuration is backed up, one of the following scenarios is possible:

- If the current Admin Domain does not own the switch and you are logged in with any role that allows configuration upload or download, the following items are saved in the configuration file:
  - Local zone configuration
  - iSCSI configuration (if any)
  - No other configuration information
- If the current Admin Domain owns the switch and you are logged in with any role that allows configuration upload or download, the following items are saved in the configuration file:
  - Local zone configuration
  - iSCSI configuration (if any)
  - All other configuration information except Admin Domain configuration information
- If you invoke Admin Domain from AD255 and you are logged in with any role that allows configuration upload/download, the following items are saved in the configuration file:
  - Configuration information for zones in all Admin Domains
  - iSCSI configuration (if any)
  - All other configuration information, including zoning from all Admin Domains

The filtering depends on the Admin Domain switch ownership, with additional access if you are in AD255. Access to the command itself is limited by Role-Based Access (RBAC), and not by whether the current user is a Physical Fabric Administrator or an admin user with enumerated access to the relevant domains.

NOTE
The ability to change Admin Domain context requires installing the EGM license. Refer to "Changing the Admin Domain context" on page 23 for complete instructions.

The EGM license is required only for 8 Gbps platforms, such as the Brocade DCX and DCX-4S enterprise-class platforms, the Encryption Switch, the 300, 5300, and 5100 switches. For non-8 Gbps platforms, all functionalities are available without the EGM license.

Uploading and downloading from USB storage

If you choose to upload or download from a USB device, you must click the USB port to launch the USB Port Management wizard (Figure 25).

Perform the following procedure to update your USB storage.

1. Select **Mount USB Device,** and select **Yes** at the confirmation prompt.

2. Right-click on a configuration file to access **Export**, **Copy**, and **Search** options (Figure 25).
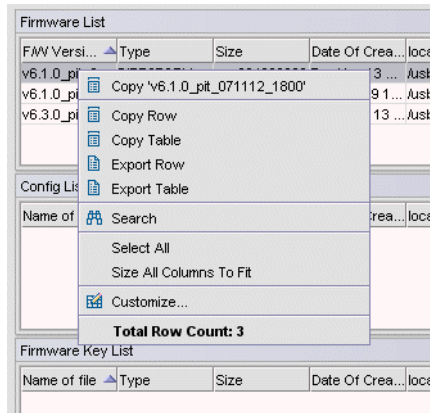
FIGURE 25   USB Port Management wizard—right-click options

3. Click **Copy** to upload and **Export** to download.

Performing a firmware download

During a firmware download, the switch restarts and the browser temporarily loses connection with the switch. When the connection is restored, the version of the software running in the browser is different from the new software version that was installed and activated on the switch. You must close all of the Web Tools windows and log in again to avoid a firmware version mismatch. Note that for chassis-based switches, you might get popup messages that imply the loss of connection is temporary and will soon be resolved. You must still close all windows and re-log in.

When you request a firmware download, the system first checks the file size being downloaded. If the compact flash does not have enough space, Web Tools displays a message and the download does not occur. If this happens, contact your switch support supplier.

NOTE
You can perform a firmware download only when the current Admin Domain owns the switch.

Perform the following procedure to download a new firmware version.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Firmware Download** tab.

3. Choose to download either the firmware or the firmware key.

   The download source can be located on the network or a USB device.

   NOTE
   When you select the **USB** button, you can specify only a firmware path or directory name. No other fields on the tab are available. The **USB** button is available if the USB is present on the switch.

4. Enter the host name or IP address, user name, password, and fully-qualified path to the file *release.plist*.

   You can enter the IP address in either IPv4 or IPv6 format.

   The path name should use the following structure:

   //<directory>/<fos_version_directory>/release.plist

   where the *<directory>* is the path up to the entry point of *<fos_version_directory>* and *<fos_version_directory>* is where the unzipped version of Fabric OS is located.

   Example
   //directory_1/my_directory/v6.4.0/release.plist

   NOTE
   You cannot downgrade to an earlier version of Fabric OS using Web Tools. You must use the Brocade Fabric OS CLI.

5. Select the protocol type in the **Protocol Type** field.

   If you select "Secure Copy Protocol (SCP)," you cannot specify "anonymous" in the **User** field.

6. Click **Apply.**

   The firmware download begins. You can monitor the progress by looking at the Firmware Download progress bar.

   NOTE
   About halfway through the download process, after the firmware key is downloaded to the switch, connection to the switch is lost and Web Tools invalidates the current session. Web Tools invalidates all windows because upfront login is always enabled and cannot be disabled.

7. Close all Web Tools windows and log in again.

    If the firmware download is in progress when you log in, you can continue to monitor its progress.

SAS and SA firmware download

If you are downloading SAS and DMM firmware directly to the blade, you have more options on the **Firmware Download** tab. Also, for the Brocade 7600 a collapsible area displays on the **Firmware Download** tab to show application firmware information.

NOTE
SAS is supported only on Brocade 7600 and FA4-16.

In addition to specifying the information described in the steps in "Performing a firmware download" on page 65, you can select from the following:

- Whether to download the firmware or the firmware key.
- The source of the firmware, network or USB.
- The type of firmware you want to be downloaded.
- The firmware path. Web Tools displays up to 15 path entries in the **Specify Firmware Path** field.
- To download the firmware even if it is not compatible (skip version check).
- To enable removal of application firmware (erase SA).
- The blade to be upgraded (by slot).

Firmware download is non-disruptive, from the Web Tools perspective.

Switch configurations for mixed fabrics

You can use Web Tools to configure switches in a mixed fabric. You do this by setting the switch to interoperability mode, which is either McDATA Open Fabric mode or McDATA Fabric mode. When you turn on interoperability mode, the Zone database is cleared.

When you turn on McDATA Fabric Mode, which supports M-EOS switches v9.6.2 and later, note the following:

- Switch View displays McDATA Fabric in the **Interoperability Mode** field.
- For configuration download:
    - When downloading the configuration file, the download process checks the compatibility, such as Admin Domain, SCC/DCC policy, domain ID and zoning and issues error messages if there are incompatible issues, such as the switch settings.
- Domain ID is calculated based on the offset selected.

When you turn on McDATA Open Fabric Mode, which supports M-EOS switches v9.6.2 and later, the following features are not available:

- Admin Domain in the **Manage** section of the **Tasks** menu.
- The FICON CUP feature.
- Zone, alias, and configuration operations.

Enabling interoperability

When you configure interoperability, Web Tools verifies that the domain ID of the switch falls within the range for the interoperability mode you select. The domain ranges are:

- The normal domain ID range is 1 through 239.
- The McDATA Fabric mode and McDATA Open Fabric mode are calculated based on the offset selected.
- Before making changes to the interoperability mode, Web Tools checks to see if the domain ID of the switch falls within the range. If not, a warning message prompts you to change the domain ID before enabling the interoperability mode.

Perform the following steps to change the interoperability configuration.

1. Open the **Switch Administration** window.

2. Disable the switch, as described in "Enabling and disabling a switch" on page 40.

   NOTE
   You can change the interoperability configurations only to a disabled (offline) switch. You are only able to disable the switch if the Admin Domain you are logged into owns the switch.

3. Select the **Configure** tab.

4. Select the **Interoperability** subtab.

5. Select the required mode:

   - Brocade Native Fabric Mode
   - For McDATA Fabric Mode and McDATA Open Fabric Mode, you can select the offset value from the **Domain Offset (Hex) / Domain ID (Hex)** drop-down list.

     The available offset values are 0 (0x00), 32 (0x20), 64 (0x40), 96 (0x60), 128 (0x80), 160 (0xA0), and 192 (0xC0). McDATA Open Fabric Mode supports 239 (0xFF) also.

6. Click **Apply.**

Managing Administrative Domains

In this chapter

Administrative Domain overview

Using Administrative Domains (Admin Domains or ADs), you can partition the fabric into logical groups and allocate administration of these groups to different user accounts so that these accounts manage only the Admin Domains assigned to them and do not make changes to the rest of the fabric. The ability to assign an Admin Domain to a specific user account is performed in the **User** tab of the **Switch Administration** window and not in the **Admin Domain** window.

You can create domains that are grouped together based on the type of members in the domain. For example, you can create Admin Domains based on the type of switches in your fabric using the WWN (not to be confused with the Admin Domain number) or put all the devices in a particular department in the same Admin Domain for ease of administering those devices.

You can have up to 256 Admin Domains in a fabric (254 user-defined and 2 system-defined), numbered from 0 through 255. Admin Domains are designated by a name and a number. This document refers to specific Admin Domains using the format "AD*n*" where *n* is a number between 0 and 255.

NOTE
ADs and Virtual Fabrics are mutually exclusive. Virtual Fabrics must be disabled if you want to use the AD feature.

Requirements for Admin Domains

The following are the requirements for using administrative domains:

• Admin Domains are supported on fabrics with switches running Fabric OS v5.2.0 or later.

• To manage Admin Domains, you must be a physical fabric administrator. A physical fabric administrator is a user with the Admin role and access to all Admin Domains (AD0 through AD255).

• The default zone mode setting must be set to No Access (refer to "Enabling Admin Domains" on page 71).

User-defined Admin Domains

AD1 through AD254 are user-defined Admin Domains. These user-defined Admin Domains can be created only by a physical fabric administrator in AD255.

System-defined Admin Domains

AD0 and AD255 are special Admin Domains and are present in every AD-capable fabric.

AD0

AD0 is a system-defined Admin Domain that, in addition to containing members you explicitly added (similar to user-defined Admin Domains), it contains all online devices, switches, and switch ports that were not assigned to any user-defined Admin Domain.

AD0 also implicitly contains all devices from switches running Fabric OS versions earlier than 5.2.0, as they can never be part of an Admin Domain unless they are upgraded to v5.2.0 or later.

Unlike user-defined Admin Domains, AD0 has both an automatic membership list and a fixed membership list. User-defined Admin Domains have only a fixed membership list.

- Automatic membership list—Contains all devices and switches that were not assigned to any other Admin Domain.
- Fixed membership list—Contains all devices and switches that you explicitly add to AD0 and can be used to force device and switch sharing between AD0 and other Admin Domains.

The Admin Domain window displays the fixed members and not the automatic members, you can use the **View** menu to display a list of the automatic members.

AD0 can be managed like any user-defined Admin Domain. The only difference between AD0 and user-defined Admin Domains is the automatic membership list.

In filtered views, the automatic members of AD0 are considered direct members.

The automatic members of AD0 change dynamically as the membership of other Admin Domains changes. The fixed members of AD0 are not deleted unless you explicitly remove them.

For example, if you explicitly add DeviceA to AD0 and it is not a member of any other Admin Domain, then DeviceA is both an automatic and a fixed member of AD0. If you add DeviceA to AD2, then DeviceA is deleted from the AD0 automatic membership list, but is *not* deleted from the AD0 fixed membership list. If you then remove DeviceA from AD2, DeviceA is added back to the AD0 automatic membership list (assuming DeviceA is not in any other Admin Domains).

AD0 is useful if you want to share its zone database (called "root zone database") with a legacy fabric.

AD255 or physical fabric

AD255 is a virtual domain that contains all devices, switches, and switch ports in the fabric. AD255 presents an unfiltered view of the fabric and is also referred to as the physical fabric.

You can use AD255 to do the following:

- Manage other Admin Domains.

- Get an unfiltered view of the fabric.

- Manage ACL and distribution (this can be managed in AD0 if no other Admin Domains are present).

- Manage Advanced Performance Monitoring (this can be managed in AD0 if no other Admin Domains are present, but only if you are using Web Tools with the EGM license).

The EGM license is required only for 8 Gbps platforms, such as the following:

- Brocade DCX and DCX-4S enterprise-class platforms

- Encryption Switch

- 300, 5300, and 5100 switches

For non-8 Gbps platforms, all functionalities are available without EGM license.

You cannot manage zones with AD255, because AD255 does not have a zone database associated with it.

Admin Domain membership

Switches, ports, and devices can be members of an Admin Domain. Admin Domain members can be either direct or indirect members as described below:

- Direct members—Devices, switches, and ports that you explicitly add to an Admin Domain. Direct members are listed in the Admin Domain membership list.

- Indirect port members—Ports that are implicitly added as part of an Admin Domain when any of the following occurs:

    - A device that is connected to a port was added to the Admin Domain.

    - A switch to which the port belongs is a member of the Admin Domain.

- Indirect device members—Devices that are connected to ports that are direct members of an Admin Domain.

Enabling Admin Domains

The default zone mode setting gives attached devices either All Access to all devices or No Access to all devices. To begin implementing an Admin Domain structure within a SAN, you must set the default zone mode to No Access. You must be in AD0 to change the default zone mode. After the default zone mode is set to No Access, you cannot change it from the physical fabric.

NOTE
The term "physical fabric" is used in Web Tools only.

Even though the default zone mode access is set to No Access, you can still create and enable zones within each Admin Domain. These zones are configurable only from the Admin Domain in which they were created. Indirect port members cannot be zoned.

Perform the following steps to enable Admin Domains.

1. Change the Admin Domain context to AD0. Refer to "Changing the Admin Domain context" on page 23.

   > **NOTE**
   > Changing the Admin Domain context requires using Web Tools with the EGM license; otherwise, access to this feature is denied and an error message displays. Change the Default Zone mode to No Access. Refer to "Setting the default zoning mode" on page 121 for more information.

2. Navigate to AD255 or the physical fabric and begin managing the Admin Domains.

Admin Domain window

You can view and manage Admin Domains through the Admin Domain window, shown in Figure 26.

If you are not using Web Tools with the EGM licensed installed, managing Admin Domain operations is denied and an error message displays.

The EGM license is required only for 8 Gbps platforms, such as the following:

- Brocade DCX and DCX-4S enterprise-class platforms
- Encryption Switch
- 300, 5300, and 5100 switches

For non-8 Gbps platforms, all functionalities are available without EGM license.
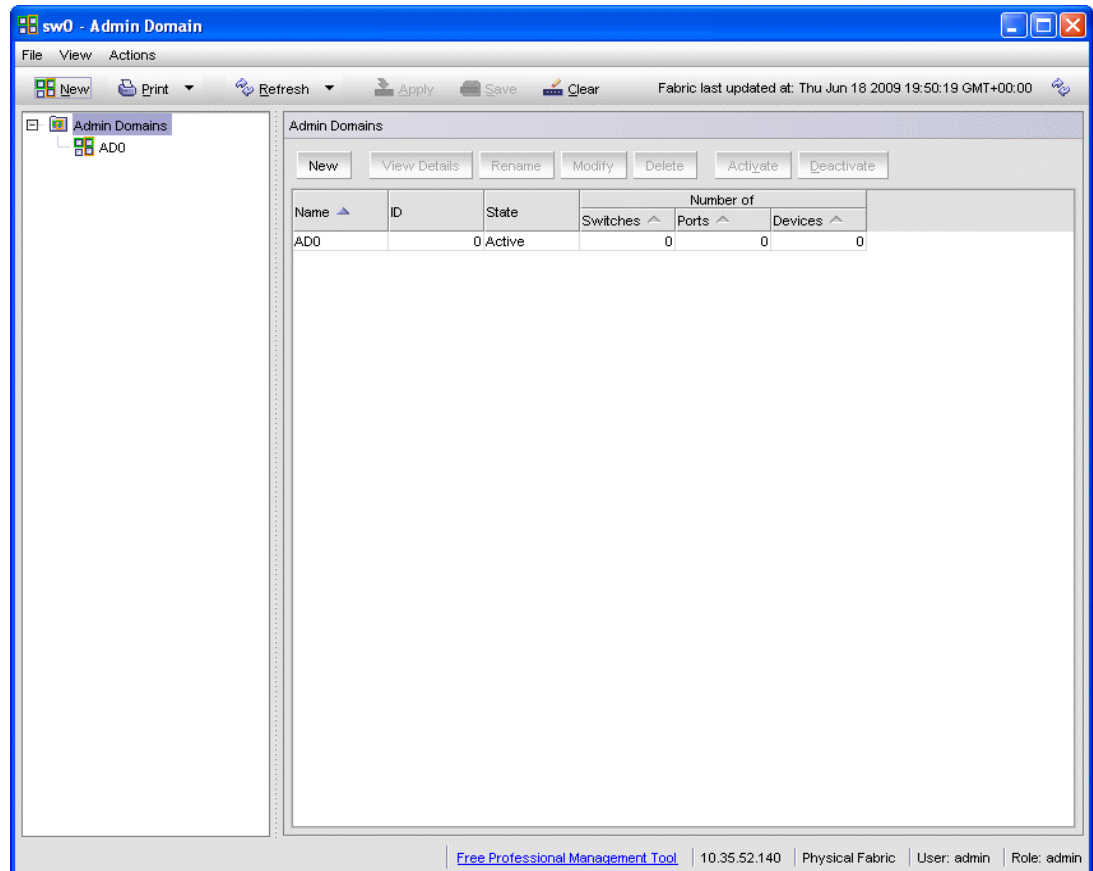
FIGURE 26   Admin Domain window, summary view

The **Admin Domain** window displays information about the Admin Domains that are defined in the fabric. If you launch the **Admin Domain** window from AD255 (physical fabric), the window contains information about the current content of all Admin Domains. If you launch the **Admin Domain** window from any other Admin Domain, the window displays the current Admin Domain only.

To manage Admin Domains, you must be logged in with the role of Admin.

ATTENTION
Any changes you make in the **Admin Domain** window are held in a buffered environment and are *not saved to persistent storage until you explicitly save the changes*. If you close the **Admin Domain** window without saving your changes, your changes are lost.

To save the buffered changes you make to persistent storage in the **Admin Domain** window, refer to

When you are logged into AD0, if a physical fabric administrator modifies the AD configuration from another session, the changes in the membership might not be visible to you.

When you launch the **Admin Domain** window and select the parent **Admin Domains** node in the tree on the left pane, the **Admin Domain** window displays summary information about all of the Admin Domains, as shown in Figure 26. You can also select a specific Admin Domain from the tree to display detailed information about that Admin Domain, as shown in Figure 27. The detailed view displays summary information as well as information about the online switch, port, and device members of the selected Admin Domain.
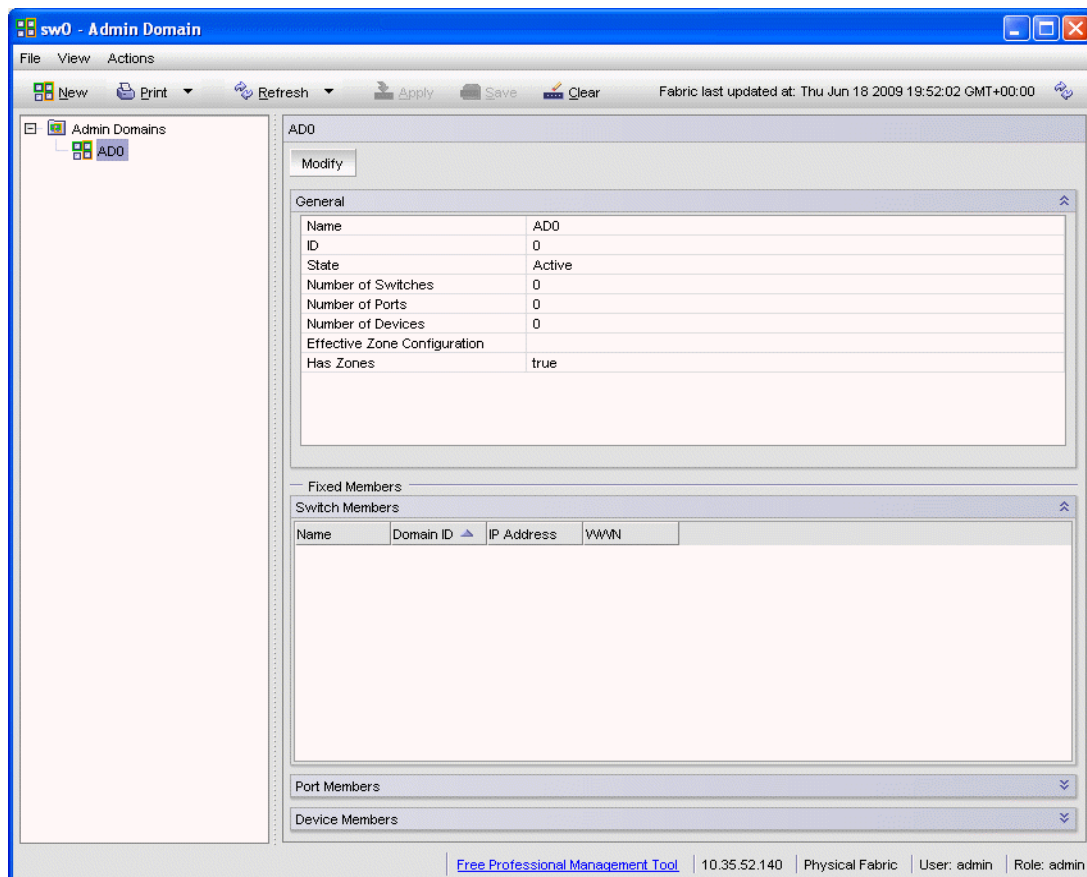
FIGURE 27   Admin Domain window, single Admin Domain detail

NOTE
The tree only displays launched switches and their ports. It also displays all the devices in the fabric. Slot and port information of other switches are not displayed in the tree.

The **Admin Domain** window has the following buttons in a task bar at the top of the window:

- **New** allows you to create a new Admin Domain.
- **Print** allows you to print the current or effective configuration.
- **Refresh** allows you to refresh the information for the entire fabric or a specific Admin Domain.
- **Apply** allows you to apply a configuration.
- **Save** allows you to save a configuration.
- **Clear** allows you to clear the configuration.

You can right-click on any of the table content in the **Admin Domain** window to access **Export**, **Copy**, and **Search** options. The options are not available if the table does not have any content.

NOTE
You must accept the Brocade Certificate at the beginning of the log in to Web Tools to enable the functionality of **Export** and **Copy**.

- Click **Export Row** or **Export Table** to save the contents to a tab-delimited file.
- Click **Copy Row** or **Copy Table** to copy the contents in tab-delimited text format to a file.
- Click **Search** to search for a specific text string in the table.

    The **Switch Members** dialog box displays, as shown in Figure 28.

    In the **Switch Members** dialog box, enter the text string and press **Enter**. This is an incremental search and allows 24 maximum characters including the wildcards question mark (?) and asterisk (*). The first row containing the text string is highlighted. To find the next match, press the down arrow. To find the previous match, press the up arrow. If the text is not found in the table, the text turns red.
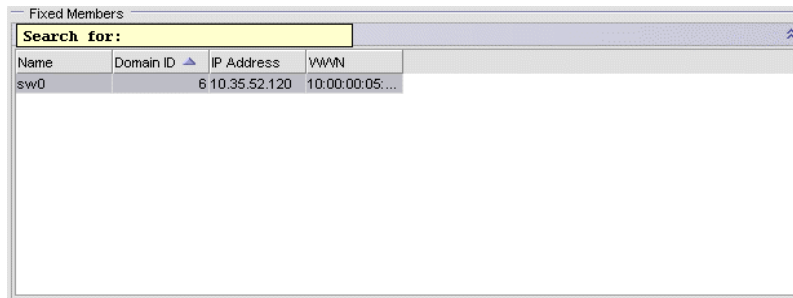
FIGURE 28   Searching for a text string in a table

Opening the Admin Domain window

Use the **Admin Domain** window to perform all Admin Domain configuration procedures.

If you want to configure Admin Domains, you must launch the **Admin Domain** window from the physical fabric context. If you are in any Admin Domain other than the physical fabric, the module launches in read-only mode.

NOTE
The switch must be running Fabric OS v5.2.0 or later.

Perform the following procedure to open an **Admin Domain** window.

1. Select a switch from the Fabric Tree and log in when prompted.

    Switch View displays information for the selected switch.

2. If you plan to modify the Admin Domain configuration, from the **Admin Domain** menu, select **Physical Fabric**.

3. Click **Admin Domain** in the **Manage** section of the **Tasks** menu.

    The **Admin Domain** window displays (Figure 26).

Refreshing fabric information

When you refresh, the system updates the display of fabric elements only (switches, ports, and devices). It does not update Admin Domain changes in the **Admin Domain** window.

This option allows you to refresh the fabric element information displayed at any time.

To refresh the fabric information open the **Admin Domain** window and click **Refresh**. The status for the fabric, including switches, ports, and devices is refreshed.

Refreshing Admin Domain information

Any changes you make in the **Admin Domain** window are saved to a local buffer. They are not applied to persistent storage until you invoke one of the transactional operations listed in the **Actions** menu.

You can refresh the Admin Domain information at any time to reflect changes that might have been made by other users or to back out of current, unsaved work and start again.

ATTENTION
When you refresh the buffered information in the **Admin Domain** window, any Admin Domain configuration changes you made and not yet saved are erased from the buffer and replaced with the currently enabled Admin Domain information that is saved on the switch.

Use the following procedure to update the information in the **Admin Domain** window with the information saved on the switch.

1. In the **Admin Domain** window, click the **Refresh** arrow.

2. Click **Refresh Admin Domains**.

    The information in the **Admin Domain** window is updated with the saved information on the switch. This action also refreshes the fabric information as described in "Refreshing fabric information" on page 76. Any unsaved Admin Domain changes are deleted.

Saving local Admin Domain changes

All information displayed and all changes made in the **Admin Domain** window are buffered until you save the changes. That means that any other user looking at the Admin Domain information for the switch does not see the changes you made until you save them.

Use the following procedure to save the local Admin Domain changes.

1. Select **Actions** > **Save AD Configuration** to save your changes to persistent storage as the defined Admin Domain configuration.

2. Select **Actions** > **Apply AD Configuration** to save your changes to persistent storage *and make your changes effective in the fabric*.

These options are not enabled until you make a change to the Admin Domain configuration.

If another user has an Admin Domain operation in progress at the time that you attempt to save changes, Web Tools displays a warning to indicate that another Admin Domain transaction is in progress on the fabric. You can select to abort the other transaction and override it with yours.

This action updates the entire contents of the **Admin Domain** window, not just the selected Admin Domain. You can save your changes at any time during the Admin Domain configuration session.

Closing the Admin Domain window

It is important to remember that any changes you make in the **Admin Domain** window are not saved automatically.

Use the following procedure to close the Admin Domain window.

1. In the **Admin Domain** window, select **File** > **Close**.

   If there are changes in the buffer that were not saved, a warning message displays. Confirm that you want to close the Admin Domain session without saving the changes.

2. Click **Yes** to close without saving changes, or click **No** to go back to the **Admin Domain** window to save the changes (refer to *"Saving local Admin Domain changes"* on page 76).

Creating and populating domains

Setting up an Admin Domain involves the following steps.

1. Creating an Admin Domain.

2. Assigning one or more administrators to the Admin Domain.

   The Admin account always has access to administer the Admin Domains, even if no other users are assigned (refer to *"Changing user account parameters"* on page 212).

When you create an Admin Domain, you can activate the Admin Domain after you finish creating it. If you activate the Admin Domain, you must click **Apply** to transfer your changes from the Web Tools database to the fabric database so that your changes are applied to the fabric. You can log in to an active Admin Domain. You cannot log in to an Admin Domain that was deactivated.

Creating an Admin Domain

Use the following procedure to create an Admin Domain.

1. Open the **Admin Domain** window, as described in *"Opening the Admin Domain window"* on page 75.

2. Click **New**.

   The **Create Admin Domain** wizard displays.

3. In the **Name** area, assign an Admin Domain name.

   You can specify a name or let the system assign the name for you.

4. In the **ID** area, assign an Admin Domain ID.

   You can specify an ID or let the system assign the ID for you.

5. In the **State** area, select the **Active** check box to activate the Admin Domain when you finish creating it.

   NOTE
   Clear the **Active** check box if you want the Admin Domain deactivated when you finish creating it.

6. Click **Next**.

7. In the **Membership** area, assign members to the Admin Domain by selecting them in the Available Members section and clicking **Add**, **Add Ports**, or **Add Devices** as described below:

    - Select a switch, port, or device in the **Available Members** tree and click **Add** to add the selected element.

        Alternatively, you can press the **Insert** key to add your selections.

    - Select a switch or slot and click **Add Ports** to add all of the ports in the selected switch or slot.

    - Select a switch, slot, or port and click **Add Devices** to add all of the devices for the selected element.

8. *Optional*: Click **Manual** to add offline devices.

    NOTE
    To add ports or other switches in the fabric, launch the **Add Member** wizard by clicking the **Manual** button.

    NOTE
    The tree only displays two FC ports and all logical ports if the 7500E upgrade license is not installed.

    ---

9. Click **Next**.

    The wizard displays a summary of the Admin Domain. Read the summary to verify that the Admin Domain setup is correctly.

10. Click **Finish** to close the wizard.

11. Click **Save** to save the new Admin Domain configuration to persistent storage.

12. Click **Apply** to enforce the new Admin Domain configuration as the effective configuration.

Adding ports or switches to the fabric

Use the following procedure to add ports or switches to the fabric.

1. From the **Create Admin Domain** wizard, click **Manual**.

    The **Add Member** window displays.

2. Select **Port** and enter the member ID in the **Member** field using the Domain Index (D,I) format.

3. Click **Apply** to enforce the added members, and then click **OK** to accept the changes.

Activating or deactivating an Admin Domain

Use the following procedure to activate or deactivate an Admin Domain.

1. Open the **Admin Domain** window.

2. From the tree on the left, select the Admin Domain you want to activate or deactivate.

3. Click **Activate** to activate the Admin Domain, or click **Deactivate** to deactivate the Admin Domain.

4. Select **Actions** > **Save AD Configuration** to save the new Admin Domain configuration to persistent storage.

5. Select **Actions** > **Apply AD Configuration** to enforce the new Admin Domain configuration as the effective configuration.

    ATTENTION
    When you deactivate an Admin Domain, the members or devices assigned to the domain can no longer access its hosts or storage unless those devices are part of another Admin Domain.

    When you deactivate an Admin Domain, no one can use this Admin Domain to log in to a switch.

Modifying Admin Domain members

Use the following procedure to modify members from an Admin Domain.

1. Open the **Admin Domain** window.

2. From the tree on the left, select the Admin Domain you want to modify.

3. Click **Modify**.

    The **Modify Admin Domain** wizard displays the Membership step.

4. Assign members to the Admin Domain by selecting them in the **Available Members** section and clicking **Add**, **Add Ports**, or **Add Devices** as described below:

    - Select a switch, port, or device in the **Available Members** tree and click **Add** to add the selected element.

      Alternatively, you can press the **Insert** key to add your selections.

    - Select a switch or slot and click **Add Ports** to add all of the ports in the selected switch or slot.

    - Select a switch, slot, or port, and click **Add Devices** to add all of the devices for the selected element.

5. *Optional*: Click **Manual** to add offline switches and devices.

6. Remove members from the Admin Domain by selecting them in the **Selected Members** section and clicking **Remove**.

    Alternatively, you can press the **Delete** key to remove selected items.

7. Click **Next**. Use the summary to verify that the Admin Domain setup is correct.

8. Click **Finish**.

9. Select **Actions** > **Save AD Configuration** to save the new Admin Domain configuration to persistent storage.

10. Select **Actions** > **Apply AD Configuration** to enforce the new Admin Domain configuration as the effective configuration.

Renaming Admin Domains

You can change the name of an Admin Domain, including an auto-assigned ID name. The Admin Domain name cannot exceed 63 characters and can contain alphabetic and numeric characters. The only special character allowed is an underscore (_).

NOTE
You cannot rename AD0 or AD255.

Use the following procedure to rename an Admin Domain.

1. Open the **Admin Domain** window.

2. From the tree on the left, select the Admin Domain.

3. Click **Rename**.

4. Enter the new name and click **OK**.

5. Select **Actions** > **Save AD Configuration** to save the new Admin Domain configuration to persistent storage.

6. Select **Actions** > **Apply AD Configuration** to enforce the new Admin Domain configuration as the effective configuration.

Deleting Admin Domains

When you delete an Admin Domain, its devices no longer have access to the members of the zones with which it was associated.

Use the following procedure to delete an Admin Domain.

1. Open the **Admin Domain** window.

2. From the tree on the left, select the Admin Domain.

3. Click **Delete**.

4. In the confirmation dialog box, click **Yes** to delete the domain.

    The system deletes the Admin Domain.

5. Select **Actions** > **Save AD Configuration** to save the new Admin Domain configuration to persistent storage.

6. Select **Actions** > **Apply AD Configuration** to enforce the new Admin Domain configuration as the effective configuration.

Clearing the Admin Domain configuration

When you clear the Admin Domain configuration, all user-defined Admin Domains are deleted and all fabric resources (switches, ports, and devices) are returned to AD0. You cannot clear the Admin Domain configuration if zone configurations exist in any of the user-defined Admin Domains.

Use the following procedure to clear the **Admin Domain** configuration.

1. Open the **Admin Domain** window.

2. Select **Actions** > **Clear AD Configuration**.

3. In the confirmation dialog box, click **Yes** to clear the Admin Domain configuration.

Managing Ports

In this chapter

Port management overview

This chapter describes how to manage FC and gigabit Ethernet (GbE) ports. Refer to "Viewing EX_Ports" on page 153 for information on how to view and configure EX_Ports.

The **Port Administration** window is refreshed automatically every sixty seconds and is refreshed immediately when you make any port changes through Web Tools.

To manage ports, you must be logged in with the role of switchadmin, admin, basicswitchadmin, operator, or fabric admin. If you are logged in with a user, securityadmin, or zoneadmin role, you can only view the port information.

Opening the Port Administration window

To open the Port Administration window, click **Port Admin** in the **Switch View** window. The window displays in Basic Mode. Refer to "Switch View" on page 25 for information about accessible ports.

The **Port Administration** window displays information about the ports on the switch. Click **Show Advanced Mode** in the upper-right corner of the window to see more port management options.

Admin Domain considerations

In fabrics with user-defined Admin Domains, the **Port Administration** window is filtered to show only ports that are direct or indirect members of the currently selected Admin Domain:

- Direct members are ports that were directly added to the Admin Domain as members.

- Indirect members are:
  - Non-owned ports on a member switch
  - Non-owned ports to which member devices are attached
- All active ports, as well as any inactive EX_Ports are shown.

Port Administration window components

The **Port Administration** window has the following two tabs in the top left corner:

- **FC Ports** tab displays all of the FC ports on the switch (physical FC ports and logical FCIP ports)
- **GigE Ports** tab displays all of the GbE ports

  If the switch does not have GbE ports, the **GigE Ports** tab does not display.

  The **GigE Ports** tab has the following three subtabs:

  - **Interfaces**—Lets you view interfaces
  - **Routes**—Lets you view routes
  - **FCIP tunnel**—Lets you view tunnels. This tab has two buttons: **Go to FCIP port** and **Show Security Policies.**

    On selecting an FCIP tunnel the following circuit details with the circuit properties are displayed:

    - ❑ Circuit Number
    - ❑ Tunnel ID
    - ❑ Administrator Status
    - ❑ Operational Status
    - ❑ GigEPort
    - ❑ Source IP
    - ❑ Gateway
    - ❑ VLAN ID
    - ❑ MTU Size
    - ❑ Compression Enabled
    - ❑ Data L2COS Value
    - ❑ DSCP Data
    - ❑ IKE Policy Number
    - ❑ IPsec Policy Enabled
    - ❑ Keep Alive Timeout
    - ❑ MaximumCommunicationRate
    - ❑ MinimumCommunicationRate
    - ❑ MaxRetransmitRate
    - ❑ MinRetransmitRate
    - ❑ Metric
    - ❑ Pre-Shared key
    - ❑ QOS Mapping
    - ❑ Selective Ack

- The Ports Explorer tree on the left side. Items in the tree are displayed as follows:
    - Switches—Switch ID, with switch name in parentheses; for example, 3(MapsSW_202)
    - Blades—Slot number of the blade, with blade ID in parentheses; for example, Slot 7(24)
    - Ports—Port number; for example, Port 2
- Button area. The button area contains buttons for all the tasks you can perform on the selected port. If you select more than one port, buttons are available for only the tasks that you can perform on all of the selected ports. Buttons are grayed (unavailable) if they are not applicable to the selected ports.
    - Port information displays in either a table of ports or information about a specific port, depending on your selection.
    - If you select a slot or switch, the system displays a table of all the ports for the slot or switch.
    - If you select a port, the system displays detailed information about the port.

      You can view either **Basic Mode** or **Advanced Mode**, and to view the subtabs that contain additional information about the port. The available subtabs depend on the type of port selected.
- When viewing detailed information about a port, **Basic Mode** provides these subtabs:
    - **General**—All ports
        - ◻ View Details
        - ◻ Rename
        - ◻ Edit Configuration
        - ◻ Enable/Disable (port)
        - ◻ Persistent Enable/Persistent Disable (port)
    - **SFP**—Physical ports only (FC, CEE, and GbE)
        - ◻ Basic information about the port equipment
    - **Port Statistics**—All ports
        - ◻ Basic port information and statistics
        - ◻ Advanced port information

      Note that on the **Port Statistics** subtab, you can view either absolute values or deltas for port statistics. Viewing the deltas is useful if you want to view current port trends. To reset the counters on the port statistics, click the **Clear Counters** button.

      FCIP statistics for a GbE port are the accumulated statistics of all the FCIP tunnels for that GbE port.
    - IP Interfaces—GbE ports only
    - IP Routes—GbE ports only

- When viewing detailed information about a port, the **Advanced Mode** provides these additional subtabs:
  - General—All ports
    - ☐ Enable/Disable Trunking
    - ☐ Enable/Disable NPIV
    - ☐ Port Swap
    - ☐ F_Port Trunking
    - ☐ Re-Authenticate
    - ☐ Bind/Un-Bind PID
    - ☐ F-Port BB Credit
    - ☐ QoS Enable/Disable (These options require Adaptive Networking License)
  - SFP—Physical ports only (FC, CEE, and GbE)
    - ☐ Advanced information about the port equipment
  - Port Statistics
    - ☐ Advanced port statistics
    - ☐ Error details
    - ☐ FCIP Tunnels—GbE ports and logical FCIP ports only (not available for the FR4-16IP)

Controllable ports

All ports have a "Controllable" attribute visible from the **Advanced Mode**, which represents a combination of the RBAC and Admin Domain permissions, as shown in Figure 29 on page 87.

The Controllable attribute is **No** in the following situations:

- If your account has read-only permission, all accessible ports display in read-only mode, regardless of the Admin Domain context. All configuration functionality is disabled.

- Non-owned E_Ports and indirect member ports on non-owned switches are accessible in read-only mode and are not controllable, regardless of RBAC permissions.

The Controllable attribute is **Yes** for ports that are directly owned by the current Admin Domain and for all ports on switches that are owned by the current Admin Domain, if your role gives you Modify permission for ports. If a port is controllable, all configuration functionality is enabled.

Ports on a non-owned switch that are not E_Ports and are neither direct nor indirect members of the current Admin Domain are inaccessible and are not displayed in the **Port Administration** window.
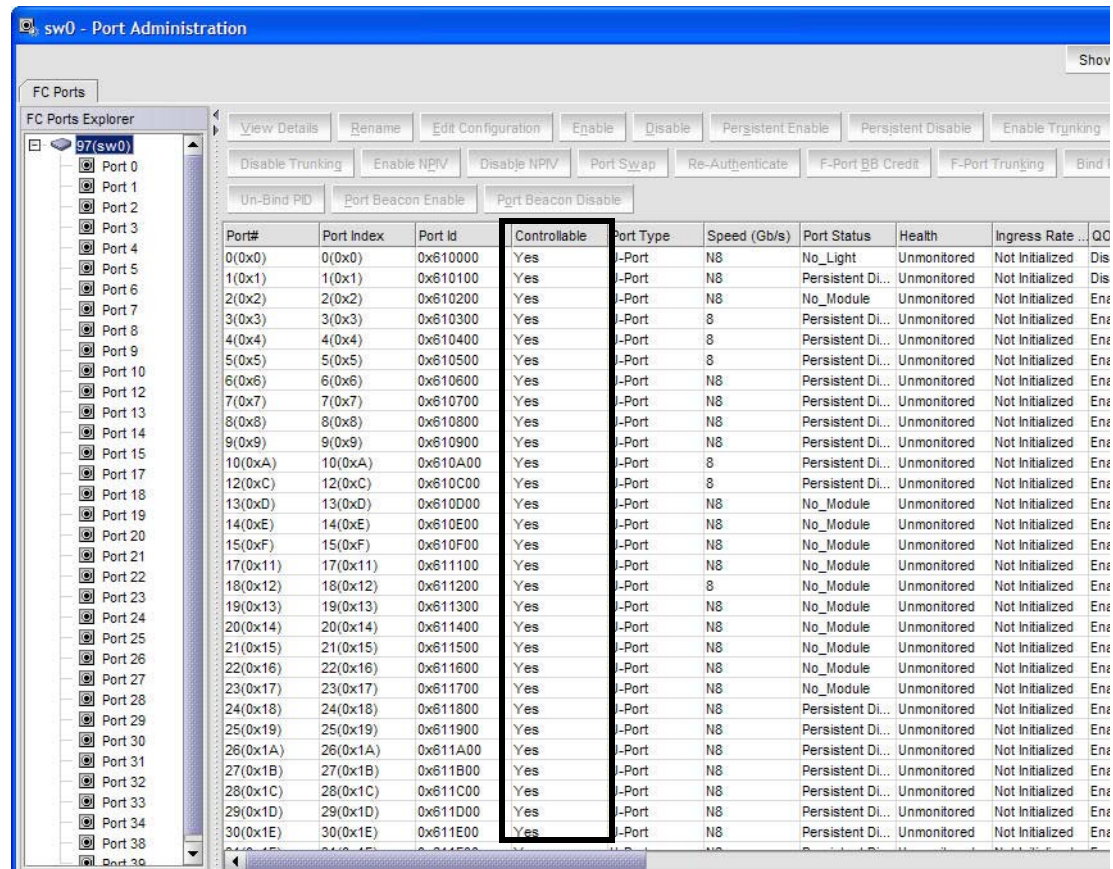
FIGURE 29   Port Administration window, Table view

Configuring FC ports

With the **FC Port Configuration** wizard, you can configure allowed port types, port speed, and long distance mode for physical ports.

You must use Web Tools with the EGM license enabled on the switch to configure long distance; otherwise, access to this feature is denied and an error message displays.

The EGM license is required only for 8 Gbps platforms, such as the following:

- Brocade DCX and DCX-4S enterprise-class platforms
- Encryption Switch
- 300, 5300, and 5100 switches

For non-8 Gbps platforms, all functionality is available without EGM license.

The following procedure describes how to open the **FC Port Configuration** wizard. The wizard is self-explanatory, so the explicit steps are not documented here.

1. Click a port in the **Switch View** to open the **Port Administration** window.

2. Click the **FC Ports** tab.

3. Select the port you want to configure from the tree on the left.

4. Click the **General** subtab.

> **NOTE**
> Long distance does not display from the **General** or **Table** subtabs if the EGM license is not enabled on the switch.

5. Click **Edit Configuration**. The **FC Port Configuration** wizard displays.

   The fields are populated with the current configuration values.

> **NOTE**
> Long distance is not displayed from the **Edit Configuration** window. You can view long distance from the **View** tab when you display the port details.

6. Follow the steps in the wizard.

> **NOTE**
> If you configure a disabled port as an EX_Port, the wizard displays the **Enable Port after configuration** check box. If you select the check box, the disabled port is automatically enabled after configuration; otherwise, the port remains in the same state after configuration.

Allowed port types

For FC ports, the **Port Administration** window displays the following values relating to port type:

Port Type   This is the actual or current port type. If the port is offline, this value is the allowed types (or U_Port, if no type constraint is specified). If the port is online, this value is the type to which the port has been configured.

Allowed Port Type   The allowed or configured port type.

The allowed port types indicate any constraints on what types the port can be configured when it comes online. For normal (that is, non-EX_Port) ports, the following are the allowed port types:

L_Port   The port can be used to connect a loop device.

F_Port   The port can be used to connect a non-loop device.

E_Port   The port can be used to connect to another switch.

U_Port   For a physical FC port: the port can be any one of E_Port, F_Port, or L_Port. For a logical FC port: the port can be either VE_Port or VEX_Port.

When the wizard prompts you to select allowed port types, if all of these boxes are selected, there are no constraints on port type. The port negotiated to its preferred type when the switch comes up, depending on what type of device or switch to which it is connected.

Clearing a check box guarantees that the port does not attempt to function as a port of the unchecked type. At least one type must remain selected. L_Port and F_Port, cannot both be cleared. An FC port cannot be configured as an E_Port and L_Port.

NOTE
To configure a port as an EX_Port, the switch must be capable of supporting FCR or FCIP features. The EX_Port option is disabled in the wizard if the switch does not meet these requirements.

Long distance mode

Port long distance configurations can be performed in the **Switch Admin Extended Fabric** tab if the link is used over long distances. To configure the long-distance settings, the EGM license must be enabled on the switch. Otherwise, access to this feature is denied and an error message displays. For information about long-distance mode settings, refer to Chapter 14, "Administering Extended Fabrics".

The EGM license is required only for 8 Gbps platforms, such as the following:

- Brocade DCX and DCX-4S enterprise-class platforms
- Encryption Switch
- 300, 5300, and 5100 switches

For non-8 Gbps platforms, all functionality is available without EGM license.

Ingress rate limit

Ingress rate limiting is a licensed feature that requires the Adaptive Networking license. Ingress rate limiting restricts the speed of traffic from a particular device to the switch port.

Ingress rate limiting is applicable only to F_Ports and FL_Ports and is available only on the following platforms:

- Brocade 300, 5100, and 5300
- Brocade Encryption Switch
- Brocade DCX or DCX-4S

FC Fastwrite

FC Fastwrite reduces the number of round-trip times required to write data.

FC Fastwrite is available as a feature on the 7500 and 7500E Extension switches, Brocade 48000, and Brocade DCX and DCX-4S enterprise-class platforms with FR4-18i blades. When FC Fastwrite is enabled, all GbE ports and FCIP features are disabled.

Assigning a name to a port

Port names are optional. You can assign a name to an FC or FCIP port to make port grouping easier. You can also rename FC and FCIP ports to new names. You cannot rename GbE ports. The **Port Name** column in the **Ports** tab displays the port name, if one exists.

Port names can be from 1 through 32 alphanumeric characters, unless FICON Management Server (FMS) mode is enabled. If FMS mode is enabled, port names should be limited from 1 through 24 alphanumeric characters. The comma (,), semicolon (;), and "at" symbol (@) are not allowed.

Perform the following procedure to assign a name to a port.

1. Click a port in the **Switch View** to open the **Port Administration** window.

2. Select the **FC Ports** tab.

3. From the tree on the left, select the switch or slot that contains the port you want to rename.

4. From the table, select the port you want to rename.

5. Click **Rename**.

6. Enter a name for the port and click **Rename**.

   NOTE
   To delete the existing port name, leave the field blank and click **Rename**.

Port beaconing

Individual FC ports can be set to beacon using the Web Tools Port Admin dialog box. Port beaconing status displays in the Port Beaconing column. The Switch View reflects the port beaconing status by flashing the port amber and green for 2.5 seconds each, in an alternating pattern.

Use the following procedure to configure beaconing for an FC port.

1. Open the **Port Admin** window.

2. Click **Show Advanced Mode**, if the **Port Admin** window is in **Basic Mode**.

3. Select the switch in the **FC Ports Explorer** list.

4. Select a port from the list in the main window.

   The **Port Beacon Enable** or **Port Beacon Disable** button becomes active.

   NOTE
   You may select all the ports on the switch, but if you select a port that is not valid for beaconing, the Port Beacon buttons are disabled.

There is an optional procedure for configuring a single FC port.

1. Open the **Port Admin dialog box**.

2. Click **Show Advanced Mode**, if the **Port Admin** window is in **Basic Mode**.

3. Select a port from the list in the main window.

   The **Port Beacon Enable** or **Port Beacon Disable** button becomes active.

Enabling and disabling a port

Use the following procedure to enable or disable a port.

1. Click a port in the **Switch View** to open the **Port Administration** window.

2. Select the **FC Ports** or **GigE Ports** tab.

3. From the tree on the left, select the switch or slot that contains the port you want to enable or disable.

4. From the table, select one or more ports.

   NOTE
   Use Shift+click and Ctrl+click to select multiple ports. You can select multiple ports from the table. You cannot select multiple ports from the tree.

5. Click either the **Enable** or **Disable** button.

   NOTE
   If the button is gray (unavailable), the port is already in the enabled or disabled state. For example, if the **Enable** button is unavailable, the port is already enabled.

   If you select multiple ports in both enabled and disabled states, both buttons are active. When you click either button, the action is applied to all selected ports.

6. *Optional:* If you are accessing a Brocade 7800 switch, you can set the media type for the GE0 and GE1 GigE ports to either copper or optical.

   a. Select the **GigE Ports** tab.

   b. Select either the **GE0** or **GE1** port.

   c. Select either **Copper** or **Optical** from the **Media Type** selection list.

7. Click **Yes** in the confirmation window.

Considerations for enabling or disabling a port?

You should understand the following limitations and conditions when enabling or disabling a port:

- On FR4-18i and FC4-16IP port blades, all ports are disabled by default. You can disable and re-enable them as needed.

- If a port is not licensed you cannot enable it until you install the appropriate license, such as a **Ports on Demand** or **N_Port ID Virtualization** license (refer to "Port activation" on page 93 for more information). The **Licensed** field located in the **General** tab in the **Port Administration** window indicates whether a port is licensed.

- If you disable a *principal* ISL port (an ISL port that is designated by the fabric to be a part of the path to communicate with the principal switch), the fabric automatically reconfigures.

- If you disable a port that was connected to a device, that device is no longer accessible from the fabric. For more information, refer to *Fabric OS Administrator's Guide*.

Persistent enabling and disabling ports

Use the following procedure to enable or disable an FC port so that it remains enabled or disabled across switch restarts.

NOTE
Ports cannot be persistently enabled or disabled when FMS is enabled.

1. Select a port in the **Switch View** to open the **Port Administration** window.

2. Select the **FC Ports** or **GigE Ports** tab.

3. From the tree on the left, select the switch or slot that contains the port.

4. From the table, select one or more ports.

   NOTE
   Use Shift-click and Ctrl-click to select multiple ports. You can select multiple ports from the table. You cannot select multiple ports from the tree.

5. Click **Persistent Enable** or **Persistent Disable**.

   NOTE
   Persistent Enable or Disable is not supported in FMS mode.

   NOTE
   If the button is gray (unavailable), the port is already in that state or FMS mode is enabled on the switch gray (unavailable), the port is already in the enabled or disabled state. For example, if the **Enable** button is unavailable, the port is already enabled.

   If you select multiple ports in both enabled and disabled states, both buttons are active. When you click either button, the action is applied to all selected ports.

6. *Optional:* If you are accessing a Brocade 7800 switch, you can set the media type for the GE0 and GE1 GigE ports to either copper or optical.

   a. Select the **GigE Ports** tab.

   b. Select either the **GE0** or **GE1** port.

   c. Select either **Copper** or **Optical** from the **Media Type** selection list.

7. Click **Yes** in the confirmation window.

Enabling and disabling NPIV ports

The NPIV license must be installed on a switch before NPIV functionality can be enabled on any port. For detailed information about understanding and configuring NPIV ports, refer to the *Fabric OS Administrator's Guide*. With Web Tools, you can only enable or disable the NPIV functionality on a port.

NOTE
NPIV feature cannot be disabled when Access Gateway mode is enabled.

Perform the following procedure to enable or disable NPIV ports.

1. Select a port in the **Switch View** to open the **Port Administration** window.

2. Select the **FC Ports** tab.

3. From the tree on the left, select the logical port you want to enable or disable.

4. Click **Enable NPIV** or **Disable NPIV**.

Port activation

Brocade switches come with a preset number of ports enabled. Additional ports can be enabled using the Ports on Demand (POD) licenses and the Dynamic Ports on Demand (DPOD) feature (for supported switches only).

Ports on Demand is ready to be unlocked in the switch firmware. The license might be part of the licensed Paper Pack supplied with the switch software, or you can purchase the license separately from your switch vendor, who will provide you with a key to unlock it. You can install up to two Ports on Demand licenses on each switch.

lists the ports that are enabled by default settings and the ports that can be enabled after you install the first and second Ports on Demand licenses for each switch type, and the ports that can be enabled with the Dynamic PODs feature.

TABLE 9   Ports enabled with POD licenses and DPOD feature

Switch name	Enabled by default	Enabled with Ports on Demand licenses	Enabled with the Dynamic Ports on Demand feature
Brocade Encryption	0-15	Not supported	Not supported
Brocade 8000	None	0-7	Not supported
Brocade 5480	1-8, 17-20	9-12, 21, 22 0, 13-16, 23	Any available ports
Brocade 5424	1-8, 17-20	9-16, 21-23 0	Any available ports
Brocade 5300	0-47	48-63 64-79	
Brocade 5100	0-23	24-31 32-39	
Brocade 5000 Brocade 4100	0-15	16-23 24-31	Not supported
Brocade 4900	0-31	32-47 48-63	Not supported
Brocade 4024	1-8, 17-20	9-12, 21, 22 0, 13-16, 23	Any available ports
Brocade 4020	0-7, 15, 16	8, 9, 17-19 10-14	Any available ports
Brocade 4018	0-11	12-17	Any available ports
Brocade 4016	0-7, 10-13	8, 9, 14, 15	Any available ports
Brocade 300	0-7	8-15 16-23	

When using the Brocade 4016, 4018, 4020, and 4024 switches, you can enable the Dynamic Ports on Demand (DPOD) feature, which allows you to select the ports to be enabled (instead of predefined sets of ports) after the POD license is installed. Web Tools allows you only to enable or disable the DPOD functionality on a port. To configure DPOD, refer to the *Fabric OS Administrator's Guide*.

In the **Port Administration** window, the **Licensed** attribute indicates whether a port is licensed (yes), whether it can be license (possible) because there are free licenses available (only applicable with the Dynamic POD feature), or whether it is not licensed and cannot be licensed because there is no available license.

After the license keys are installed, you must enable the ports. You can do so without disrupting switch operation, as described in *"Enabling and disabling a port"* on page 90. Alternatively, you can disable and re-enable the switch to activate all ports as described in *"Enabling and disabling a switch"* on page 40.

To unlock a Ports on Demand license, you can use the supplied license key or generate a license key. If you need to generate a key, open an Internet browser and go to the Brocade website at *www.brocade.com*. Select **Products** > **Software License Keys** and follow the instructions to generate the key.

Enabling Ports on Demand

Use the following procedure to enable Ports on Demand.

1. Install the Brocade Ports on Demand licensed product. For instructions, refer to *"Activating a license on a switch"* on page 47.

2. Enable the ports as described in *"Enabling and disabling a port"* on page 90.

If you remove a Ports on Demand license, the licensed ports are disabled after the next platform restart or the next port deactivation.

Enabling Dynamic Ports on Demand

You must be logged in as Admin to enable the Dynamic POD feature.

NOTE
The Dynamic PODs feature is supported on the Brocade 4018, 4020, 4024, 5460, and 5470 switches only. If you click the **Enable DPOD** button on an unsupported switch, an error message displays.

Use the following procedure to enable Dynamic Ports on Demand.

1. Select a port in the **Switch View** to open the **Port Administration** window.

2. Select the **FC Ports** or **GigE Ports** tab.

3. From the tree on the left, select the switch or the slot that contains the port.

4. Click **Enable DPOD** to enable the licensing mechanism to be dynamic. If the button is labeled **Disable DPOD**, the licensing mechanism is already set to dynamic.

    The existing POD associations and assignments are set as the initial Dynamic POD associations.

    Two fields are displayed:

- **Available Licenses** indicate the number of free licenses. These can be allocated for any port.
- **Total Licenses** indicate the total number of licenses.

Disabling Dynamic Ports on Demand

NOTE
Disabling DPODs causes traffic disruption. Any prior port associations and assignments are lost the next time the switch is restarted.

Log in as Admin and perform the following procedure to disable the Dynamic POD feature.

1. Select a port in the **Switch View** to open the **Port Administration** window.

2. Select the **FC Ports** or **GigE Ports** tab.

3. From the tree on the left, select the switch or the slot that contains the port.

4. Click **Disable DPOD** to set the licensing mechanism to static. If the button is labeled **Enable DPOD**, the licensing mechanism is already set to static.

Reserving and releasing licenses on a port basis

NOTE
If the Admin Domains feature is enabled, the Dynamic POD configuration is only applied to the ports if the switch is a member of the current Admin Domain.

The Dynamic PODs feature is supported on the Brocade 4018, 4020, and 4024 switches only.

Perform the following procedure to reserve and release licenses on a port basis.

1. Click a port in the **Switch View** to open the **Port Administration** window.

2. Click the **FC Ports** or **GigE Ports** tab.

3. From the tree on the left, click the switch or the slot that contains the port.

   The **License** column identifies the port license status:

   - If the port has a license allocated, the **License** field contains the value **Yes**.
   - If the port does not have a license allocated and there are no free licenses that can be allocated, the **License** field contains the value **No**.
   - If the port does not have a license allocated and there are licenses that can be allocated to the port, the **License** field contains the value **Possible**.

   You can reserve or release a license on any port with a a license allocated. You must be logged in as Admin to reserve and release licenses.

   To reserve a license, click **Reserve License** in the **Port Administration** window.

   To release a license, click **Release License** in the **Port Administration** window.

   NOTE
   You must disable the port or switch before reserving or releasing a license.

Port swapping index

If a port malfunctions, or if you want to connect to different devices without having to rewire your infrastructure, you can move traffic from one port to another (*swap ports*) without changing the I/O Configuration Data Set (IOCDS) on the mainframe computer.

You must disable the ports before you do a port swap.

NOTE
Port swapping is not applicable to GE ports because there are no areas assigned to these ports.

Port swapping

In the Port Admin list view and detailed view, swapped ports are indicated with the "(Swapped)" label appended to the Port Index column and field (Figure 30).

Port#	Port Index
0(0x0)	0(0x0)
1(0x1)	4(0x4) (Swapped)
2(0x2)	3(0x3) (Swapped)
3(0x3)	2(0x2) (Swapped)
4(0x4)	1(0x1) (Swapped)
5(0x5)	7(0x7) (Swapped)
6(0x6)	6(0x6)

FIGURE 30   Port swapped label

Use the following procedure to swap ports.

1. Select a port in the **Switch View** to open the **Port Administration** window.

2. Select the **FC Ports** tab.

3. Click **Advanced**.

4. From the tree on the left, select the port you want to swap.

5. Click **Disable.**

   You must disable the ports used for port swapping.

6. Click **Port Swap**.

   NOTE
   When the Port Swap dialog box is launched for a swapped port, the dialog box displays "The Selected port is already Swapped"

7. Enter the number of the port with which you want to swap the current port.

   If the port is on a blade, you must also provide the slot number.

   NOTE
   Port swap is not supported above the 16th port in a 48 port card in FMS mode.

8. Click **Swap**.
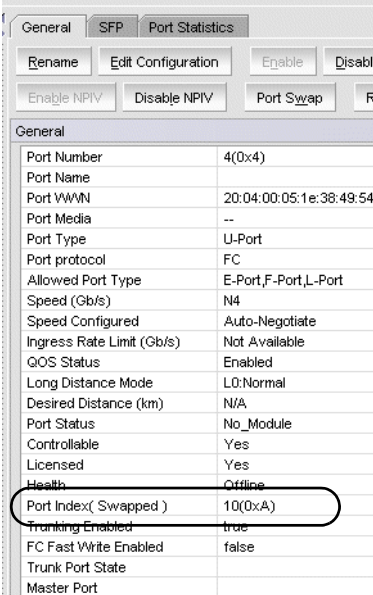
Determining if a port index was swapped with another switch port

Use the following procedure to determine whether a port was swapped.

1. Select a port in the **Switch View** to open the **Port Administration** window.

2. Select the **FC Ports** tab.

3. Click **Show Advanced Mode**.

4. From the tree on the left, select the port you want to swap.

5. Click the **General** tab.

> **NOTE**
> The **Port Index** attribute on the **General** tab indicates whether a port was swapped. For ports that were swapped, the attribute name displays as *Port Index (Swapped)*, as shown in Figure 31. The value indicates with which port index the port was swapped.

FIGURE 31   Port swapping index

Configuring BB credits on an F_Port

In Fabric OS v6.4.0 and later, you can configure the BB credits value on an F_Port. Follow the steps given below.

1. Select a port in the **Switch View** to open the **Port Administration** window.

2. Select the **FC Ports** tab.

3.   Click **Show Advanced Mode.**
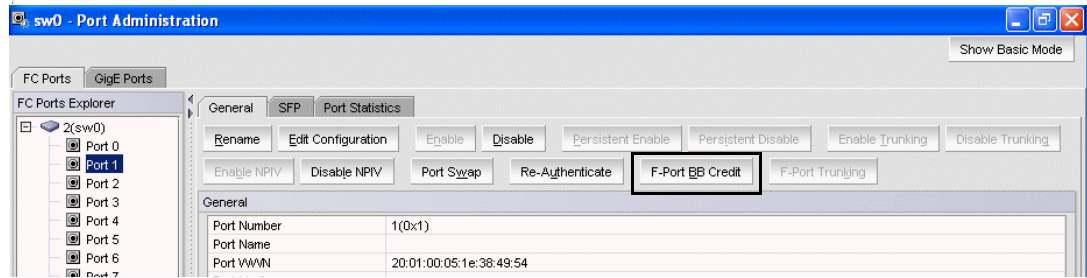
FIGURE 32   F-Port BB Credit button

4.   Click **F-Port BB Credit** as shown in Figure 32.

5.   Enter the BB credit value in the **Enter BB Credit** field (the default value is 8).

NOTE
You cannot modify the default BB credit value for VE and ICL ports.

6.   Click **Ok.**

The value displays in the table of the **Port Administration** window. If no value is configured, the **F-Port BB Credit** column displays the default value.

Enabling ISL Trunking

In this chapter

ISL Trunking overview

Inter-Switch Link (ISL) Trunking optimizes network performance by forming trunking groups that can distribute traffic between switches across a shared bandwidth.

A trunking license is required on each switch that participates in the trunk. For details on obtaining and installing licensed features, refer to "Licensed feature management" on page 47. For additional information about ISL Trunking, refer to the *Fabric OS Administrator's Guide*.

You must use Web Tools with the EGM license to create ISL trunk groups and to manage F_Port trunks.

The EGM license is required only for 8 Gbps platforms, such as the following:

- Brocade DCX and DCX-4S enterprise-class platforms
- Encryption Switch
- 300, 5300, and 5100 switches

For non-8 Gbps platforms, all functionality is available without EGM license.

For detailed information about ISL Trunking configurations and criteria, refer to the *Fabric OS Administrator's Guide*.

Disabling or enabling ISL Trunking

The trunking feature requires using Web Tools with the EGM license. If you attempt to use this feature without the EGM license, an error message displays.

When the trunking license is activated, trunks are automatically established on eligible ISLs and trunking capability is enabled by default on all ports. Trunking is not supported on logical ports or GbE ports.

Use the following procedure to disable trunking on a port, or to re-enable trunking if it has been disabled.

1. Select a port in the **Switch View** to open the **Port Admin** window.

2. Select the **FC Ports** tab.

3. From the tree on the left, select the switch name or slot name.

4. From the table, select the port that you want to trunk.

   You can select multiple ports from the table. You cannot select multiple ports from the tree.

5. Click the **Show Advanced Mode** button on **Ports Admin**.

   Click either the **Trunking Enable** or **Trunking Disable** button.

   If the button is unavailable, then the selected port is already in that state.

6. 6. Click **Yes** in the confirmation dialog box.

Admin Domain considerations

You can only enable and disable trunking for a port when the current Admin Domain owns the switch. You can log in to a switch that is not in your Admin Domain, but most of the functionality is unavailable. F_Port trunking should not be configured in physical fabric mode.

Viewing trunk group information

Use the **Trunking** tab on the **Switch Administration** window to view trunk group information.
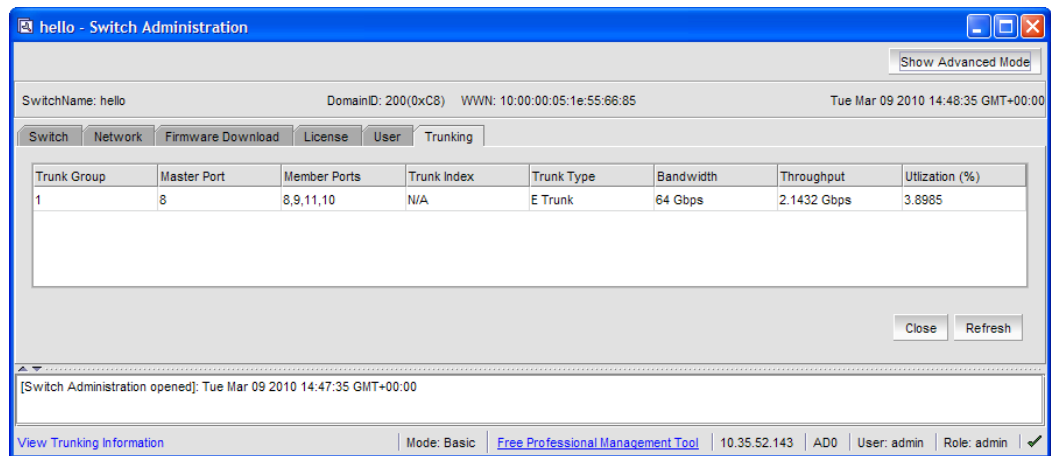
FIGURE 33   **Trunking tab**

The following trunking attributes can be displayed from the **Switch Admin** view:

- Trunk port state, either master or slave.
- Trunk master port
- Trunk index (applies only to F_Port trunking).
- Trunk type
- Bandwidth (shown only for E_Port, Ex_Port, F_Port, and N_Port).
- Throughput (shown only forE_Port, Ex_Port, F_Port, and N_Port).
- Utilization (shown only for E_Port, Ex_Port, F_Port, and N_Port).

Additionally, the following trunking attributes can also be displayed from the **Port Admin** view by clicking the **Show Advanced Mode** button:

- Trunk port state, either master or slave.
- Master Port
- Trunk Index (applies only to F_Port trunking).
- Trunking Enabled

F_Port trunk groups

F_Port trunking provides extra bandwidth and robust connectivity for hosts and targets connected by switches in Access Gateway mode. There are five general criteria for establishing F_Port trunking:

- The F_Port trunking feature requires installing the EGM license; otherwise if you attempt to use this feature in Web Tools without the license, an error message displays.
- Trunking must be enabled on the ports.
- The trunking license must be enabled on the switch in Access Gateway mode.
- The ports should not be configured for long distance connections.
- The ports should not be port-swapped.

When you create an F_Port trunk you create a logical entity called a trunk index (TI), which represents the physical ports. The TI represents all ports in the trunk. If a master port fails, and a slave port takes over, the TI remains the same.

The EGM license is required only for 8 Gbps platforms, such as the following:

- Brocade DCX and DCX-4S enterprise-class platforms
- Encryption Switch
- 300, 5300, and 5100 switches

For non-8 Gbps platforms, all functionality is available without EGM license.

NOTE
If F_Port trunking is configured, a firmware downgrade is not allowed.

Creating and maintaining F_Port trunk groups

The FS8 -18 Encryption blade provides trunk groups with a maximum of eight ports per trunk group. The trunk groups are in the blade port ranges 0-7 and 8-15, which are applicable to front end ports. On the Brocade Encryption Switch, the trunk groups are in the port ranges 0-7, 8-15, 16-23, and 24-31, which are applicable on the front end ports.

User this procedure to create an F_Port trunk group, and to add or remove member ports.

1. Select **Port Admin**.
2. Click **Show Advanced Mode**.
3. Select any port from the port group in which you want to create the trunk group.
4. Select **F_Port Trunking**.

The **F_Port Trunking** dialog box displays.

5. Select one or more ports in the **Ports for trunking** pane.

    A dialog box displays, asking you to select a trunk index.

6. Select the trunk index from the drop-down list populated with the index for all the ports.

    A trunk group is created, identified by the trunk index, and containing the port you selected.

7. Select the trunk group you just created.

    Add Members becomes active.

8. Additional ports can be added by selecting a port from **Ports for trunking** table and then clicking **Add Members**.

NOTE
To remove a port from the trunk group, select the port from **Trunk Groups** table and then click **Remove Members**.

---

9. Click **OK** to save your changes.

Monitoring Performance

In this chapter

Performance Monitor overview

The Web Tools Performance Monitoring tool graphically displays throughput (in megabytes per second) for each port and for the entire switch. To utilize performance monitoring, the EGM license must be enabled on the switch. Otherwise, when you select **Performance Monitor** tab, access to this feature is denied and an error messages displays

Basic monitoring

The **Basic Monitoring** menu is standard in the Web Tools software. Any user logged into Web Tools with an associated role of zoneadmin or securityadmin cannot open **Performance Monitor**. The roles user, operator, and basicswitchadmin are allowed to perform basic monitoring tasks, except save or display canvas operations in any Admin Domain context. Only users with the admin, switchadmin and fabricadmin roles associated with their login accounts are able to save or display a canvas. Use the **Basic Monitoring** option in the **Performance Graphs** window to do the following:

• Create user-definable reports.

• Display a performance canvas for application-level or fabric-level views.

• Save persistent graphs across restarts (saves parameter data across restarts).

Advanced monitoring

The **Advanced Monitoring** menu is an optionally licensed feature. To utilize the **Advanced Monitoring** feature you must have a Performance Monitor license installed and you must log in using an account with an admin, switchadmin, or fabricadmin role.

The **Advanced Monitoring** option in the **Performance Graphs** window displays predefined reports and filter-based performance monitoring. You can use this feature to track the following:

- The number of words received and transmitted in Fibre Channel frames with a defined SID/DID pair.
- The number of times a particular filter pattern in a frame is transmitted by a port.

For detailed information on performance monitoring, refer to the *Fabric OS Administrator's Guide*.

Performance graphs

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

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

When you have multiple graphs open in the **Performance Monitoring** window, you can perform the following tasks:

- Select **Window** > **Tile** to view all graphs at once, tiled in the **Performance Monitoring** window.
- Select **Window** > **Cascade** to view one graph at a time.
- Select **Window** > **Close All** to close all open Performance Monitor graphs in the **Performance Monitoring** window.

In addition, the **Window** menu lists all open graphs. You can click **Window**, and then select a graph name to view that graph.

The **Tunnel and TCP Graph** option in the **Performance Graphs** window displays real time performance monitoring charts for Brocade 7800 Extension Switch and FX8-24 DCX Extension Blade. This option is not available on other platforms.

Admin Domain considerations

You must consider the following when configuring Admin Domain:

- If you are not the switch owner, only directly and indirectly-owned E_Ports, including EX_Ports are available.
- You can use the Advanced Performance Monitoring feature only in AD255 if there are user defined Admin domains or in AD0 if there are no other user-defined Admin Domains. Otherwise, access to Advanced Monitoring features in the **Performance Graphs** menu are unavailable.
- It is recommended that you define a user with a switchadmin role and give that user access to AD255 for the purpose of data collecting using the **Advanced Performance Monitor**.

Predefined performance graphs

Web Tools predefines basic graph types to simplify performance monitoring. A wide range of end-to-end fabric, LUN, device, and port metrics graphs are included.

Table 10 lists the basic monitoring graphs available. Table 11 lists the advanced monitoring graphs.

The advanced monitoring graphs give more detailed performance information to help you manage your fabric. You can access the basic monitoring graphs on all switches; advanced monitoring graphs are available only on switches that have a Brocade Advanced Performance Monitoring license activated.

TABLE 10   Basic performance graphs

Graph type	Display description
Port Throughput	The performance of a port, in bytes per second, for frames received and transmitted.
Switch Aggregate Throughput	The aggregate performance of all ports on a switch.
Blade Aggregate Throughput	The aggregate performance of all ports on a port card. This graph is available only for the Brocade 48000 and Brocade DCX and DCX-4S enterprise-class platforms.
Switch Throughput Utilization	The port throughput, in Gbps at the time the sample is taken. For the Brocade 48000 and Brocade DCX and DCX-4S enterprise-class platforms, this graph displays the throughput for each slot. You can customize this graph to display information for particular ports.
Port Error	CRC errors for a given port.
Switch Percent Utilization	The percentage utilization for each port in a switch. For the Brocade 48000 and Brocade DCX, this graph displays the percent utilization for each slot. You can customize this graph to display information for particular ports.
Port Snapshot Error	The CRC error count between sampling periods for all the ports on a switch. For the Brocade 48000 and Brocade DCX and DCX-4S enterprise-class platforms, this graph displays the CRC error rate for each slot. You can customize this graph to display information for particular ports.

TABLE 11   Advanced performance monitoring graphs

Graph type	Display description
SID/DID Performance	The traffic between the SID-DID pair on the switch being managed. The member selection list in the LHS displays the port in the current switch only. The **All Devices** tab lists all the devices in the fabric. SID/DID Performance can be used to select the source and destination. For more information, refer to "Creating SID-DID Performance graphs" on page 111.
SCSI vs. IP Traffic	The percentage of SCSI versus IP frame traffic on each individual port. For more information, refer to "Creating the SCSI vs. IP Traffic graph" on page 112.
SCSI Commands by port and LUN (R, W, R/W)	The total number of read/write commands on a given port to a specific LUN. For more information, refer to "Creating SCSI command graphs" on page 112.

The Brocade Encryption Switch, FS8-18 blade, FR4-18i blade, FX8-24 DCX Extension blade, and the Brocade 7500, 7500E, and 7800 Extension switches include physical FC ports, logical FC ports, and GbE ports. The Brocade FC4-16IP blade includes physical FC ports and GbE ports. Not all of the performance monitoring graphs support the logical FC ports and GbE ports.

Table 12 lists each graph and indicates the supported port types for each graph. The port selection columns for each graph displays the supported ports.

TABLE 12   Supported port types for Brocade switches

Graph type	Physical FC ports	Logical FC ports	GbE ports
Port Throughput	P	P	P
Switch Aggregate Throughput	N/A	N/A	N/A
Blade Aggregate Throughput[1]	N/A	N/A	N/A
Switch Throughput Utilization	P	N/A	P
Port Error	P	P	P
Switch Percent Utilization	P	N/A	P
Port Snapshot Error	P	P	N/A
SID/DID Performance	P	P	N/A
SCSI Commands	P	N/A	N/A
SCSI vs. IP Traffic	P	N/A	N/A

1. The Blade Aggregate Throughput graph is supported only on the Brocade 48000 and Brocade DCX and DCX-4S enterprise-class platforms.

The labeling of the axes in the graphs depends on the switch type:

- For the Brocade 48000 and DCX and DCX-4S enterprise-class platforms, slot numbers are displayed with expansion arrows next to them, as shown in Figure 34 on page 107. Click the arrows to expand and contract the list of ports per slot.

- Switches such as the Brocade 200E, 300, 4012, 4016, 4018, 4020, 4024, 4100, 4900, 5000, 5100, 5300, 7600, 7500, the 7500E, and 7800 Extension switches, and the Brocade Encryption Switch do not have slot numbers because they have no blade FRUs, and therefore there is no need for slot numbering.

- For Brocade 48000 and Brocade DCX and DCX-4S enterprise-class platforms, the X-axis scales up to 409.6 Gbps in multiples of 2.

- For Brocade 200E, 300, 4012, 4016, 4018, 4020, 4024, 4100, 4900, 5000, 5100, 5300, 7600, the 7500, 7500E, and 7800 Extension switches, and the Brocade Encryption Switch, the X-axis scales up to 8.0 Gbps in increments of 0.8 Gbps.

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

NOTE
Virtual ports on logical switches cannot be graphed.

Figure 34 shows how to access the list of Advanced Performance Monitoring graphs using Web Tools with the EGM license. This example displays the graphs available in the Performance Monitoring window for a Brocade 48000 director with the Advanced Performance Monitoring license installed. Note that the slot number is identified.
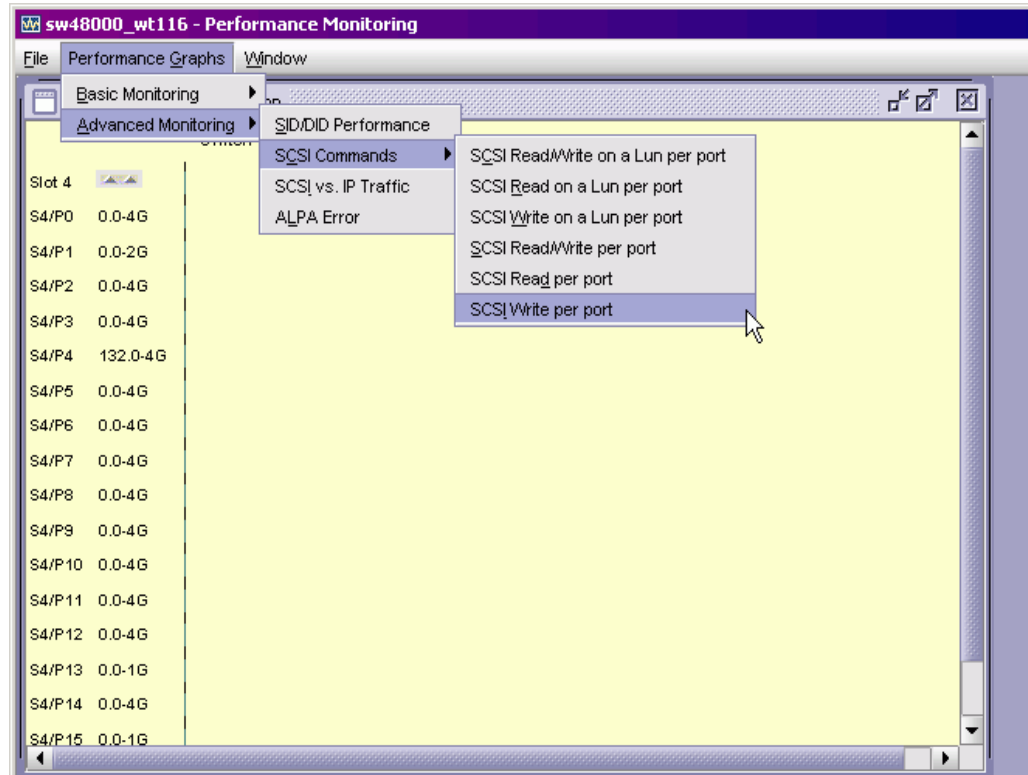
FIGURE 34   Accessing performance graphs

User-defined graphs

You can modify the predefined graphs to create your own customized graphs (refer to "Customizing basic monitoring graphs" on page 109 for more information). These user-defined graphs can be added and saved to canvas configurations.

Canvas configurations

A canvas is a saved configuration of graphs. The graphs can be either the Web Tools predefined graphs or user-defined graphs. Each canvas can hold up to eight graphs per window, with six shown in Figure 35. Up to 20 canvases can be set up for different users or different scenarios. Each canvas is saved with a name and an optional brief description.
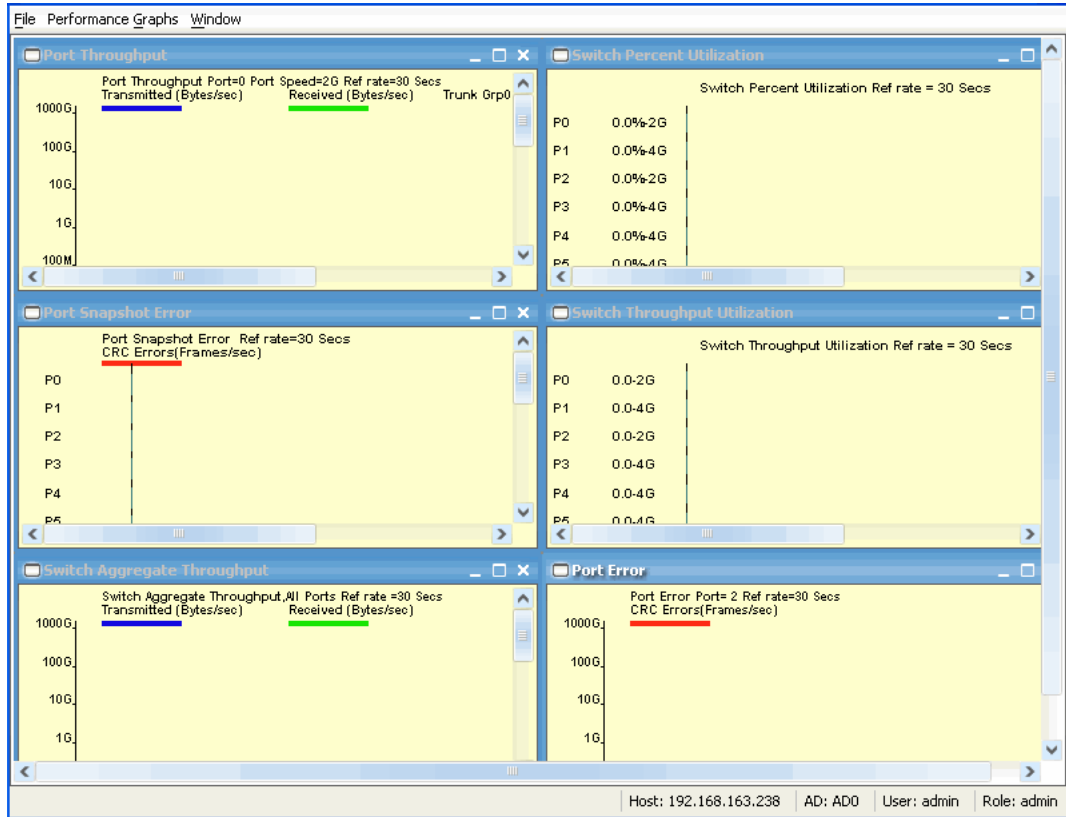
FIGURE 35   Canvas of six performance monitoring graphs

Opening the Performance Monitoring window

To perform performance monitoring, you must use Web Tools with the EGM license; otherwise, when you click on the **Performance Monitor** tab, access to this feature is denied and an error messages displays.

Use the following procedure to open the **Performance Monitoring** window.

1. Select a switch from the Fabric Tree and log in when prompted.

2. In the **Monitor** area under **Tasks**, click **Performance Monitor**. The **Performance Monitor ing** window displays.

Creating basic performance monitor graphs

Use the following procedure to create the basic performance monitor graphs listed in Table 10 on page 105.

1. Open the **Performance Monitoring** window.

2. Select **Performance Graphs** > **Basic Monitoring** > *Graph Type*.

   Depending on the type of graph you select, you might be prompted to select a slot or port for which to create a graph.

3. If prompted, drag the port into the **Enter/drag slot,port** field, or manually enter the slot and port information in the field, in the format *slot,port*.

   NOTE
   For Brocade 200E, 300, 4012, 4016, 4018, 4020, 4024, 4100, 4900, 5000, 5100, 5300, 7600 switches, the 7500, 7500E and 7800 Extension switches, the 8000, and the Encryption Switch, enter only a port number.

4. Click **OK**.

   The graph is displayed in a window in the **Performance Monitoring** window.

Customizing basic monitoring graphs

You can customize some of the basic performance monitoring graphs to display information for particular ports. For the Brocade 48000 and Brocade DCX and DCX-4S enterprise-class platforms, you can also customize these graphs to display information for a slot.

You can customize the following graphs:

- Switch Throughput Utilization
- Switch Percent Utilization
- Port Snapshot Error

The following procedure assumes that you already created one of these customizable graphs.

1. Create or access the graph you want to customize.

   Refer to *"Creating basic performance monitor graphs"* on page 109 for instructions on creating a graph.

2. For Brocade 48000 and Brocade DCX and DCX-4S enterprise-class platforms, display the detailed port throughput utilization rates for each port in a slot by clicking the arrows next to a slot. The port information for that slot displays in the graph.

   NOTE
   For Brocade 200E, 300 4012, 4016, 4018, 4020, 4024, 4100, 4900, 5000, 5100, 5300, 7600, the 7500, 7500E, and 7800 Extension switches, the 8000, and the Brocade Encryption Switch, proceed to step 3.

3. To display detailed port throughput utilization rates for particular ports only, right-click anywhere in the graph and click **Select Ports**.

   The setup dialog box displays, as shown in Figure 36.

The title of the dialog box varies, depending on the type of graph you are customizing, but the layout of the dialog box is the same. Figure 36 displays an example of the setup dialog box for the **Edit Switch Throughput Utilization** graph.
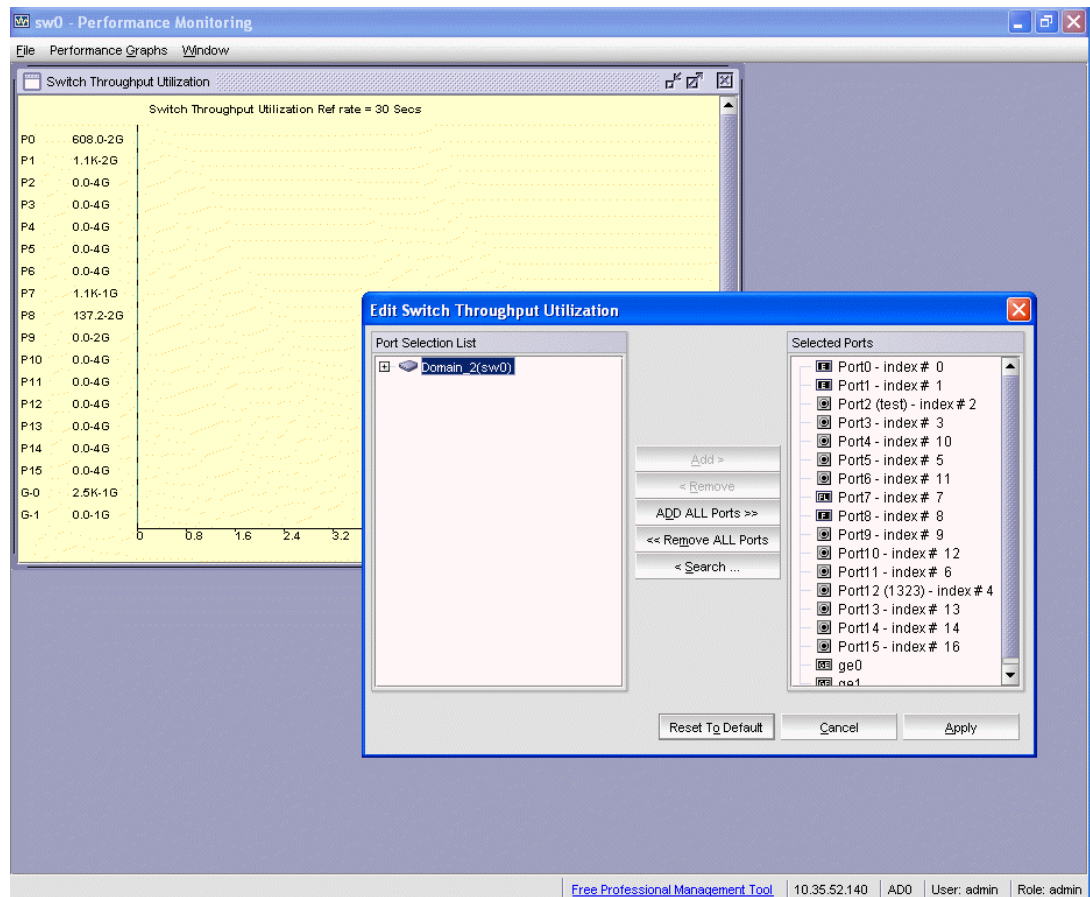
FIGURE 36   Select Ports for customizing the Switch Throughput Utilization graph

You can perform the following in the dialog box:

a. Double-click the domain to expand the slot or port list.

NOTE
For the Brocade 48000, Brocade DCX and Brocade DCX 4S enterprise-class platforms, click the plus (+) signs to expand the ports under each slot, as shown in Figure 36.

b. Click the port you want to monitor in the graph in the **Port Selection List**.

Use Shift+click and Ctrl+click to select multiple ports.

c. Click **Add** to move the selected ports to the Selected Ports list.

d. *Optional*: Click **ADD ALL Ports** to add all of the ports in the **Port Selection List** to the Selected Ports list.

e. *Optional*: Click **Search** to open the **Search Port Selection List** dialog box, from which you can search for all E_Ports, all F_Ports, or all port names with a defined string. Select the ports you want to add and click **Search** in the **Search Port Selection List** dialog box.

f. Click **Apply**.

Only the selected ports are displayed in the graph.

Advanced performance monitoring graphs

This section describes how to create the advanced performance monitor graphs listed in Table 11 on page 105. Because the procedure for creating these graphs differs depending on the type of graph, each type is described separately in the sections that follow.

The advanced monitoring graphs are not supported for GbE ports.

NOTE
You must have an Advanced Performance Monitoring license installed to use the Advance Performance Monitor features. If user-defined Admin Domains are configured, Advanced Performance Monitoring works only in AD255.

Creating SID-DID Performance graphs

The SID/DID Performance graph displays the traffic between a SID-DID pair on the switch being managed.

Use the following procedure to create a SID-DID performance graph.

1. Open the **Performance Monitoring** window.

2. Select **Performance Graphs** > **Advanced Monitoring** > **SID/DID Performance**. The **SID/DID Performance Setup** dialog box displays.

   - To see the end-to-end (EE) monitors that are currently set up on a particular port, proceed to step 3.

   - To specify the port, Source ID and Domain ID, skip to step 4.

   NOTE
   Only the FC ports of the launched switch display in the tree. The **All Devices** tab lists all the devices in the fabric and lets you select the source and destination. Slot and port information of other switches is not displayed in the tree.

3. Click a port from the **Slot/Port** or **Sid/Did Selection List**.

   a. Drag the selected port into the **Enter/drag slot, port number** field.

   b. Click **Retrieve preset EE monitors**. The current end-to-end monitors for that port are displayed in the "Current EE monitors set for selected port" table.

   c. *Optional*: To display a performance graph for the current EE monitors set for the selected port, click a SID-DID pair in the table. You can select multiple Source ID and Destination IDs. Click **Select**. If you selected multiple SID/DID monitors, click **OK** in the confirmation dialog box that displays. Skip to step 6. If you do not want to display a performance graph for the current EE monitors set for the selected port, continue with step 4.

4. Select a source ID from the **Port or Sid/Did Selection List**, and click **Add Sid**.

   You can also enter a source ID in the **Enter/drag SID number** field.

5. Select a destination ID from the **Port or Sid/Did Selection List**, and click **Add Did**.

You can also enter a destination ID in the **Enter/drag DID number** field.

6. Click **OK**.

   If you selected multiple EE monitors, SIDs, or PIDs, a confirmation dialog box displays, reminding you that one graph is opened for each selection.

7. Click **Yes** to display the graphs.

8. When you close a graph, a dialog box asks if you want to save the monitor.

   If you click **OK**, the monitor is saved, and persists if the switch is restarted.

Creating the SCSI vs. IP Traffic graph

The SCSI vs. IP Traffic graph displays the SCSI versus IP traffic for selected ports. For Brocade 48000 and Brocade DCX and DCX-4S enterprise-class platforms, the slot and port name are identified in the graph.

In a trunk group, the SCSI vs. IP Traffic graph displays only the master port and not the slave ports.

Use the following procedure to create a SCSI vs. IP Traffic graph.

1. Open the **Performance Monitoring** window.

2. Select **Performance Graphs** > **Advanced Monitoring** > **SCSI vs. IP Traffic**.

   The **SCSI vs. IP Traffic Setup** dialog box displays. This dialog box is similar to that shown in Figure 36 on page 110.

3. Double-click the domain to expand the slot/port list.

   NOTE
   For Brocade 48000 and Brocade DCX and DCX-4S enterprise-class platforms, click the plus (+) signs to expand the ports under each slot, as shown in Figure 36.

4. Click the port you want to monitor in the graph in the **Port Selection List**. Use Shift+click and Ctrl+click to select multiple ports.

5. Click **Add** to move the selected ports to the Selected Ports list.

6. *Optional*: Click **ADD ALL Ports** to add all of the ports in the **Port Selection List** to the Selected Ports list.

7. *Optional*: Click **Search** to open the **Search Port Selection List** dialog box, from which you can search for all E_Ports, all F_Ports, or all port names with a defined string. Select the ports you want to add and click **Search** in the **Search Port Selection List** dialog box.

8. Click **Apply** in the **SCSI vs. IP Traffic Setup** dialog box.

   Only the selected ports are displayed in the SCSI vs. IP traffic graph.

Creating SCSI command graphs

This graph displays the total number of read or write (or both) commands on a given port or to a specific LUN on a given port.

Use the following procedure to create a SCSI command graph.

1. Open the **Performance Monitoring** window.

2. Select **Performance Graphs** > **Advanced Monitoring** > **SCSI Commands** > *Graph Type*.

   The applicable setup dialog box displays.

3. Navigate to a **switch** > **slot** > **port** in the **Port Selection List**.

4. Click the port from the **Port Selection List** and drag it into the **Enter/drag port** field.

5. *Optional*: For the LUN per port graphs, enter a LUN number, in hexadecimal notation.

   For the Brocade Encryption Switch, you can enter up to eight LUN masks
   For the Brocade 5100, 5300, 300, 7500, 7800, and 8000, you can enter up to eight LUN masks
   For the Brocade 4100 or 5000 Switch, you can enter up to eight LUN masks.
   For the Brocade 48000 director, you can enter up to four LUN masks.
   For all other switches running Fabric OS 4.x or v5.x, you can enter up to two LUN masks.
   For switches running Fabric OS 3.x, you can enter up to three LUN masks.

6. Click **OK**.

   The selected graph displays in the canvas.

Tunnel and TCP performance monitoring graphs

This section describes how to generate the Tunnel and TCP performance monitor graphs. You can launch maximum of four Tunnel and TCP graphs for a switch at a time. A total of 16 TCP connection graphs can be launched for a switch.

The TCP graphs available are:

- Sender RoundTrip
- Sender RoundTripVariance
- TCP DupAck
- TCP OOS
- TCP SlowStart
- TCP FastRetransmit
- TCP Tx(MB/sec)
- TCP Rx(MB/sec)

The Tunnel graphs available are:

- Throughput(MB/sec)
- Effective Throughput(MB/sec)
- CompressionRatio

For TCP connection graphs, tool tip is displayed only for all selected connections.

Use the following procedure to create a Tunnel and TCP graph.

1. Select **Monitor > Performance Monitoring**.

   The **Performance Monitoring** window displays.

2. Select **Performance Graphs** > **Tunnel and TCP Graph**.

The **Tunnel and TCP Graph** dialog box displays.

3. Select the tunnel from the **Tunnels** drop-down list for which you want to generate the graphs.

    For Brocade 7800 extension switch, you can have maximum four circuit connections in a tunnel and for FX8-24 DCX extension blade, you can have maximum 10 circuit connections in a tunnel.

4. In the **Tunnel and TCP** area at the bottom of the screen, select the required check boxes for the statistic you want to graph.

    Note that each column represents a different graph.

5. Click **Options** to set the display options for the graphs.

    *   **Range:** The range is 3-30 seconds. The X axis is limited to 30 minutes. The graph scale starts with 0 minutes and auto scales to draw the statistics. Once the 30 minutes graph is drawn, the first minute data is removed to accommodate the 31st minute values.

    *   **Global auto scaling**: By default, this option is in disabled mode. User can either enable or disable this option. If enabled, the graph's X-axis scale up to 30 minutes and if it is disabled, the X-axis will scale up to 10 minutes

    *   **Number of graphs per row**: Designate how many graphs you wish to appear in each row.

6. Click **Generate.**

7. Click **Reset** to reset all the graphs.

Tunnel and TCP graph chart properties

When a Tunnel and TCP graph displays, you can right-click the graph to access the display properties.

These properties include:

*   Font selection
*   Background color selection
*   Title text
*   Display zoom

These value selections are not persistent. When you close the graph, these values reset to the default settings.

In addition, you can print the graph and save the graph to a file

Saving graphs to a canvas

Saving graphs is useful when you create customized graphs and do not want to recreate them each time you access the **Performance Monitoring** window. When you save graphs, you must save them to a canvas.

The following procedure describes how to save graphs to a new canvas.

1. Open the **Performance Monitoring** window.

2. Create basic or advanced Performance Monitor graphs, as described in "Creating basic performance monitor graphs" on page 109 and "Advanced performance monitoring graphs" on page 111.

   The graphs display in the **Performance Monitor** window.

3. Select **File** > **Save Current Canvas Configuration**.

   The **Save Canvas Configuration** dialog box displays.

4. Enter a name and description for the configuration and click **Save Canvas**. A message displays, confirming that the configuration was successfully saved to the switch.

Adding graphs to an existing canvas

The following procedure assumes that a canvas is already created.

To create a new canvas, you must first create graphs, as described in "Creating basic performance monitor graphs" on page 109 and "Advanced performance monitoring graphs" on page 111, and then save those graphs to a canvas, as described in "Saving graphs to a canvas" on page 114.

Use the following procedure to add a graph to an existing canvas.

1. Select **File** > **Display Canvas Configurations**.

   The **Canvas Configuration List** displays. The error message "No Canvas configuration to display" displays if there are no saved canvas configurations.

2. Click a canvas in the list.

3. Click **Edit**.

   The **Edit Canvas** dialog box displays.

4. Click **Add**.

   A list of graphs displays.

5. Click a graph to add it to the canvas, and click **Save**.

Printing graphs

You can print a single graph or all the graphs displayed on the selected canvas configuration. Only one canvas configuration can be opened at a time.

Use the following procedure to print a graph.

1. Open the **Performance Monitoring** window.

2. Create a basic or advanced Performance Monitor graph as described in "Creating basic performance monitor graphs" on page 109 and "Advanced performance monitoring graphs" on page 111.

3. To print a single graph, right-click the graph and choose **Print**. To print all the graphs displayed on the selected canvas configuration, select **File** > **Print All Graphs**.

4. In the print dialog box, click **OK**.

Modifying graphs

Use the following procedure to modify an existing graph that is saved in a canvas.

1. Open the **Performance Monitoring** window.

2. Select **File** > **Display Canvas Configurations**.

   The **Canvas Configuration List** displays. A message "No Canvas configuration to display" displays if there are no saved canvas configurations.

3. Select a canvas from the list and click **Edit**.

   The **Performance Monitor Canvas:** *Canvas Name* dialog box displays.

4. Select a graph from the list and click **Edit**.

   NOTE
   The **Edit** button is enabled only for the graphs that are configurable or editable.

   ---

5. Make changes in the **Edit** dialog box, as necessary.

6. Click **OK** to close the **Edit** dialog box.

7. Click **Save** to save the changes and close the **Performance Monitor Canvas** dialog box.

8. Click **Close** to close the **Canvas Configuration List**.

Administering Zoning

In this chapter

Zoning overview

This chapter describes zoning and provides the procedures for managing zoning. The **Zoning Administration** window provides two zoning options on the left pane:

- Basic zones

- Traffic isolation zones

You can perform basic zoning and traffic isolation zones using Web Tools and Web Tools with the EGM license.

Table 13 describes the zoning features that are no longer supported in Web Tools, but are supported in DCFM Professional and Enterprise Edition.

TABLE 13   Zoning features supported in DCFM Professional and Enterprise Edition

Description	DCFM Professional Edition	DCFM Enterprise Edition
Live Fabric library scope	yes	yes
QoS support	yes	yes
Remove offline devices	no	yes
Device to zone, zoneset participation analysis	yes	yes
Rolling back an already activated zone DB	no	yes
Importing/Exporting of a zone DB to/from file system in XML format	no	yes

Basic zones

Basic zoning enables you to partition a storage area network (SAN) into logical groups of devices that can access each other. For example, you can partition a SAN into two zones, *winzone* and *unixzone*, so that the Windows servers and storage do not interact with UNIX servers and storage.

Zones can be configured dynamically. They can vary in size, depending on the number of fabric-connected devices, and devices can belong to more than one zone. Because zone members can access only other members of the same zone, a device not included in a zone is not available to members of that zone.

Traffic Isolation zones

A traffic isolation zone (TI zone) is a special zone that creates a dedicated path for a specific traffic flow. TI zones are primarily for shaping and controlling traffic rather than partitioning access to storage.

LSAN zone requirements

An LSAN zone enables device connectivity between fabrics connected in Fibre Channel Routing (FCR) configurations without forcing you to merge fabrics. Extension switches like the 7500 provide multiple mechanisms to manage interfabric device connectivity. Zones that contain hosts and targets that are shared between the two fabrics need to be explicitly coordinated. To share devices between any two fabrics, you must create an LSAN zone in both fabrics considering the following:

- The name of an LSAN begins with the prefix LSAN_. The prefix is not case sensitive.
- Members must be identified by their port WWN because port IDs are not necessarily unique across fabrics.

QoS zone requirements

A QoS zone is a special zone that assigns a Quality of Service (QoS) level for traffic flow between a given host or target pair. The members of a QoS zone are WWNs of the host or target pairs. QoS zones can contain only WWN members. A QoS zone has a special prefix, to differentiate it from a regular zone. The formats and meaning of the QoS zone name prefix are shown in Table 14 (the names are not case dependent).

TABLE 14   QoS zone name prefixes

QoS name prefix	Priority	Bandwidth assignment
QosH_	High	Five virtual circuits, 60% of available bandwidth
QosM_	Medium	Four virtual circuits, 40% of available bandwidth
QosL_	Low	Two virtual circuits, 10% of available bandwidth

Zoning configurations

The **Zone Administration** window is where all of the zoning tasks are performed.

When performing zoning tasks for switches in a mixed fabric—that is, a fabric containing two or more switches running different fabric operating systems—you should use the switch with the highest Fabric OS level. Refer to "Best practices for zoning" on page 138 for more recommendations about zoning.

Opening the Zone Administration window

You cannot open the **Zone Administration** window from AD255 (physical fabric).

Use the following procedure to open a Zone Adminstration window.

1. Select a switch from the Fabric Tree.

2. Click **Zone Admin** in the **Manage** section of the **Tasks** menu.

   The **Zone Administration** dialog box displays.

Setting the default zoning mode

The default zoning mode has two options:

- **All Access**—All devices within the fabric can communicate with all other devices.
- **No Access**—Devices in the fabric cannot access any other device in the fabric.

Web Tools supports default zoning on switches running firmware v5.1.0 or later. Default zoning on legacy switches (switches running firmware versions prior to v 5.1.0) are not supported. Legacy switches can use default zoning; however, they cannot manipulate the default zone or default configuration.

NOTE
To use Admin Domains, you must set the default zoning mode to No Access prior to setting up the Admin Domains. To use the Admin Domain feature, the EGM license must be enabled on the switch; otherwise, access to this feature is denied. You cannot change the default zoning mode to All Access if user-specified Admin Domains are present in the fabric.

Use the following procedure to set the default zoning mode.

1. Open the **Zone Administration** window (refer to "Opening the Zone Administration window" on page 121).

2. Select **Zoning Actions** > **Set Default Mode**, and then select the access mode.

Zoning management

You can monitor and manage basic and traffic isolation zoning through the Web Tools Zone Administration. The information in the **Zone Administration** window is collected from the selected switch.

If the FCS policy is activated in the fabric, zoning can be administered only from the primary FCS switch. If the selected switch has an Advanced Zoning license installed, but is not the primary FCS switch, the **Zone Admin** option is displayed, but not activated.

You must be logged into the switch using a user name with one of the following roles associated with it to make changes to the zoning: zoneAdmin, admin, or fabricAdmin. All other roles allow only a view or read-only access. Most of the zoning operations are disabled in read-only mode.

A snapshot is taken of all the zoning configurations at the time you launch the **Zone Administration** window; this information *is not updated automatically* by Web Tools. To update this information, refer to *"Refreshing Zone Administration window information"* on page 124.

When you log in to a virtual switch, or select a virtual switch using the drop-down list under **Fabric Tree** section in the Switch Explorer window, only the ports that are associated with the Virtual Fabric ID you selected are displayed in the member selection list, as shown in Figure 37. You can use the **Add Other** button to add ports of other switches in the fabric.
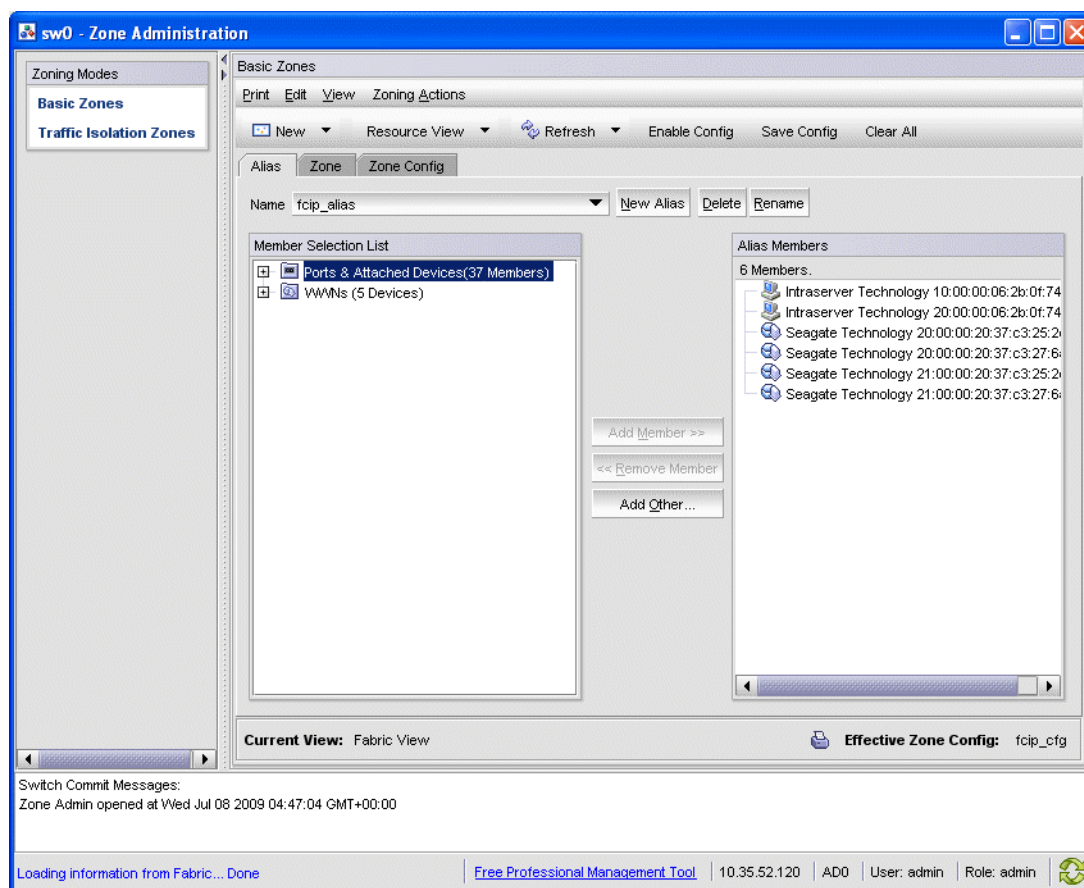
FIGURE 37   Zone Administration window

ATTENTION

Any changes you make in the **Zone Administration** window are held in a buffered environment and are not updated in the zoning database until you save the changes. If you close the **Zone Administration** window without saving your changes, your changes are lost. To save the buffered changes you make in the **Zone Administration** window to the zoning database on the switch, refer to *"Saving local zoning changes"* on page 124.

Note the following:

- "Saving" means updating the zoning database on the switch with the local changes from the Web Tools buffer.
- "Refreshing" means copying the current state of the zoning database on the switch to the Web Tools buffer, overwriting its current contents.

In the **Zone Administration** window, all WWNs also display vendor names.

NOTE
The **Member Selection List** only lists the ports of the current switch and the devices of all the switches in the fabric. Slot and port information of other switches are not displayed in the tree.

Click the **Alias** tab to display which aliases the port or device is a member of. Also, you can right-click the device nodes and click **View Device Detail** to display detailed information about the selected device.

The **Member Selection List** panel displays only physical FC ports. To verify whether you have any unzoned devices, you must use DCFM Professional or Enterprise Edition to analyze zone configurations.

The **Member Selection List** displays virtual initiators if the chassis has an FC4-16IP blade in it; they are shown under a separate folder icon called Virtual Initiators. If the chassis has the Brocade 7500E Extended Switch license installed, the tree displays only two FC ports; otherwise all logical ports display if you have the 7500 without the extended license installed.

Admin Domain considerations: The **Member Selection List** panel displays a filtered list of ports that are:

- Direct port members that are zoneable and are displayed in the tree.
- Indirect port members to which owned devices are attached are displayed in the tree, but cannot be added to a zone or alias.
- Direct device members that are zoneable and are displayed in the tree.
- Indirect device members (devices that are currently attached to owned ports) that are also zoneable and displayed in the tree. But if such a device is later moved to a non-owned port it will no longer be displayed or zoneable.
- Switches and blades that are displayed only if they contain owned ports or devices, regardless of switch ownership, such as the FS8-18 Encryption blade.
- Ports that are indirect members only because the switch is owned are not displayed.

NOTE
When no user- defined Admin Domains are present on the switch, AD0 displays the port count. If there are user-defined Admin Domains, AD0 does not show the port count and the user-defined AD displays the port count.

Refreshing fabric information

This function refreshes the display of fabric elements only (switches, ports, and devices). It does not affect any zoning element changes or update zone information in the **Zone Administration** window. You can refresh the fabric element information displayed at any time.

To refresh fabric information.

1. Open the **Zone Administration** window.

2. Select **View** > **Refresh From Live Fabric**.

    This refreshes the status for the fabric, including switches, ports, and devices.

    NOTE
    Depending on the role associated with your user name or if the switch is owned by the current Admin Domain you are logged in to, you may not be able to modify zones or ports in other Admin Domains.

Refreshing Zone Administration window information

The information displayed in the **Zone Administration** window is initially a snapshot of the contents of the fabric zoning database at the time the window is launched. Any changes you make to this window are saved to a local buffer; but they are not applied to the fabric zoning database until you invoke one of the transactional operations listed in the **Zoning Actions** menu.

Any local zoning changes are buffered by the **Zone Administration** window until explicitly saved to the fabric. If the fabric zoning database is independently changed by another user or from another interface (for example, the CLI) while Web Tools zoning changes are still pending, the refresh icon starts to blink (after a 15–30 second polling delay). You can then decide to refresh the current Web Tools zoning view to reflect the new, externally changed contents of the fabric zoning database, in which case any pending local changes are lost, or you can ignore the blinking refresh icon and save your local changes, overwriting the external changes that triggered the icon to blink.

You can refresh zoning to back out of current, unsaved work and start over.

You can refresh the zoning information at any time, either using the refresh icon (whether it is flashing or not) or from the **View** menu.

The following procedure updates the information in the **Zone Administration** window with the information saved in the zoning database on the switch.

ATTENTION
When you refresh the buffered information in the **Zone Administration** window, any zoning configuration changes you made *and not yet saved* are erased from the buffer and replaced with the currently enabled zone configuration information that is saved on the switch.

Use the following procedure to refresh the **Zone Administration** window.

1. Launch the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select **View** > **Refresh Zoning** or click **Refresh**.

This re-displays the information in the **Zone Administration** window with the information in the switch's zoning database. This action also refreshes the fabric information as described in "Refreshing fabric information" on page 123. Any unsaved zoning changes are deleted.

Saving local zoning changes

All information displayed and all changes made in the **Zone Administration** window are buffered until you save the changes. In that case any other user looking at the zone information for the switch do not see the changes you have made until you save them.

Saving the changes propagates any changes made in the **Zone Administration** window (buffered changes) to the zoning database on the switch. If another user has a zoning operation in progress at the time that you attempt to save changes, a warning displays that indicates that another zoning transaction is in progress on the fabric. You can select to abort the other transaction and override it with yours.

If the zoning database size exceeds the maximum allowed, you cannot save the changes. The zoning database summary displays the maximum zoning database size.

This action updates the entire contents of the **Zone Administration** window, not just the selected zone, alias, or configuration. You can save your changes at any time during the Zone Administration session.

Use the following procedure to save the local zone changes.

1. Make the zoning changes in the **Zone Administration** window.

2. Select **Zoning Actions** > **Save Config**.

NOTE
If you have made changes to a configuration, you must enable the configuration before the changes are effective. To enable the configuration, refer to "Enabling zone configurations" on page 134.

Selecting a zoning view

You can define how zoning elements are displayed in the **Zone Administration** window. The zoning view you select determines how members are displayed in the **Member Selection List** panel (Figure 37). The views filter the fabric and device information displayed in the **Member Selection List** for the selected view, making it easier for you to create and modify zones, especially when creating "hard zones."

Depending on the method you use to zone, certain tabs might or might not be available in the **Zone Administration** window.

There are two views of defining members for zoning:

- **Fabric View**—Displays the physical hierarchy of the fabric, a list of the attached and imported physical devices (by WWN), and a list of the FC Virtual Initiators on switches that support iSCSI. In the Fabric View, you can select ports for port-based zoning or devices for WWN-based zoning.

- **Devices Only**—Displays a list of the attached and imported physical devices by WWN. You cannot select ports for port-based or mixed zoning schemes, nor can you select virtual initiators for iSCSI FC Zone creation.

Use the following procedure to define the view of the fabric resource.

1. Launch the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select **View** > **Choose Fabric Resources View**.

3. Define the way you want to view the fabric resource and click **OK**.

Creating and populating zone aliases

An alias is a logical group of port index numbers and WWNs. Specifying groups of ports or devices as an alias makes zone configuration easier, by enabling you to configure zones using an alias rather than inputting a long string of individual members. You can specify members of an alias using the following methods:

- Identifying members by switch domain and port index number pair, for example, 2, 20.
- Identifying members by device node and device port WWNs.

For more information on enabling the configuration, refer to "Enabling zone configurations" on page 134.

Use the following procedure to create a zone alias.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select a format to display zoning members in the **Member Selection List** as described in "Selecting a zoning view" on page 125.

3. Select the **Alias** tab and click **New Alias**.

   The **Create New Alias** dialog box displays.

4. In the **Create New Alias** dialog box, enter a name for the new alias and click **OK**.

   The new alias displays in the **Name** list.

5. Expand the **Member Selection List** to view the nested elements.

   The choices available in the **Member Selection List** depend on the selection in the **View** menu.

6. Click elements in the **Member Selection List** that you want to include in the alias. The **Add Member** button becomes active.

7. Click **Add Member** to add alias members.

   Selected members move to the **Alias Members** window.

8. *Optional*: Repeat steps 6 and 7 to add more elements to the alias.

9. *Optional*: Click **Add Other** to include a WWN or port that is not currently a part of the fabric.

10. Select **Actions** > **Save Config** to save the configuration changes.

Adding and removing members of a zone alias

For more information on enabling the configuration, refer to "Enabling zone configurations" on page 134.

Use the following procedure to add or remove zone alias members.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select the **Alias** tab.

3. Select the alias you want to modify from the **Name** list.

4. Select an element in the **Member Selection List** that you want to add to the alias, or select an element in the **Alias Members** list that you want to remove.

5. Click **Add Member** to add the selected alias member, or click **Remove Member** to remove the selected alias member.

    The alias is modified in the Zone Admin buffer.

6. Select **Zoning Actions** > **Save Config** to save your configuration changes.

Renaming zone aliases

The new alias name cannot exceed 64 characters and can contain alphabetic, numeric, and underscore characters.

For more information on enabling the configuration, refer to "Enabling zone configurations" on page 134.

Use the following procedure to change the name of a zone alias.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select the **Alias** tab and select the alias you want to rename from the **Name** list.

3. Click **Rename**.

    The **Rename an Alias** dialog box displays.

4. Enter a new alias name and click **OK**.

    The alias is renamed in the Zone Admin buffer. At this point, you can either save your changes or save and enable your changes.

5. Select **Zoning Actions** > **Save Config** to save the configuration changes.

Deleting zone aliases

You can remove a zone alias from the Zone Admin buffer. When a zone alias is deleted, it is no longer a member of the zones of which it was once a member.

NOTE
If you delete the only member zone alias, an error message is issued when you attempt to save the configuration.

Use the following procedure to delete the zone aliases.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select the **Alias** tab.

3. Select the alias you want to delete from the **Name** list. and click **Delete**.

    The **Confirm Deleting Alias** dialog box displays.

4. Click **Yes**.

    The selected alias is deleted from the Zone Admin buffer. At this point, you can either save your changes or save and enable your changes.

5. Select **Zoning Action** > **Save Config** to save the configuration changes.

To enable the configuration, refer to "Enabling zone configurations" on page 134.

Creating and populating zones

A zone is a region within the fabric where specified switches and devices can communicate. A device can communicate only with other devices connected to the fabric within its specified zone.

Use the following procedure to create a zone.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select a format to display zoning members in the **Member Selection List** as described in "Selecting a zoning view" on page 125.

3. Select the **Zone** tab.

4. Click **New Zone**.

    The **Create New Zone** dialog box displays.

5. In the **Create New Zone** dialog box, enter a name for the new zone, and click **OK**.

    LSAN zones and QoS zones have specific naming requirements:

    - For LSAN zones, refer to "LSAN zone requirements" on page 120.
    - For QoS zones, refer to "QoS zone requirements" on page 120.

    The new zone displays in the **Name** list.

6. Expand the **Member Selection List** to view the nested elements. The choices available in the list depend on the selection made in the **View** menu.

7. Select an element in the **Member Selection List** that you want to include in your zone.

    Note that LSAN zones should contain only port WWN members. The **Add Member** button becomes active.

8. Click **Add Member** to add the zone member.

    The selected member is moved to the **Zone Members** window.

9. *Optional*: Repeat steps 7 and 8 to add more elements to your zone.

10. *Optional*: Click **Add Other** to include a WWN or port that is not currently a part of the fabric. At this point, you can either save your changes or save and enable your changes.

11. Select **Zoning Actions** > **Save Config** to save the configuration changes.

To enable the configuration, refer to "Enabling zone configurations" on page 134.

Adding and removing members of a zone

For information on enabling the configuration, refer to "Enabling zone configurations" on page 134.

Use the following procedure to add or remove zone members.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select the **Zone** tab.

3. Select the zone you want to modify from the **Name** list.

    The zone members for the selected zone are listed in the **Zone Members** list.

4. Highlight an element in the **Member Selection List** that you want to include in your zone, or highlight an element in the **Zone Members** list that you want to delete.

5. Click **Add Member** to add a zone member, or click **Remove Membe**r to remove a zone member. The zone is modified in the Zone Admin buffer.

6. Select **Zoning Actions** > **Save Config** to save the configuration changes.

Renaming zones

For information on enabling the configuration, refer to "Enabling zone configurations" on page 134.

Use the following procedure to change the name of a zone.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Click the **Zone** tab.

3. Select the zone you want to rename from the **Name** list.

4. Click **Rename**.

5. In the **Rename a Zone** dialog box, enter a new zone name and click **OK**. The zone is renamed in the Zone Admin buffer.

6. Select **Zoning Actions** > **Save Config** to save the configuration changes.

Cloning zones

To perform clone operations for zoning, the EGM license must be installed on the switch; otherwise, access to this feature is denied and an error message displays.

The EGM license is required only for 8 Gbps platforms, such as the following:

- Brocade DCX and DCX-4S enterprise-class platforms
- Encryption Switch
- 300, 5300, and 5100 switches

For non-8 Gbps platforms, all functionalities are available without EGM license.

Use the following procedure to clone a zone configuration.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Click the **Zone** tab.

3. Select the zone you want to clone from the **Name** list.

4. Click **Clone**

5. In the **Clone an Existing Zone** dialog box, enter a name for the copied zone.

6. Click **OK**. The selected zone is copied from the Zone Admin buffer.

7. Select **Zoning Actions** > **Save Config** to save the configuration changes. Because no changes were made to the effective configuration, you do not need to enable the configuration.

Deleting zones

For information on enabling the configuration, refer to *"Enabling zone configurations"* on page 134.

Use the following procedure to delete a zone.

1. Open the **Zone Administration** window as described in *"Opening the Zone Administration window"* on page 121.

2. Click the **Zone** tab.

3. Select the zone you want to delete from the **Name** menu and click **Delete**.

4. On the confirmation dialog box, click **Yes**.

   The selected zone is deleted from the Zone Admin buffer. At this point, you can either save your changes or save and enable your changes.

5. Select **Zoning Actions** > **Save Config** to save the configuration changes.

Creating and populating enhanced traffic isolation zones

An enhanced traffic isolation zone (TI zone) is a special zone that creates a dedicated path for a specific traffic flow. When an enhanced TI zone is activated, inter-switch traffic from a zone member is directed to E_Ports that are included in the TI zone. Traffic from outside the TI zone is excluded. A maximum of 255 TI zones can be configured. LSAN devices can be added only in TI zones created in the backbone switch.

A port may be assigned to more than one enhanced TI zone in a fabric. A port can be part of more than one enhanced TI zone provided following conditions are satisfied:

- All the switches in the fabric should have Fabric OS v 6.4 or later.
- A port can be assigned to multiple TI Zones that have the same failover state.
- The fabric is composed entirely of Condor-2 or GoldenEye-2 switches:

Use the following procedure to create and populate an enhanced TI zone.

1. Open the **Zone Administration** window as described in *"Opening the Zone Administration window"* on page 121.

2. Under **Zoning Modes**, select **Traffic Isolation Zones**.

   The **Traffic Isolation Zones** view displays.

3. Click **New** on the menu bar.

   The **Add TI Zone** dialog box displays.

4. Expand the **Member Selection List** to view the nested elements.

5. Select an element in the **Member Selection List** that you want to include in your zone.

   The **Add Member** button becomes active.

6. Click **Add Member** to add the zone member.

   The selected member is moved to the **Zone Members** window.

   NOTE
   All switches in the fabric must be running Fabric OS v6.4.0 or later and all the ports in the TI zones must be in the same failover mode.

7. *Optional*: Repeat steps 5 and 6 to add more elements to your TI zone.

8. When you are finished, click **OK**. The **Traffic Isolations Zones** window displays.

9. Click **Apply** to save the TI zone configuration.

Zone configuration and zoning database management

A zone configuration is a group of zones; zoning is enabled on a fabric by enabling a specific configuration. You can specify members of a configuration using zone names.

Figure 38 displays a sample zoning database and the relationship between the zone aliases, zones, and zoning configuration. The database contains one zoning configuration, *myconfig*, which contains two zones: *Zone A* and *Zone B*. The database also contains four aliases, which are members of Zone A and Zone B. Zone A and Zone B also have additional members other than the aliases.
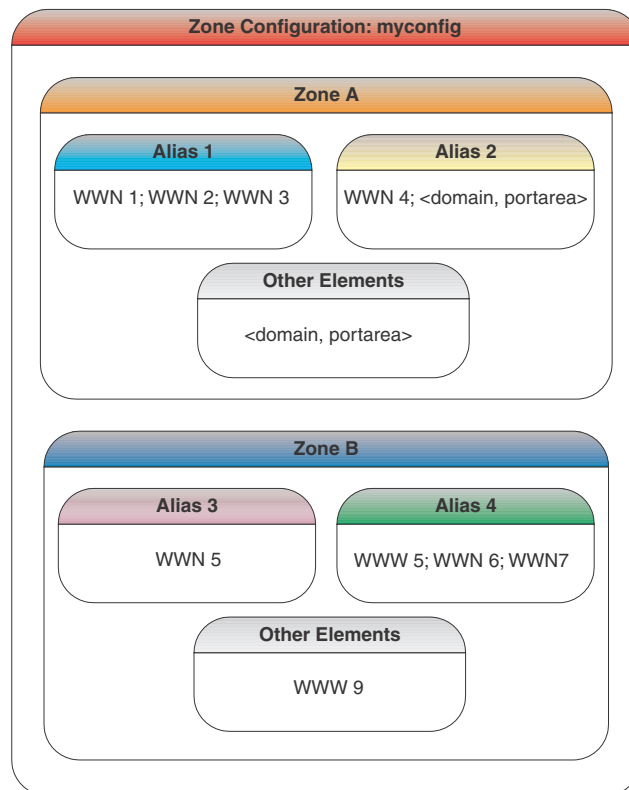
FIGURE 38   Sample zoning database

Creating zone configurations

Use the following procedure to create a zone configuration. After creating a zone configuration, you must explicitly enable it for it to take effect.

For information on enabling the configuration, refer to "Enabling zone configurations" on page 134.

NOTE
Any changes made to the currently enabled configuration does not display until you re-enable the configuration.

Use the following procedure to create zone configurations.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select a format to display zoning members in the **Member Selection List** as described in "Selecting a zoning view" on page 125.

3. Select the **Zone Config** tab and click **New Zone Config**.

4. In the **Create New Config** dialog box, enter a name for the new configuration and click **OK**.

   The new configuration displays in the **Name** list.

5. Expand the **Member Selection List** to view the nested elements.

   The choices available in the list depend on the selection made in the **View** menu.

6. Select an element in the **Member Selection List** that you want to include in your configuration.

   The **Add Member** button becomes active.

7. Click **Add Member** to add configuration members.

   Selected members are moved to the **Config Members** window.

8. Repeat steps 6 and 7 to add more elements to your configuration.

9. Select **Zoning Actions** > **Save Config** to save the configuration changes.

Adding or removing zone configuration members

For information on enabling the configuration, refer to "Enabling zone configurations" on page 134.

Use the following procedure to add or remove members of a zone configuration.

NOTE
You can make changes to a configuration that is currently enabled; however, changes do not display until you re-enable the configuration.

Use the following procedure to configure the zone members.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select the **Zone Config** tab.

3. Select the configuration you want to modify from the **Name** list.

4. Click an element in the **Member Selection** list that you want to include in your configuration or select the element in the **Config Members** list that you want to delete.

5. Click **Add Member** to add a configuration member or **Remove Member** to remove a configuration member.

6. Select **Zoning Actions** > **Save Config** to save the configuration changes.

Renaming zone configurations

The new name cannot exceed 64 characters and can contain alphabetic, numeric, and underscore characters. Use the following procedure to change the name of a zone configuration.

NOTE
You cannot rename the currently enabled configuration.

Use the following procedure to rename the zone configuration.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select the **Zone Config** tab.

3. Select the configuration you want to rename from the **Name** list and click **Rename**.

4. In the **Rename a Config** dialog box, enter a new configuration name and click **OK**.

    The configuration is renamed in the configuration database.

5. Select **Zoning Actions** > **Save Config** to save the configuration changes.

Cloning zone configurations

You must use Web Tools with the EGM license to perform cloning operations for zone configurations; otherwise, access to this feature is denied and an error message displays.

Use the following procedure to clone a zone configuration.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select the **Zone Config** tab.

3. Select the zone configuration you want to clone from the **Name** list.

4. Click **Clone**.

5. In the **Copy An Existing Zone Config** dialog box, enter a name for the copied zone and click **OK**.

    The selected zone is copied from the Zone Admin buffer.

6. Select **Zoning Actions** > **Save Config** to save the configuration changes.

    No changes were made to the effective configuration. You do not need to enable the configuration.

Deleting zone configurations

Use the following procedure to delete a zone configuration.

NOTE
You cannot delete a enabled configuration.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select the **Zone Config** tab.

3. Select the configuration you want to delete from the **Name** list and click **Delete**.

4. On the confirmation dialog box, click **Yes**. The selected configuration is deleted from the configuration database.

5. Select **Zoning Actions** > **Save Config** to save the configuration changes.

Enabling zone configurations

Several zone configurations can reside on a switch at the same time, and you can quickly alternate between them. For example, you might want to have one configuration enabled during the business hours and another enabled overnight. However, only one zone configuration can be enabled at a time.

When you enable a zone configuration from Web Tools, the entire zoning database is automatically saved, and then the selected zone configuration is enabled.

If the zoning database size exceeds the maximum allowed, you cannot enable the zone configuration. The zoning database summary displays the maximum zoning database size.

Use the following procedure to enable the zone configuration.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select **Zoning Actions** > **Enable Config**.

3. On **Enable Config**, select the configuration to be enabled from the menu.

4. Click **OK** to save and enable the selected configuration.

Disabling zone configurations

When you disable the active configuration, the Advanced Zoning feature is disabled on the fabric, and all devices within the fabric can communicate with all other devices. This does not mean that the zoning database is deleted, however, only that there is no configuration active on the fabric.

When you disable a zone configuration from Web Tools, keep in mind that the entire zoning database is automatically saved, and then the selected zone configuration is disabled.

NOTE
When you disable the active configuration, Advanced Zoning is disabled on the fabric, and according to the default zone set, devices within the fabric can or cannot communicate with other devices.

Use the following procedure to disable a zone configuration.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select **Zoning Actions** > **Disable Zoning**.

    The **Disable Config** warning message displays.

3. Click **Yes** to save and disable the current configuration.

Displaying enabled zone configurations

The enabled zone configuration screen displays the actual content of the single zone configuration that is currently enabled on the fabric, whether it matches the configuration that was enabled when the current Zone Administration session was launched or last refreshed. The zones are displayed, and their contents (ports, WWNs) are displayed next to them. Aliases are not displayed in the enabled zone configuration. If there is no active zone configuration enabled on the switch, a message displays to that effect.

NOTE
The enabled configuration is listed in the lower-right corner of the **Zone Administration** window.

Viewing the enabled zone configuration name without opening the Zone Administration window

Use the following procedure to view the enabled zone configuration name:

1. Select a logical switch using the drop-down list under **Fabric Tree** section in the Switch Explorer window.

   The selected switch displays in the Switch View.

2. You can view the current zone configuration name (if one is enabled) in the lower portion of the Switch Events and Switch Information window.

   If no zone configuration is enabled, the field displays "No configuration in effect".

Viewing detailed information about the enabled zone configuration

Use the following procedure to view detailed information about the enabled zone configuration.

1. Open the **Zone Administration** window, as described on

   The zone configuration in effect at the time you launched the **Zone Administration** window is identified in the lower-right corner. It is also updated if you manually refresh the **Zone Administration** window contents by clicking the refresh icon at the lower-right corner of the **Zone Administration** window, or when you enable a configuration through the **Zone Administration** window.

⚠️ CAUTION

Clicking the refresh icon overwrites all local unsaved zoning changes. If anyone has made any changes to the zones outside of your Zone Admin session, those changes are applied.

2. To identify the most recently effective zone configuration *without* saving or applying any changes you made in the **Zone Administration** window, select **Print > Print Effective Zone Configuration** in the **Zone Administration** window.

> **NOTE**
> If no zone is enabled, a message displays, indicating that there is no active zoning configuration on the switch.

3. *Optional*: Click **Print** located in the **Print Effective Zone Configuration** dialog box to print the enabled zone configuration details.

> **NOTE**
> You must use DCFM Professional or Enterprise Edition to print the zone database summary configurations, display zone configuration summaries and create configuration analysis reports.

Adding a WWN to multiple aliases and zones

This procedure enables you to configure a WWN as a member in a zone configuration prior to adding that device to the fabric. Specifically, it is useful if you want to add a WWN to all or most zoning entities. The added WWN does not need to currently exist in the fabric.

Use the following procedure to add a WWN.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.
2. Select **Edit** > **Add WWN**.

   The **Add WWN** dialog box displays.
3. Enter a WWN value in the **WWN** field and click **OK**.

   The **Add WWN** dialog box displays all the zoning elements that include the new WWNs. All of the elements are selected by default.
4. Click items in the list to select or unselect, and click **Add** to add the new WWN to all the selected zoning elements.

   The WWN is added to the Zone Admin buffer and can be used as a member.

Removing a WWN from multiple aliases and zones

Use this procedure if you want to remove a WWN from all or most zoning entities.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.
2. Select **Edit** > **Delete WWN**.

   The **Delete WWN** dialog box displays.
3. Enter a WWN value in the **WWN** field and click **OK**.

   The **Delete WWN** dialog box displays all the zoning elements that include the WWN.
4. Click items in the list to select or unselect, and click **Delete** to delete the WWN from all the selected zoning elements.

   The WWN is deleted from the selected items in the Zone Admin buffer.

Replacing a WWN in multiple aliases and zones

This procedure enables you to replace a WWN throughout the Zone Admin buffer. This is helpful when exchanging devices in your fabric and helps you to maintain your current configuration.

Use the following procedure to replace a WWN in multiple aliases and zones.

1. Launch the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select **Edit** > **Replace WWN**.

    The **Replace WWN** dialog box displays.

3. Enter the WWN to be replaced in the **Replace** field.

4. Enter the new WWN in the **By** field and click **OK**.

    The **Replace WWN** dialog box displays. It lists all the zoning elements that include the WWN.

5. Click an item in the list to select or unselect, and click **Replace** to replace the WWN in all the selected zoning elements.

    The former WWN is replaced in the Zone Admin buffer by the new WWN, including within any alias or zone in which the old WWN was a member.

Searching for zone members

You can search zone member selection lists for specified strings of text. If you know some identifying information about a possible member of a zoning entity, you can select the tab and view for that entity and then search through its member selection list using the **Search for Zone Member** option. If the target entity is an alias or zone, then the search domain includes elements like switch names and domain numbers, port names and "domain, port" addresses, device WWNs and manufacturer names, and also any aliases that might already have been defined. If the target entity is a configuration, then zones are also included, along with the elements they contain.

The search starts from the top of the list, and when the target element is found, it is also selected in the **Member Selection List** so it can be added or its parent or children can be found. By default, the **Member Selection List** is searched from beginning to end one time. If you select the wraparound option, the search continues to loop from the beginning to the end of the **Member Selection List**.

Use the following procedure to search for zone members.

1. Open the **Zone Administration** window as described in "Opening the Zone Administration window" on page 121.

2. Select **Edit** > **Search Member**.

3. Enter the zone member name in the **Member Name** field.

    Optional: Narrow the search by selecting one or more of the check boxes, such as **Match Case**.

4. Click **Next** to begin the zone member search.

Clearing the zoning database

Use the following procedure to disable the active zoning configuration, if one exists, and delete the entire zoning database. You must disable any active configuration before you can delete the zoning database.

ATTENTION
This action not only disables zoning on the fabric, but also deletes the entire zoning database. This results in all devices being able to communicate with each other.

Use the following procedure to clear the zone database.

1. Open the **Zone Administration** window as described in

2. Select **Actions** > **Clear All**.

   The **Disable Config** wizard displays.

3. Click **Yes** to do *all* of the following in the wizard:

   - Disable the current configuration.
   - Clear the entire contents of the current Web Tools Zone Admin buffer.
   - Delete the entire persistent contents of the fabric zoning database.

   The wizard allows you to define one and only one name for each device port (WWN). Devices with one or more aliases are considered already named and are not displayed.

Zone configuration analysis

You must use DCFM Professional or Enterprise Edition to analyze the following zone configurations:

- Add unzoned devices
- Remove offline or inaccessible devices
- Replace offline devices
- Define device alias

Best practices for zoning

The following are recommendations for using zoning:

- Always zone using the highest Fabric OS-level switch.

  Switches with lower Fabric OS versions do not have the capability to view all the functionality that a newer Fabric OS provides as functionality is backwards compatible but not forwards compatible.

- Zone using the core switch versus an edge switch.

- Zone using a director over a switch.

  A director has more resources to handle zoning changes and implementations.

- Zone on the switch you connect to when bringing up Web Tools (the proxy switch).

Working with Diagnostic Features

In this chapter

Trace dumps

A trace dump is a snapshot of the running behavior within the Brocade switch. The dump can be used by developers and troubleshooters at Brocade to help understand what might be contributing to a specific switch behavior when certain internal events are seen. For example, a trace dump can be created each time a certain error message is logged to the system error log. Developers can then examine what led up to the message event by studying the traces.

Tracing is always "on." As software on the switch executes, the trace information is placed into a circular buffer in system RAM. Periodically, the trace buffer is "frozen" and saved. This saved information is a "trace dump."

A trace dump is generated when:

- It is triggered manually (use the **traceDump** command).
- A critical-level LOG message occurs.
- A particular LOG message occurs (use the **traceTrig** command to set up the conditions for this).
- A kernel panic occurs.
- The hardware watchdog timer expires.

(For information about the **traceDump** and **traceTrig** commands, refer to the *Fabric OS Command Reference*.)

The trace dump is maintained on the switch until either it is uploaded to the FTP host or another trace dump is generated. If another trace dump is generated before the previous one is uploaded, the previous dump is overwritten.

When a trace dump is generated, it is automatically uploaded to an FTP host if automatic FTP uploading is enabled.

Using the **Trace** tab of the **Switch Administration** window, you can view and configure the trace FTP host target and enable or disable automatic trace uploads.

How a trace dump is used

The generation of a trace dump causes a CRITICAL message to be logged to the system error log. When a trace dump is detected, issue the **supportSave** command on the affected switch. This command packages all error logs, the **supportShow** output, and trace dump, and moves these to your FTP server. You can also configure your switch to automatically copy trace dumps to your FTP server (refer to "Setting up automatic trace dump transfers").

In addition to automatic generation of trace dumps on faults, you can also generate a trace dump manually or when certain system error messages are logged. This is normally done with assistance from Brocade customer support when diagnosing switch behavior.

For details on the commands, refer to the *Fabric OS Command Reference*.

Setting up automatic trace dump transfers

You can set up a switch so that diagnostic information is transferred automatically to a remote server. Then, if a problem occurs you can provide your customer support representative with the most detailed information possible. To ensure the best service, you should set up for automatic transfer as part of standard switch configuration, before a problem occurs.

Setting up for automatic transfer of diagnostic files involves the following tasks:

- Specifying a remote server to store the files.
- Enabling the automatic transfer of trace dumps to the server. (Trace dumps overwrite each other by default; sending them to a server preserves information that would otherwise be lost.)

Specifying a remote server

The switch must belong to your current Admin Domain before you can perform this task.

Use the following procedure to specify a remote server.

1. Open the **Switch Administration** window.

2. Click **Show Advanced Mode**, if it is not selected.

3. Select the **Trace** tab.

4. Enter the FTP host IP address, path of the remote directory for the trace dump files, FTP user name, and FTP password in the appropriate fields.

   The IP address can be IPv4 or IPv6 format, or a DNS name.

   The password is optional if you log in as an anonymous user.

5. Click **Apply**.

Enabling automatic transfer of trace dumps

The switch must belong to your current Admin Domain before you can perform this task.

Use the following procedure to enable the automatic transfer of trace dumps.

1. Open the **Switch Administration** window.

2. Click **Show Advanced Mode**, if it is not selected.

3. Select the **Trace** tab.

4. Select **Enable** in the **Auto FTP Upload** section to enable automatic uploading of the trace dump to the FTP host.

5. Click **Apply.**

Disabling automatic trace uploads

If automatic uploading of a trace dump is disabled, you must manually upload the trace dump or else the information is overwritten when a subsequent trace dump is generated.

The switch must belong to your current Admin Domain before you can perform this task.

Use the following procedure to disable automatic trace uploads.

1. Open the **Switch Administration** window.

2. Click **Show Advanced Mode**, if it is not selected.

3. Select the **Trace** tab.

4. Select **Disable** in the **Auto FTP Upload** section to disable automatic uploading of the trace dump to the FTP host.

5. Click **Apply.**

Displaying switch information

You can right-click in the table content of **Fan**, **Temperature**, and **Power Status** windows to find Export, Copy, and Search options. These options are not available if the table does not have any content.
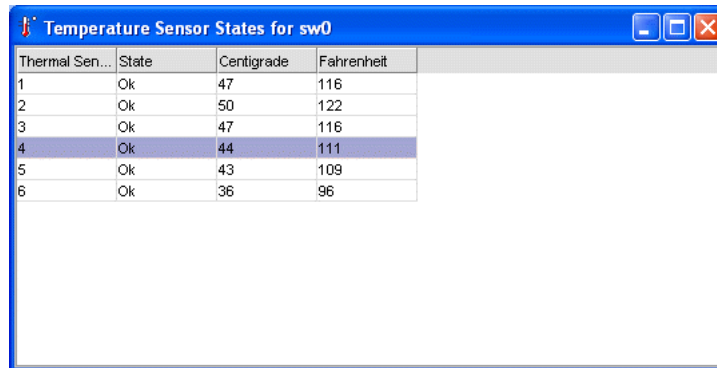
*   Click **Export Row** or **Export Table** to save the contents to a tab-delimited file.

*   Click **Copy Row** or **Copy Table** to copy the contents in tab-delimited text format to a file.

*   Click **Search** to search for a specific text string in the table.

NOTE
You must accept the Brocade Certificate at the beginning of the login to Web Tools to enable the functionality of Export and Copy.

Enter the text string in the box that displays on the table, as shown in Figure 39, and press **Enter**. This is an incremental search and allows 24 maximum characters including wildcards question mark (?) and asterisk (*). The first row containing the text string is highlighted. To find the next match, click the down arrow. To find the previous match, click the up arrow. If the text is not found in the table, the text turns red.
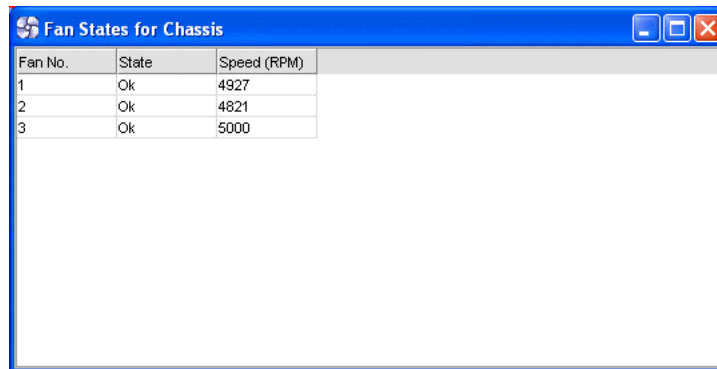
FIGURE 39   Temperature Sensor States window

Viewing detailed fan hardware status

The icon on the **Fan** button indicates the overall status of the fans. For more information about the switch fan, refer to the appropriate hardware documentation.

You can display status information about the fans, as shown in Figure 40.

FIGURE 40   Fan States window

The **Fan No.** column indicates either the fan number or the fan FRU number, depending on the switch model. A fan FRU can contain one or more fans.

- For Brocade 4100, 4900, 5000, 5100, 5300, 7600, the 7500, 7500E, and 7800 Extension switches, the 8000, and the Brocade Encryption Switch, the **Fan No.** column indicates the fan FRU number.

- The Brocade 200E, 300, 4012, 4016, 4018, 4020, and 4024 switches do not contain fan FRUs, so for these switch models, the **Fan No.** column indicates the fan number.

NOTE
For these switches, if the Fan States window has no "Fan Speed" column, *the speed is not monitored*.

Use the following steps to view the detailed fan status of a switch.

1. Select a logical switch using the drop-down list under **Fabric Tree** section in the Switch Explorer window.

   The selected switch displays in the Switch View. The icon on the **Fan** button indicates the overall status of the fan.

2. Click **Fan**.

   The detailed fan status for the switch displays, as shown in Figure 40.

Viewing the temperature status

The icon on the **Temp** button indicates the overall status of the temperature. For more information regarding switch temperature, refer to the appropriate hardware documentation.

Use the following procedure to view the temperature status.

1. Select a logical switch using the drop-down list under **Fabric Tree** section in the Switch Explorer window.

   The selected switch displays in the Switch View.. The icon on the **Temp** button indicates the overall status of the temperature.

2. Click **Temp** on the Switch View.

   The detailed temperature sensor states for the switch are displayed, as shown in Figure 39.

Viewing the power supply status

The icon on the **Power** button indicates the overall status of the power supply status. For more information regarding switch power modules, refer to the appropriate hardware documentation.

Use the following procedure to view the power supply status.

1. Select a logical switch using the drop-down list under **Fabric Tree** section in the **Switch Explorer** window.

1. The selected switch displays in the Switch View. The icon on the **Power** button indicates the overall status of the power supply.

2. Click **Power** on the Switch View. The detailed power supply states are displayed (Figure 41).
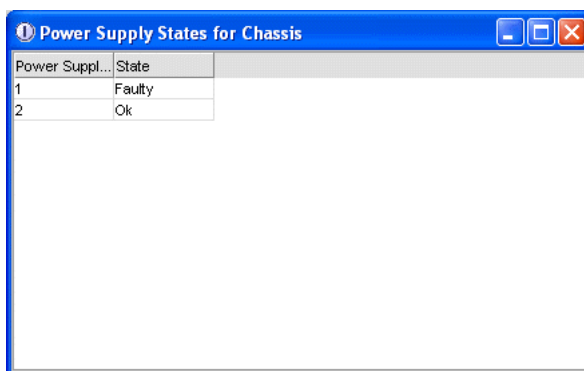
FIGURE 41   Power States window

Checking the physical health of a switch

The **Status** button displays the operational state of the switch. The icon on the button displays the real-time status of the switch.

If no data is available from a switch, the most recent background color remains displayed.

Any error-based status messages that is based on a per time interval cause the status to show faulty until the entire sample interval has passed.

If the switch status is marginal or critical, information on the trigger that caused that status displays in the **Switch Information** view.

Click the **Status** button to display a detailed, customizable switch status report, shown in Figure 42. Note that this is a static report and not a dynamic view of the switch.
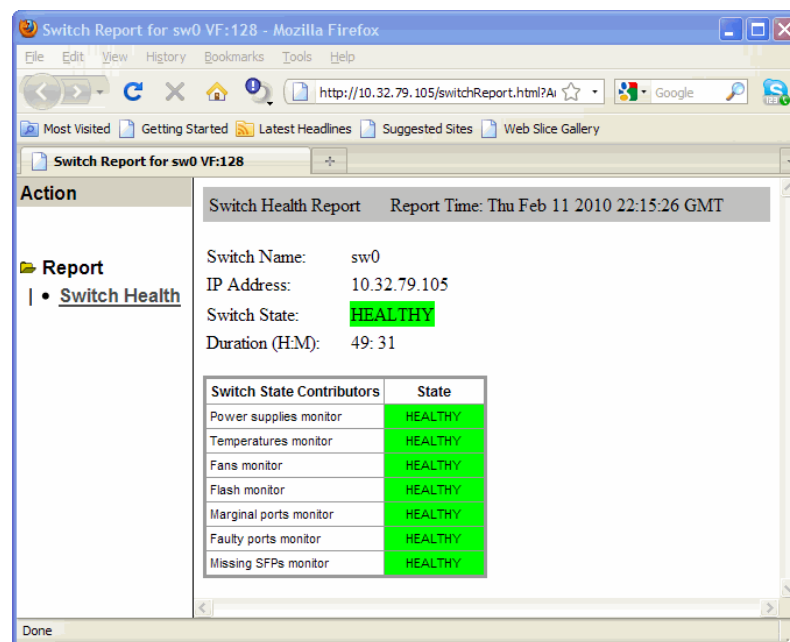
FIGURE 42   Switch Report window

Use the following procedure to check the physical health of the switch.

1. Select a logical switch using the drop-down list under **Fabric Tree** section in the Switch Explorer window.

    The selected switch displays in the Switch View. The icon on the **Status** button indicates the overall status of the switch.

2. Click **Status** on the Switch View.

    The detailed switch health report displays, as shown in Figure 42.

3. *Optional*: Click the underlined links in the left panel to display detailed information about ports and Switch Availability Monitoring (SAM).

> **NOTE**
> The Port Detail Report and Switch Availability Monitor (SAM) reports display the details of only those ports which are members of the current Admin Domain context and the E_Ports of the switch.

4. *Optional*: Hold the cursor on the **Action** bar and click an action to perform one of the following options:

    - Refresh the information displayed in the report
    - Customize the report
    - View the data in raw XML format
    - View the style sheet for the report
    - View the XML schema for the report

Defining Switch Policy

The Switch Policy dialog box lets you define the values for what you consider a healthy switch. The parameters for Switch Policy define whether the unit is listed as being "Healthy", "Marginal", or "Down".

Use this dialog box to set policy parameters for calculating the overall status of the switch. 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". The existence of policies such as WWN, CP, and Blade might differ from platform to platform.

Any single contributor can force the overall status of the switch to MARGINAL or DOWN. For example, assuming that the switch contributor values are set to the default values, if there is one faulty port in a switch, then this contributor would set the overall switch status to MARGINAL. If two ports were faulty, then this contributor would set the overall switch status to DOWN.

Note that the value, 0, for a parameter, means that it is NOT used in the calculation. In addition, if the range of configurable values in the prompt is (0..0), the policy parameter is NOT applicable to the switch.

Use the following procedure to define the Switch Status Policy.

1. Open the Web Tools main page.

2. Click the **Switch Status Policy** button.

    The **Switch Status Policy** dialog box displays, as shown in Figure 43.

NOTE
The options available in the dialog box may differ, depending on the options available on your switch, including CP, core blades, blades, and WWN.
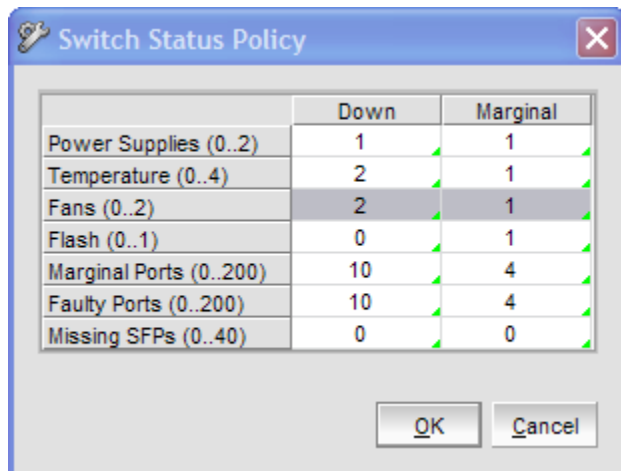
FIGURE 43   Switch Status Policy dialog box

3. Configure the numerical values to conform to your definition of a healthy switch.

   Numerical values that are above "Marginal" are considered to be "Healthy."

4. *Optional:* Right-click a row in the table to access options to copy the values to your clipboard, or to export the values to a file.

5. Click **OK**.

Port LED interpretation

The Switch View displays port graphics with blinking LEDs, simulating the physical appearance of the ports. One of the LEDs indicates port status; the other indicates port speed. For LED information, refer to the hardware documentation for the switch you are viewing. (The blink rate of the LEDs in the Switch View does not necessarily match the blink rate of the LEDs on the physical switch.)

NOTE
All 8G platform Brocade switches and port blades do not have port speed LEDs, but only port status LEDs.

Port icon colors

The background color of the port icon indicates the port status, as follows:

- Green (healthy)
- Yellow (marginal)
- Red (critical)
- Gray (unmonitored)

- Blue (buffer-limited)
- Dimmed (unlicensed)

LED representations

The port icons are different for different switch models. Figure 44 displays E_Port port icons and associated LEDs from a Brocade 4100 switch.

For the Brocade 4100, the top row of LEDs corresponds to the upper port, and the bottom row of LEDs corresponds to the lower port.
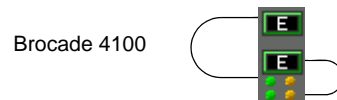
Brocade 4100

FIGURE 44   **Port and LED status color-coded display in the Port icon in Switch View**

Brocade 48000 Director LEDs

For the Brocade 48000 director, the representation of the port LEDs on the FC4-32 port blade is not the same as the LEDs on the physical blade. Figure 45 on page 148 compares the LEDs on the physical port card and the Web Tools display.
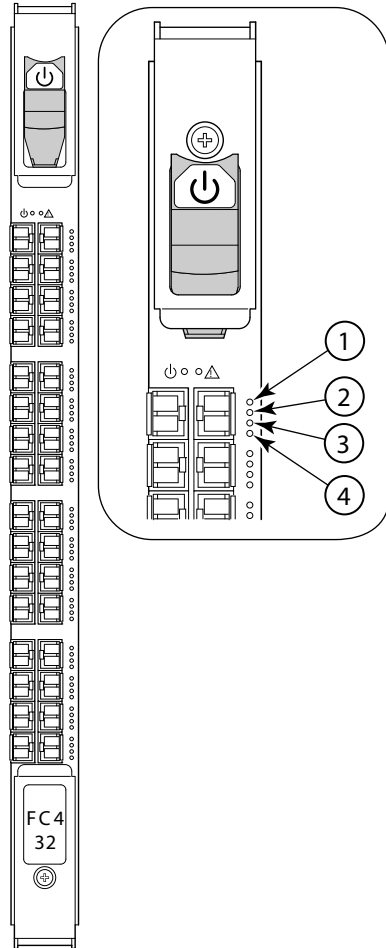
Physical Port Card Web Tools Representation
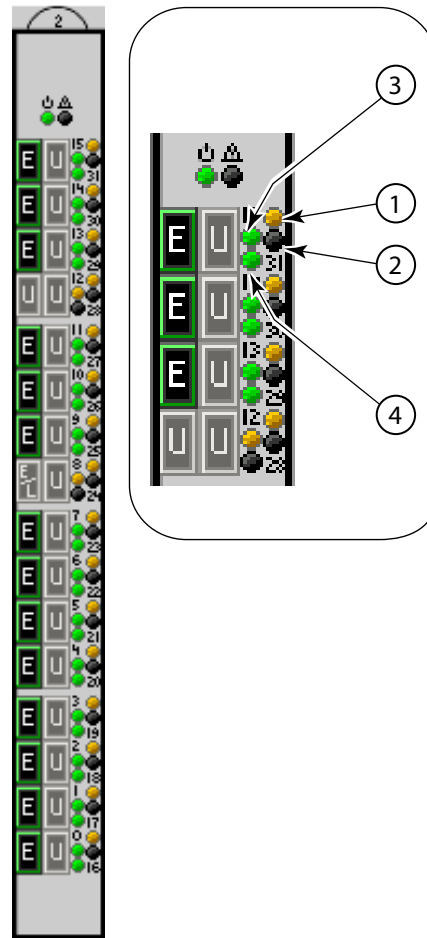
1. Port Speed LED for the right port
2. Port Status LED for the right port
3. Port Speed LED for the left port
4. Port Status LED for the left port

FIGURE 45   Port LEDs for the FC4-32 port blade in the Brocade 48000 Director

Using the FC-FC Routing Service

In this chapter

Fibre Channel Routing overview

Fibre Channel Routing provides connectivity to devices in different fabrics without merging the fabrics.

For example, Fibre Channel Routing allows you to share tape drives across multiple fabrics without the administrative problems, such as change management, network management, scalability, reliability, availability, and serviceability that might result from merging the fabrics.

Fibre Channel Routing lets you create logical storage area networks (LSANs) that can span fabrics. These LSANs allow Fibre Channel zones to cross physical SAN boundaries without merging the fabrics while maintaining the access controls of zones.

Note the following terminology for Fibre Channel Routing:

backbone fabric	An FC Router can connect two edge fabrics; a *backbone fabric* connects FC Routers. The FC Router fabric is the backbone fabric. A backbone fabric consists of at least one FC Router and possibly a number of Fabric OS-based Fibre Channel switches. Initiators and targets in the edge fabric can communicate with devices in the backbone fabric through the FC Router.
edge fabric	A standard Fibre Channel fabric with targets and initiators connected through an FC Router to another Fibre Channel fabric.
EX_Port	A type of port that functions somewhat like an E_Port, but does not propagate fabric services or routing topology information from one fabric to another.
FC Router	A switch running FC-FC Routing Service.
interfabric link (IFL)	The link between an E_Port and an EX_Port, or a VE_Port and a VEX_Port.
metaSAN	The collection of all SANs interconnected with FC Routers.

> **VEX_Port**   A virtual port that enables routing functionality through an FCIP tunnel. A
>   VEX_Port is similar to an EX_Port.

A device is shared between:

- The backbone fabric and edge fabric 1
- Edge fabric 1 and edge fabric 2
- Edge fabric 2 and edge fabric 3

Supported switches for Fibre Channel Routing

The FC-FC Routing Service is supported only on the following switch models:

- Brocade 5100 and 5300 switches, and the 7500, 7500E and 7800 Extension Switches.
- Brocade 48000 director, when configured with an FR4-18i blade.
- Brocade DCX and DCX-4S enterprise-class platforms, when configured with FR4-18i, FC8-16, FC8-32, FC8-48, or FX8-24 blades.

Setting up FC-FC routing

The following procedure provides the basic steps for setting up FC-FC Routing using an FC Router.

1. Ensure that the backbone fabric ID of the FC Router is the same as that of other FC Routers in the backbone fabric. Refer to "Configuring the backbone fabric ID" on page 156.

2. On the FC Router, ensure that the ports to be configured as EX_Ports are either not connected or are disabled.

3. Configure EX_Ports by clicking the **EX Ports** tab and then clicking **New**.

   Follow the instructions in the wizard. Refer to "Viewing EX_Ports" on page 153.

4. Connect the cables from the EX_Ports on the FC Router to the edge fabrics, if they were not connected before.

   NOTE
   For a multi-FC Router backbone fabric, make sure that each FC Router is connected to a switch in the backbone fabric.

5. Configure LSAN zones on the fabrics that share devices.

   Refer to "Viewing LSAN zones" on page 156.

6. View the information in the **EX Ports**, **LSAN Fabrics**, **LSAN Zones**, and **LSAN Devices** tabs to make sure that your configuration succeeded.

FC-FC routing management

You can perform Fibre Channel Routing operations using Web Tools, Web Tools with the EGM license, and Integrated Routing license. You can manage FC-FC Routing through the FC Routing module, shown in Figure 46 on page 152. The FC Routing module has tabbed panes that display EX_Ports, LSAN fabrics, LSAN zones, LSAN devices, and general FCR information.

The FC Routing module provides a dynamic display. Any changes in the FCR configuration on the switch are automatically updated in the FC Routing module within 30 to 90 seconds, depending on the network traffic.

The switch must be FC Router-capable, as described in "Fibre Channel Routing overview" on page 149.

The only things you need to configure on the FC Router are the EX_Ports and the backbone fabric ID. You configure LSAN zones on the fabrics from where devices need to be shared. You can configure LSAN zones on the backbone fabric to allow edge fabrics to share devices in the backbone fabric.

You can log in with any role and launch the FC routing module. To modify the data, you must log in as switchadmin, fabricadmin, basicswitchadmin, or operator. If you log in as user, zoneadmin, or securityadmin, you can only view the data.

If the FC-FC Routing service is disabled, the LSAN zones, LSAN fabric, and devices tabs continue to display the existing entries, but display the entries related to the *backbone fabric* only. All of the EX_Ports are disabled and you cannot enable them until FC-FC routing is enabled.

Opening the FC Routing module

The **FCR** button in the **Switch View** launches the FC Routing module. This button is displayed only for the following switches:

- Brocade 5100 and 5300 switches, 7500, 7500E, and 7800 Extension Switches.
- Brocade 48000 director, when configured with an FR4-18i blade.
- Brocade DCX and DCX-4S enterprise-class platforms, when configured with FR4-18i, FC8-16, FC8-32, FC8-48, or FX8-24 blades.

NOTE
When the Virtual Fabrics capability is enabled on the switch, Fabric ID cannot be set using the **Set Fabric ID** button.

Use the following procedure to open the FC Routing module.

1. Select a logical switch using the drop-down list under **Fabric Tree** section in the Switch Explorer window.

   The selected switch displays in the Switch View.

2. Click **FCR** in the **Manage** section of the **Tasks** menu.

The FC Routing module displays (as shown in Figure 46). If FC-FC Routing is disabled, a message to that effect displays on all the tabs in the module.
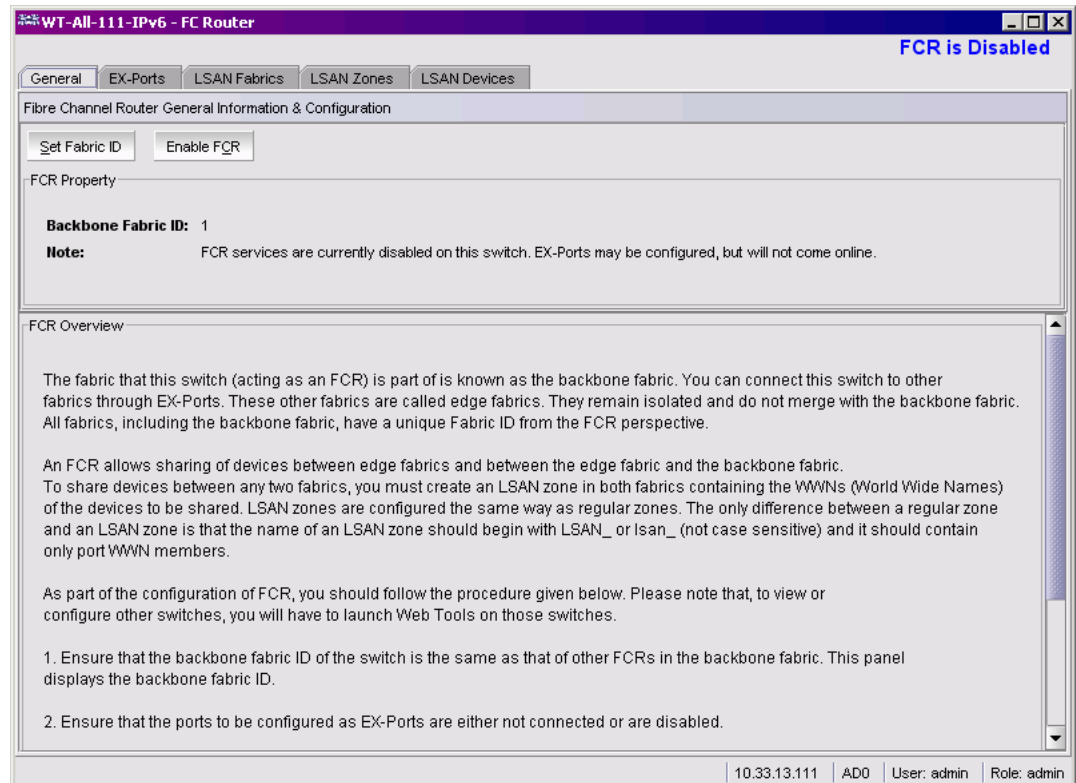
FIGURE 46   FC Routing module in Disabled mode with General tab selected

Viewing and managing LSAN fabrics

The **LSAN Fabric** tab (Figure 47 on page 153) displays all the LSAN fabrics visible to your switch, in both a tabular and tree form. (If FC-FC Routing is disabled, the table and tree nodes in this tab are empty and the tree displays only the backbone switch.)

For more detailed information about a specific LSAN fabric, click a fabric name in the table and then click **View Details** in the task bar. You can also click the fabric name in the tree on the left side of the window.

When there is more than one router present in the backbone fabric with different backbone Fabric IDs, the routers with the conflicting IDs are shown in a separate table on the LSAN Fabric tab.

To manage an LSAN fabric, select the fabric to manage and click **Manage LSAN Fabric** in the task bar. A browser window is launched with the following URL:

    http://ip-address-of-lsan-fabric-switch

For Brocade switches, this launches Web Tools. For non-Brocade fabrics, this launches the Element Manager for that switch.
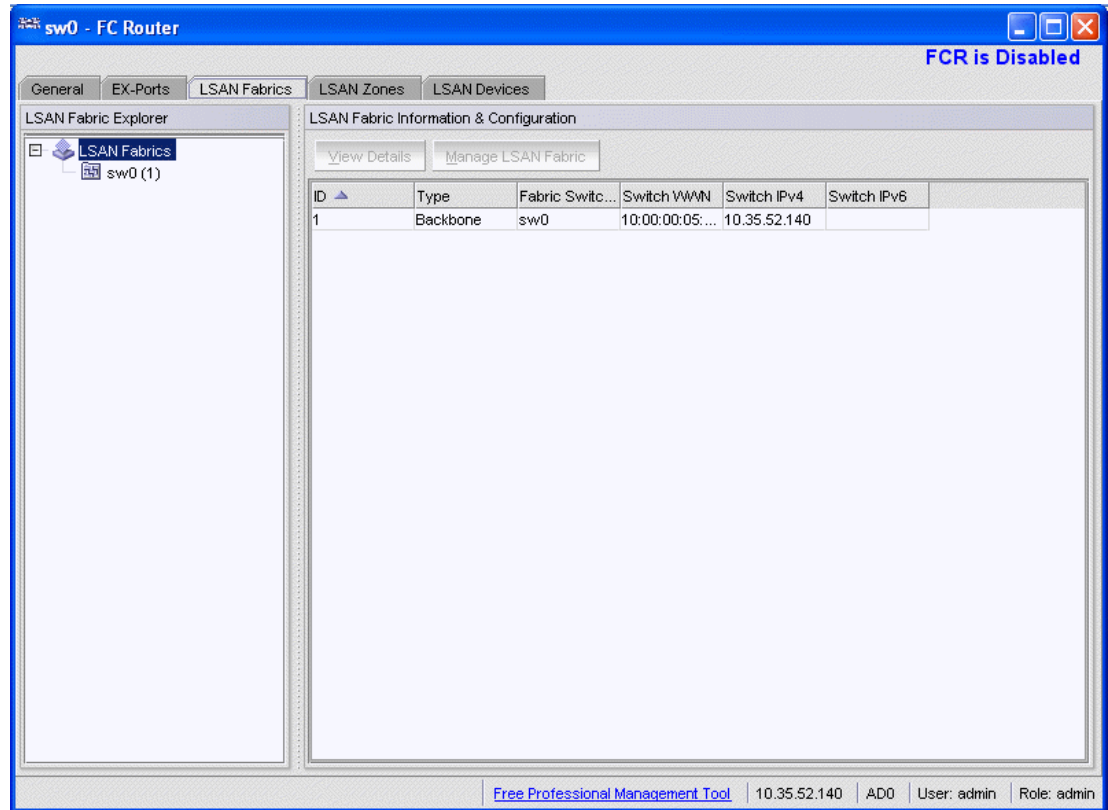
FIGURE 47   FC Routing module with LSAN Fabrics tab selected

Viewing EX_Ports

The **EX_Ports** tab (Figure 48 on page 154) displays all of the EX_Ports on the switch, including configuration and status information. The ports are sorted by slot number, and then by row number within each slot. IP address information is displayed in IPv4 and IPv6 formats.

NOTE
If FC Routing is disabled you must disable all of the EX_Ports. You cannot enable them until FC Routing is enabled.

For more detailed information about a specific port, click a port name in the table, and then click **View Details** in the task bar. You can also click the port name in the tree on the left side of the window.

From the **EX_Ports** tab, you can perform the following port management tasks by selecting a port in the table, and then clicking a task in the task bar:

- Configure EX_Ports
- Edit an EX_Ports configuration
- Rename an EX_Port
- Swap the Port Index of an EX_Port (described in "Port swapping index" on page 96)
- Enable or disable an EX_Port
- Persistently enable or disable an EX_Port
- Enable or disable trunking
- Configure router port cost

ATTENTION
During EX_Port configuration, the port is automatically disabled, and then re-enabled when the changes are applied. Be sure that you do not physically connect a port to a remote fabric before configuring it as an EX_Port; otherwise, the two fabrics merge and you lose the benefit of Fibre Channel Routing.

You can enable or disable multiple ports at one time. Use Shift-click and Ctrl-click to select multiple ports in the table, and then click one of the enable or disable tasks in the task bar.

You can select multiple ports in the table, but you can select only one port at a time in the tree.
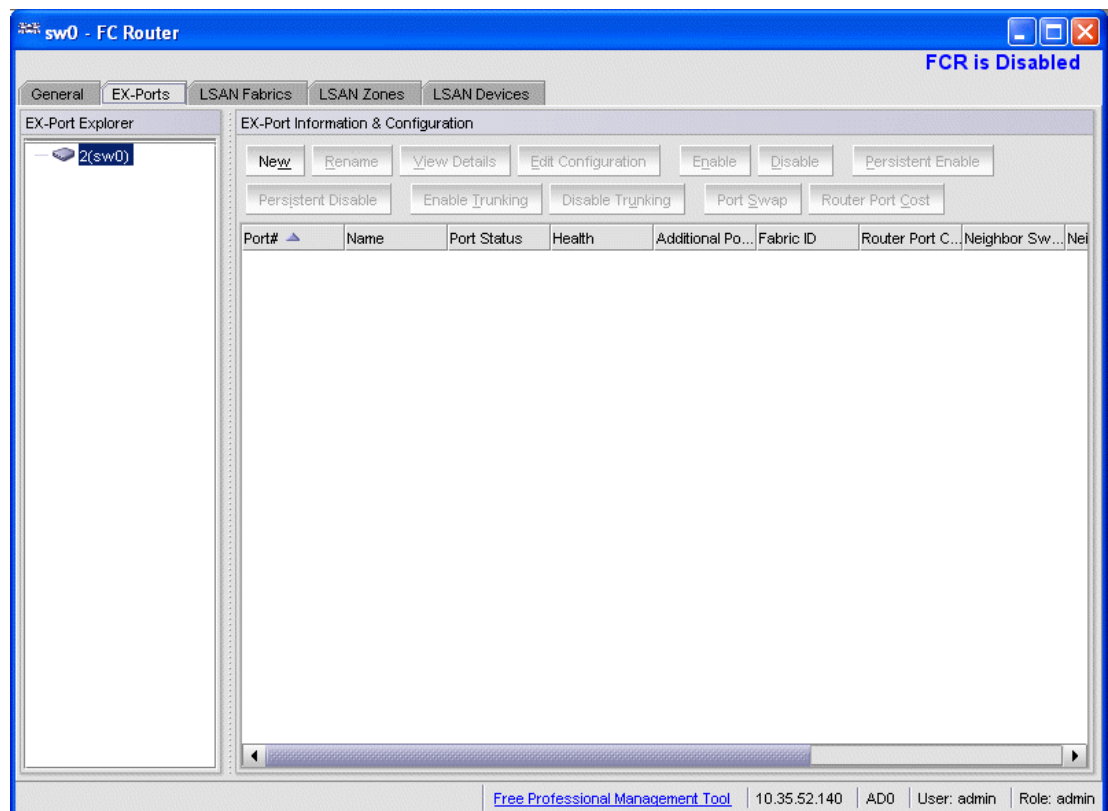
FIGURE 48   FC Routing module with EX_Ports tab selected

Configuring an EX_Port

Use the following procedure to configure an EX_Port.

1. Select **Tasks** > **Manage** > **FCR**.

2. Select the **EX_Ports** tab.

3. Click **New** in the task bar to configure one or more EX_Ports.

    This launches the port configuration wizard, which guides you through the port configuration process.

    NOTE
    The Brocade 7500, Brocade 7500E Extension , and Brocade FR4-18i switches do not require Integrated Routing License for configuring EX_Port. Other FCR Capable switches and blades require Integrated Routing License for configuring EX_Port.
    Follow the instructions in the wizard to configure the EX_Port.

    You must specify the Fabric ID and, if configuring an FC port, the speed and long distance mode. You can select any unique fabric ID as long as it is consistent for all EX_Ports that connect to the same edge fabric.

Editing the configuration of an EX_Port

Use the following procedure to edit the configuration of an EX_Port.

1. Select **Tasks** > **Manage** > **FCR**.

2. Select the **EX_Ports** tab.

3. Select a port to configure, by clicking the row.

4. Click **Edit Configuration** in the task bar. This launches the port configuration wizard, which guides you through the port configuration process. The current configuration values are displayed in the wizard steps.

    NOTE
    If you decide to configure a disabled port, the wizard provides the **Enable Port after configuration** check box. If you select this check box, the disabled port is automatically enabled after configuration. If you leave this box cleared, the port remains in the same state after configuration.

Configuring FCR router port cost

In FCR, EX_Ports can be assigned router port cost. The cost of the link is a positive number. The router port path or tunnel path is chosen based on the minimum cost per connection. If multiple paths exist with the same minimum cost, there will be load sharing over these paths. If multiple paths exist where one path costs lower than the others, then the lowest cost path is used.

Every link has a default cost. For an EX_Port 1 Gbps, 2 Gbps, 4 Gbps, and 8 Gbps links, the default cost is 1000. For a VEX_Port, the default cost is 10000. If the cost is set to 0, the default cost are be used for that link.

Use the following procedure to configure the FCR router port cost.

1. Open the **Switch View** window.

2. Click **FCR** in the **Manage** section of the **Tasks** menu.

3. Click the **Ex_Ports** tab.

4. Click the **Router Port Cost** button.

Viewing LSAN zones

The **LSAN Zones** tab displays all the LSAN zones, in both a tabular and tree form. If FC-FC Routing is disabled, the table and the tree node in this tab display only the LSAN zones present in the backbone fabric.

For more detailed information about a specific LSAN zone, click a zone name in the table and then click the **View Details** button in the task bar. You can also click the zone name in the tree on the left side of the window.

The LSAN matrix is mapping of LSAN Zones with the edge fabric they are going to communicate with. When an LSAN matrix is created in the backbone fabric, only the LSAN zones mapped in the edge fabrics are displayed in the LSAN Zones tab.

Follow the procedure described in *"Creating and populating zones"* on page 128 to create LSAN zones.

Viewing LSAN devices

The **LSAN Devices** tab displays information about the physical and proxy devices and displays these devices in a tree on the left side of the window. (If FC-FC Routing is disabled, the tables and tree nodes in this tab are empty.)

Click the **LSAN Devices** element in the tree to display a count of all the physical and proxy LSAN devices. Note that this count is for all of the LSAN fabrics.

Click the **Physical Devices** or **Proxy Devices** element in the tree to see a detailed list of the physical or proxy devices. Click the device name in the tree for more detailed information about a specific device.

Configuring the backbone fabric ID

Web Tools automatically disables FC-FC Routing before setting the fabric ID, and then re-enables it afterwards. However, you must first disable all of the EX_Ports before you begin this operation. After the fabric ID is changed, you can enable these ports again.

Use the following procedure to configure the backbone fabric ID.

1. Open the **Switch View** window.

2. Select **FCR** in the **Manage** section of the **Tasks** menu.

3. Select the **EX-Ports** tab.

4. Select all the EX_Ports in the table, and click **Disable**.

5. Select the **General** tab.

6. Click **Set Fabric ID** in the task bar.

   The **Configure Backbone Fabric ID** window displays.

7. Select a fabric ID from the drop-down menu.

   NOTE
   The fabric ID is a number from 1 through 128. Web Tools warns you if you select a fabric ID that is already in use.

   ---

8. Click **OK**.

9. After Web Tools automatically re-enables the FC-FC Routing Service, select all the EX_Ports in the table, and click **Enable..**

Using the Access Gateway

In this chapter

Access Gateway overview

Access Gateway is a software feature that allows multiple host bus adapters (HBAs) to access the fabric using fewer physical ports. You can set a switch in Access Gateway mode to transform them into a device management tool that is compatible with different types of fabrics, including Brocade Enterprise OS (EOS), and Cisco-based fabrics.

When a switch is in Access Gateway mode, it is logically transparent to the host and the fabric. Brocade Access Gateway mode allows hosts to access the fabric without increasing the number of switches and simplifies configuration and management in a large fabric by reducing the number of domain IDs and ports.

Access Gateway mode is available on the following Brocade switches:

- 5100
- 5410
- 5424
- 5450
- 5460
- 5470
- 5480
- 8000
- 300,
- Brocade Encryption Switch

For detailed descriptions of the Access Gateway, refer to the *Brocade Access Gateway Administrator's Guide*.

NOTE
When Access Gateway mode is enabled on switches managed through Web Tools, only a limited subset of menus and options related to device management are available. A switch in Access Gateway mode is considered a device management tool and not a fabric switch, therefore fabric related options are disabled, fabric management menus are unavailable, and fabric-related service requests are forwarded to the fabric switches.

Viewing Switch Explorer for Access Gateway mode

The Switch Explorer for Access Gateway mode displays as shown in Figure 49.
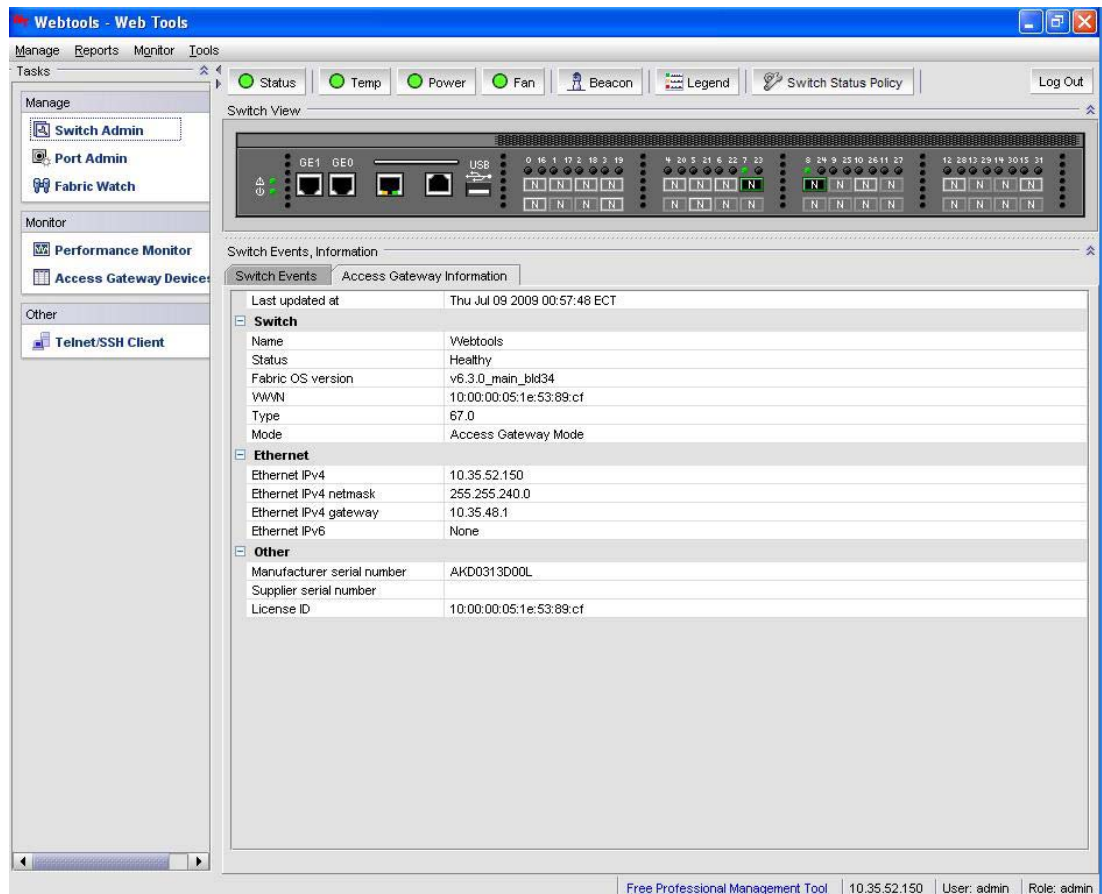
FIGURE 49   Switch Explorer view for Access Gateway mode

The Access Gateway mode Switch Explorer is divided into the following areas:

- Menu bar
- Tasks
- Switch View buttons
- Switch Events and Access Gateway information
- Indicator bar
- Professional Management Tool offering

- Switch View

Access Gateway mode

The Access Gateway feature on the Brocade Encryption switch and the Brocade 8000 enables interoperability with the Cisco fabrics. The Access Gateway mode of the switch presents standard F_Ports to the hosts, but it connects to the Enterprise fabric as N_Port (rather than as E_Port in case of a regular switch).

Restricted access in the Port Administration window

When Access Gateway mode is enabled, the following options are disabled in the **Port Administration** window:

- **Port Configuration Policy** — Configuring the Auto or Advanced mode under Port Configuration Policy is disabled as shown in Figure 50.

- **Enable Trunking and Disable Trunking** — Enabling and Disabling of N_Port trunking is disabled as shown in Figure 50.
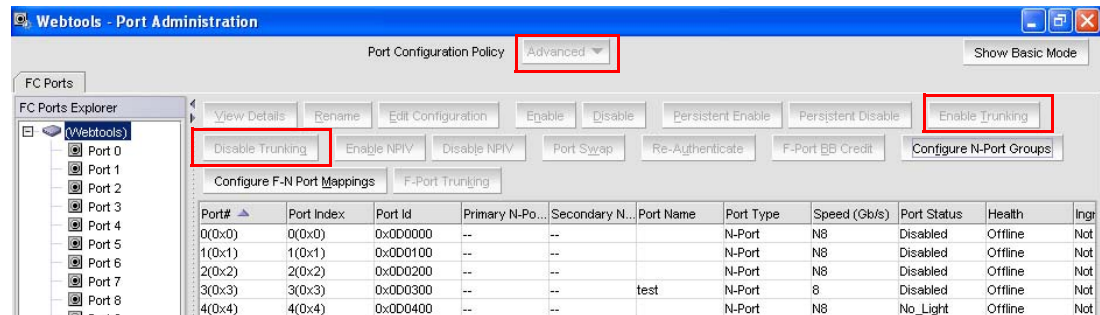
FIGURE 50   Port Configuration Policy disabled

- **Configure N-Port Groups** — You can only view the port group details from the **Port Group Configuration** window. The following options are disabled as shown in Figure 51:
    - Disable N-port Grouping
    - Add
    - Edit/View
    - Delete
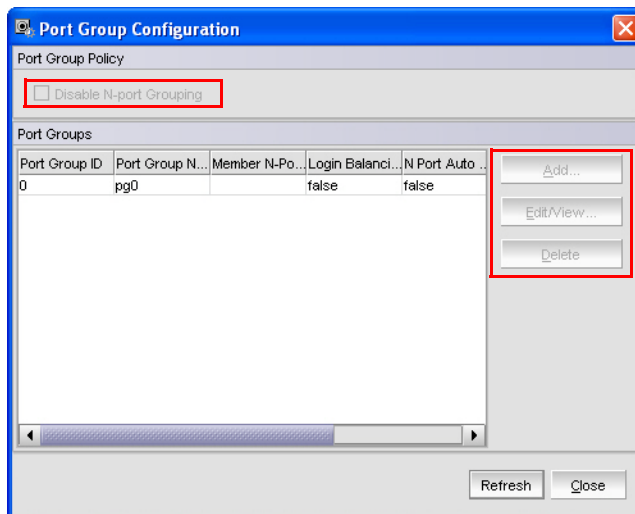
FIGURE 51   Port Group Configuration—view only

- **Configure F-N Port Mappings** — **Add** and **Remove** buttons are disabled for primary mappings and secondary failover mapping.
- **N Port Configuration** — By default all the ports are set to N_Ports and failover and fallback are disabled. You can edit the speed. The following options are disabled in the **N Port Configuration** window:
    - Lock as N Port
    - Allow as F, U Port
    - Enable N Port Failover Policy
    - Enable N Port Fallback Policy

Enabling Access Gateway mode

When you enable Access Gateway mode some fabric information, such as the zone and security databases, is erased. To recover this information, save the switch configuration before enabling Access Gateway mode.

To save the switch configuration using Web Tools, click **Switch Admin** in the **Manage** section under **Tasks**, and then select the **Configure** > **Upload/Download** subtab and upload the configuration file.

NOTE
You cannot enable Access Gateway mode if Management Server is enabled. To disable Management Server, enter the **MsplmgmtDeactivate** command.

Use the following procedure to enable Access Gateway mode.

1. Select a switch.

2. Click **Switch Admin** in the **Manage** section under **Tasks**.

   The **Switch Administration** dialog box displays.

3. Click **Disable** in the **Switch Status** section.

   You can enable Access Gateway mode only after the switch is disabled.

4. Click **Enable** in the **Access Gateway Mode** section.

5. Click **Apply**.

6. Click **Yes** to restart the switch in Access Gateway mode.

Disabling Access Gateway mode

Use the following procedure to disable Access Gateway mode.

1. Select a switch.

2. Click **Switch Admin** in the **Manage** section under **Tasks**.

   The **Switch Administration** dialog box displays.

3. Click **Disable** in the **Switch Status** section. You can disable Access Gateway mode only after the switch is disabled.

4. Click **Disable** in the **Access Gateway Mode** section.

5. Click **Apply**.

6. Click **Yes** to restart the device in native switch mode.

Viewing the Access Gateway settings

You can view the effective Access Gateway settings for the selected switch. The view can be customized. To view the Access Gateway settings select **Tasks > Monitor > Access Gateway Devices**. The **Access Gateway Device Display** dialog box displays.

Port configuration

You can configure the port types (N_Port, F_Port) on each individual port on an Access Gateway enabled switch. When you configure ports, you can specify a global configuration policy using the **Port Configuration Policy** button. By default, Advanced is selected and sets the initial defaults for port types, groups, and the F_Port-to-N_Port mappings. When the policy is Automatic, the port type assignments and mappings are configured automatically based on device and switch connections and internal load-balancing and grouping; user controls are disabled.

When you configure ports, perform the tasks in the following order:

1. Configure N_Ports, if necessary. Use the **Edit Configuration** button to configure a port.

2. Configure N_Port groups.

3. Configure F_Port-to-N_Port mappings. You can set up primary and secondary mappings. The secondary mapping is the N_Port to which an F_Port is mapped when the primary N_Port mapping goes offline.

4. Configure WWN-N_Port mappings

Creating port groups

You can group a number of N_Ports (and its mapped F_Ports) together to connect to multiple independent fabrics or to create performance optimized ports. To group a number of ports, you must create a new port group and assign desired N_Ports to it. The N_Port grouping option is enabled by default, and all N_Ports are members of a default port group 0 (pg0). Access Gateway prevents failover of F_Ports across N_Port groups.

NOTE
If you want to distribute F_Ports among groups, you can leave all ports in the default port group 0, or you can disable N_Port grouping.

Use the following procedure to create port groups.

1. Click a port in the **Switch View** to open the **Port Administration** window.

2. Make sure that you have selected **Advanced** from the **Port Configuration Policy** drop-down list.

3. Click **Configure N_Port Groups**.

> **NOTE**
> **Configure N_Port Groups** is disabled if you select **Automatic** from the **Port Configuration Policy** drop-down list.

4. In the **Port Group Configuration** dialog box, click **Add.**

    The **Add Port Group** window displays.

5. Enter the ID for the new port group in the **Port Group ID*** field.

6. Enter the name for the new port group in the **Port Group Name** field.

7. Select the **Login Balancing** check box to enable login balance for the port group.

8. Select the **Fabric Name Monitoring** check box to manually configure the managed fabric name monitoring.

9. Under the **Select Members(N-Port)*** section, select the required ports you want to group.

10. Click **Save**.

11. Click **Close on the Port Group Configuration dialog box**.

Editing or viewing port groups

Use the following procedure to edit port groups.

1. Click a port in the **Switch View** to open the **Port Administration** window.

2. Click **Configure N_Port Groups**.

3. On **Port Group Configuration** dialog box, select the group that you want to edit and then click **Edit/View.**

The **Edit/View Port Group** window displays.

4. Edit the name of the port group in the **Port Group Name** field.

5. Select the **Login Balancing** check box and the **Fabric Name Monitoring** check box if you want to enable these features. Clear the check boxes to disable these features.

   Upon selecting the **Login Balancing** check box, the **F Port Auto Rebalancing** and **N-Port Auto Rebalancing** check boxes and **Manual Balancing** button become enabled.

6. Click **Failover Enable**.

   A confirmation dialog box displays.

7. Click **Yes** to enable failover to all the ports in the port group or click **No** if you do not want to enable failover.

8. Click **Failover Disable.**

   A confirmation dialog box displays. Click **Yes** to disable failover to all the ports in the port group or click **No** if you do want to disable failover.

9. Click **Failback Enable**.

   A confirmation dialog box displays.

10. Click **Yes** to enable failback to all the ports in the port group or click **No** if you do not want to enable failback.

11. Click **Failback Disable.**

    A confirmation dialog box displays. Click **Yes** to disable failback to all the ports in the port group or click **No** if you do not want to disable failback.

12. Under the **Select Members(N-Port)*** section, select the required ports you want to group and clear the check boxes for the ports you want to remove from the port group.

13. Click **Save**.

14. Click **Close on the Port Group Configuration** dialog box.

Deleting port groups

NOTE
You cannot delete the default port group 0 (pg0).

Use the following procedure to delete port groups.

1. Click a port in the **Switch View** to open the **Port Administration** window.

2. Click **Configure N_Port Groups**.

3. In the **Port Group Configuration** dialog box, select the group that you want to delete and then click **Delete**. A confirmation dialog box displays.

4. Click **Yes** to confirm the action.

5. Click **Close**.

Defining custom F-N port mapping

Use the following procedure to manually change F-N port mappings.

1. Click a port in the **Switch View** to open the **Port Administration** window.

2. Click the **FC Ports** tab.

3. Click **Configure F_N Port Mappings**.

4. In the **Primary Mappings** area, select ports and use the **Add** (right arrow) button to map F_Ports or U_Ports to N_Ports.

5. *Optional:* Use the **Remove** (left arrow) button to delete an F_Port mapping from an N_Port.

6. *Optional:* Define a secondary N_Port in the **Secondary Failover Mappings** area, by selecting the ports using the **Add** and **Remove** buttons to set up the secondary mappings.

    The secondary mappings must be to a different port in the same group as the primary mapping. If a secondary port is not defined, the failover moves to any online ports within the same port group.

7. After you have made the appropriate changes, click **Save**.

Defining custom WWN-N port mappings

Use the following procedure to manually change WWN-N port mappings.

1. Open the **Port Administration** window.

2. Click the **FC Ports** tab.

3. Click **Configure WWN-N Port Mappings**.

4. In the **Primary Mappings** area, select a WWN from the left pane and a group or port from the right pane.

5. Click the **Add** (right arrow) button to map the WWN to the port or port group.

6. *Optional:* Expand the port in the right page and select the WWN and then use the **Remove** (left arrow) to remove the mapping.

7. *Optional:* Define a failover in the **Secondary Failover Mappings** area, by selecting the ports using the **Add** and **Remove** buttons to set up the secondary mappings.

    The WWN fails over to the secondary mapping if the primary mapped port is offline. If a secondary port is not defined, the failover moves to any online ports.

8. *Optional:* To create a detached WWN-N port mapping, enter the WWN value into the **WWN** field and click **Add**.

    The detached WWN port is now available for mapping.

9. After you have made the appropriate changes, click **Save**.

    Any unused WWNs are discarded.

Access Gateway policy modification

Although you can control a number of policies on switches in Access Gateway mode, Web Tools only provides the ability to enable and disable the policies. For more information on these policies please refer to *Access Gateway Administrator's Guide*.

Path Failover and Failback policies

The Path Failover and Failback policies determine the behavior of the F_Port if the primary mapped N_Port they are mapped to goes offline or is disabled. The Path Failover and Failback policies are attributes of the N_Port. By default, the Path Failover and Failback policies are enabled for all N_Ports.

Modifying Path Failover and Failback policies

Use the following procedure to modify Path Failover and Failback policies.

1. Click a port in the **Switch View** to open the **Port Administration** window.

2. Select the N_Port for which you want to modify the policy.

3. Click **Edit Configuration**.

4. Select the appropriate check box to modify the policy.

5. Click **Save**.

Enabling the Automatic Port Configuration policy

The Automatic Port Configuration (APC) policy is a global configuration policy for a switch in Access Gateway mode. By default, this policy is disabled. If you created an N_Port grouping and switching over to the automatic mode, those port groups are lost. After you enable the APC policy, you cannot define custom port type configurations, port mappings, Path Failover, and Failback settings.

NOTE
When port configuration is in auto mode, the **Configure N port groups**, **Configure F-N port mapping**, and **Configure WWN-N port mapping** buttons are disabled.

Use the following procedure to enable auto rebalancing from the **Switch Administration** window.

1. Click a port in the **Switch View** to open the **Port Administration** window.

2. Select **Automatic** from the **Port Configuration Policy** drop-down list.

   NOTE
   When **Port Configuration Policy** is set to **Advanced**, you can enable the auto rebalancing options from the **Configure N_Port Groups** dialog box through the **Port Administration** window.

3. Click **Yes** in the confirmation window.

4. In the **Switch Explorer** window, select **Switch Admin**. The **Switch Administration** window displays (Figure 52).
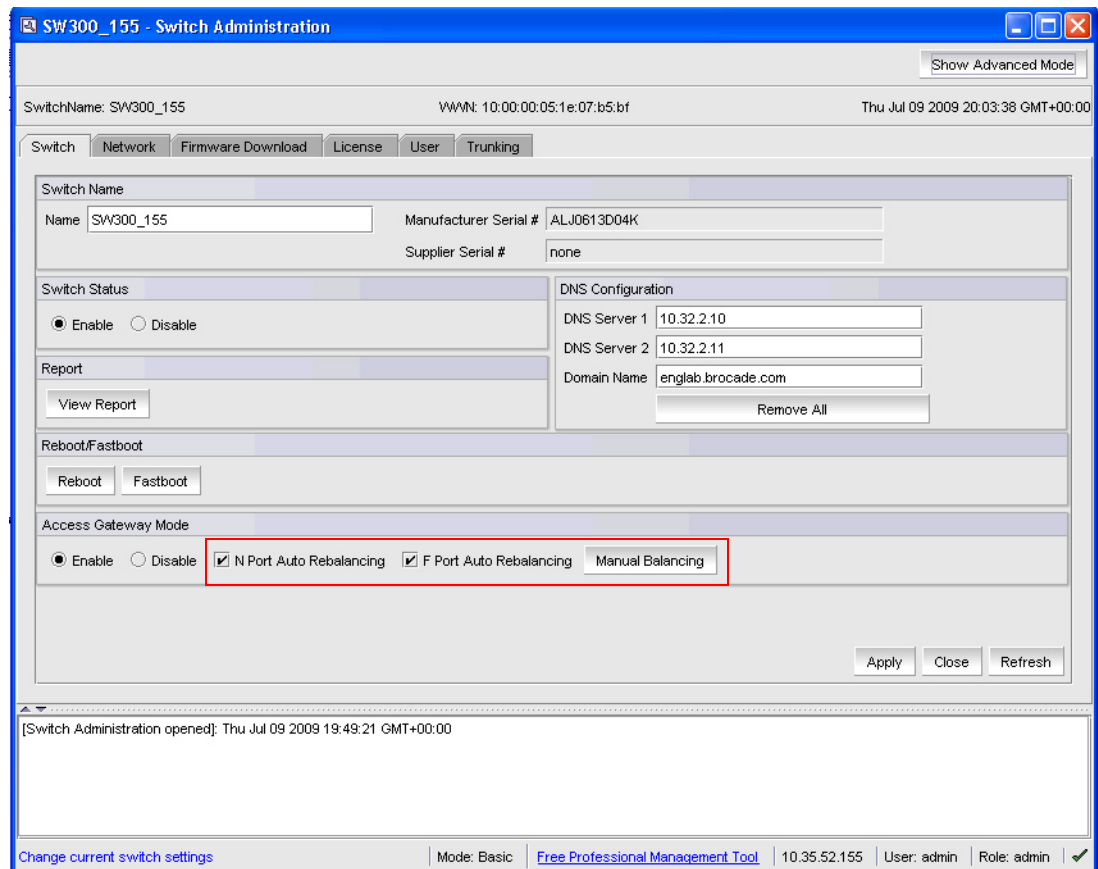
FIGURE 52   **Access Gateway Auto Rebalancing**

5. Click **Refresh.**

6. Under the **Access Gateway Mode** section, do the following:

    • Select the **N Port Auto Rebalancing** check box to enable N_Port rebalancing.

    • Select **F Port Auto Rebalancing** check box to enable F_Port rebalancing.

    • Click **Manual Balancing** and a confirmation dialog box displays. Click **Yes** to change F Port-N Port Mapping or click **No** to cancel the changes.

7. Click **Apply** to apply the changes.

Access Gateway limitations on the Brocade 8000

The following list the is a compilation of the limitations of using Access Gateway with the Brocade 8000 switch:

• Only the Port Grouping (PG) policy is supported. When the Access Gateway mode is first enabled, the default PG policy is enforced.

- When the Brocade 8000 runs in Access Gateway mode, all the FCoE ports are F-Ports and all the FC port are N-Ports.

- When Access Gateway is enabled, F Ports mapping to N Ports is allowed and all 4 FCOE ports in the trunk group are mapped to the same N Port.

- F Ports mapping to the Port Group level is not allowed.

- You cannot map individual FCoE Ports within the same trunk group to different FC ports.

- All the four FCoE ports in a set will failover or failback to one FC N Port.

- Enabling or disabling of N-port Port Grouping is not allowed

- Port Group modification is allowed

- Login Balancing and Fabric Name Monitoring is not allowed in Add Port Group configuration.

- Login Balancing, Fabric Name Monitoring, F Port Auto Rebalancing, N Port Auto Rebalancing, and Manual Balance options are not allowed in the Edit Port Group or View Port Group configuration.

Administering Fabric Watch

In this chapter

Fabric Watch overview

Fabric Watch is an optional Brocade licensed feature that monitors the performance and status of switches. Fabric Watch can automatically alert you when problems arise, before they become costly failures.

NOTE
Fabric Watch is view-only if you do not own the switch. Owning ports on a switch is not enough to enable Fabric Watch on that switch.

To use Fabric Watch, you must have a Fabric Watch license installed on the switch.

Fabric Watch tracks a number of SAN fabric elements, events, and counters. For example, Fabric Watch monitors the following:

- Fabric resources, including fabric reconfigurations, zoning changes, and new logins.

- Switch environmental functions, such as temperature, power supply, and fan status, along with security violations.

- Port state transitions, errors, and traffic information for multiple port classes as well as operational values for supported models of Finisar "Smart" GBICs/SFPs.

Fabric Watch lets you define how often to measure each switch and fabric element and allows you to specify notification thresholds. Whenever fabric elements exceed these thresholds, Fabric Watch automatically provides notification using several methods, including e-mail messages, SNMP traps, and log entries.

For detailed information regarding Fabric Watch, refer to the *Fabric Watch Administrator's Guide*.

Using Fabric Watch with Web Tools

You do not need the EGM license to perform Fabric Watch operations using Web Tools.

NOTE
Unless the switch is a member of the current Admin Domain context, Fabric Watch is view-only.
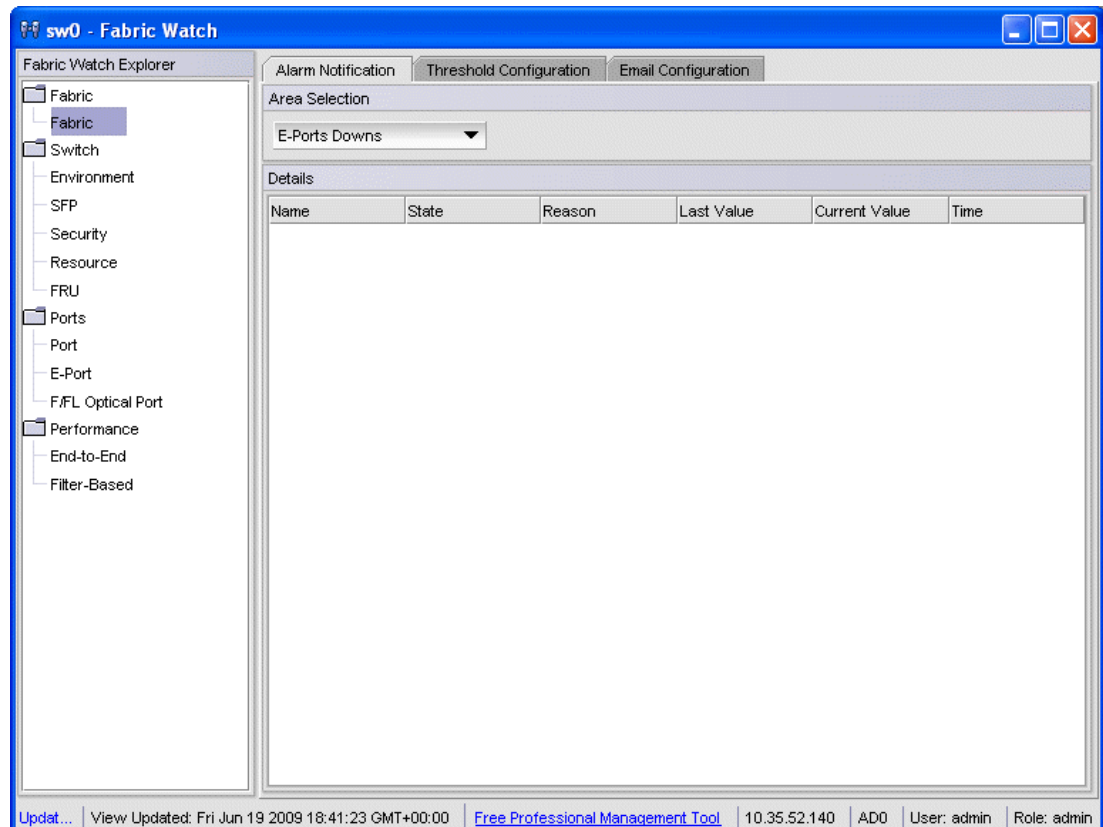
FIGURE 53   Fabric Watch window

Fabric Watch Explorer, on the left side of the window, displays the available classes. Not all classes are available for all switches. The status bar at the bottom of the window provides you with a summary of recent actions, and the date and time the module was last updated.

You can use Fabric Watch to define the following configurations:

- Configure custom threshold values on particular elements.
- Place limits on the acceptable values of those elements and enable the custom limits (configure threshold boundaries).
- Configure Fabric Watch to alert you to errant values.
- Configure Fabric Watch to identify unacceptable values (threshold traits).

Opening the Fabric Watch window

Use the following procedure to open the **Fabric Watch** window.

1. Select a switch from the **Fabric Tree** and log in if necessary.

2. Select **Tasks** > **Manage** > **Fabric Watch**.

   The Fabric Watch window displays, as shown in Figure 53.

System Monitor

The Fabric Watch license must be installed to view and modify the System Monitor details. Select **Monitor > System Monitor** to display the System Monitor.

When switch exceeds the configured usage limit, an alarm triggers. You can configure the alarm with the **Alarm Configuration** tab. The alarm can be configured for either SNMP trap, RAS log, or both.

There are three Trait and Alarm configurable values for System Monitor:

- Polling Interval
- Usage Limit
- No. of Retries

Use this procedure to configure the usage limits for System Monitor.

1. Open the **Fabric Watch** window.

2. Select either **Memory Usage** or **CPU usage**.

3. Modify the values in the **Trait Configuration** tab.

   When these values are exceeded, the alarm triggers.

4. Click the **Alarm Configuration** tab.

5. Select **SNMP Trap**, **RAS log**, or both options.

6. Click **Apply**.

Fabric Watch threshold configuration

The **Threshold Configuration** tab enables you to configure event conditions. From this tab, you configure threshold traits, alarms, and e-mail configurations.

NOTE
Use the procedures in this section to configure threshold traits for all classes except for the FRU class. Use the procedure described in "Configuring alarms for FRUs" on page 176 for the FRU class.

Configuring threshold traits

Configure threshold traits to define a threshold for a particular class and area. You can configure the following traits for a threshold:

- **Time Base**—The time base (second, minute, hour, day) for the area
- **Low Boundary**—The low threshold for the event-setting comparisons
- **High Boundary**—The high threshold for the event-setting comparisons
- **Buffer Size**—The size of the buffer zone used in event-setting comparisons

Use the following procedure to configure threshold traits.

1. Click **Fabric Watch** in the **Manage** section of the **Tasks** menu.

2. Select the **Threshold Configuration** tab.

3. Select the **Trait Configuration** subtab.

4. In **Fabric Watch Explorer,** select a class.

5. Under **Area Selection**, select an area from the list.

   This sets the units in the **Units** field.

   The module displays two columns of trait information, labeled **System Default** and **Custom Defined**. You cannot modify the information in the **System Default** column.

6. In the **Activate Level** area:

   - Use the system default settings and proceed to step 11.

     or

   - Click **Custom Defined** to specify new settings and proceed to the next step.

7. If necessary, select a time to record the event in the **Time Base** field.

8. Enter the lowest boundary of the normal zone in the **Low Boundary** field.

9. Enter the highest boundary of the normal zone in the **High Boundary** field.

10. Enter the size of the buffer zone in the **Buffer Size** field.

11. Click **Apply.**

Configuring threshold alarms

After you update the threshold information, use the **Alarm Configuration** subtab to customize the notification settings for each event setting.

The alarm naming convention is modified for Port, E-Port, and F/FL class types:

- "Above" is called "High Above"
- "Below" is called "Low Below"
- "In Between" is called "High Below"

The **Low Above** action alarm supports all port class types for these options:

- CRC errors
- Invalid words
- Protocol errors

- State change

- Trunk utilization

- C3 discards

- RX performance

- TX performance

- Loss of signal

- Link failures

- Link resets

- Packet loss

- Utilization

Use the following procedure to configure threshold alarms.

1. Open the **Fabric Watch** window.

2. Select the **Threshold Configuration** tab.

3. Select the **Alarm Configuration** subtab.

4. In **Fabric Watch Explorer**, select a class.

5. Under **Area Selection**, select an area from the drop-down list.

NOTE
The module displays two tables of alarm configuration information, labeled **System Default** and **Custom Defined**. You cannot modify the information in the **System Default** table.

---

6. In the **Activate Level** area:

   - Click **System Default** o use the system default settings and proceed to step 8.

     or

   - Click **Custom Defined** to specify new settings and proceed to the next step.

7. Select the check box for the type of notification method you want to use for each event type (Changed, Below, Above, Inbetween).

   The available alarm actions are:

   - ERROR_LOG

   - SNMP_TRAP

   - RAPI_TRAP

   - PORT_LOG_LOCK

   - EMAIL_ALERT

8. Click **Apply**.

Enabling or disabling threshold alarms for individual elements

Use the following procedure to configure element-specific alarm settings.

1. Open the **Fabric Watch** window.

2. In **Fabric Watch Explorer**, select a class.

   You can set alarms for information on a switch only if that information is monitored by Fabric Watch for that switch; not all alarm options are available for all switches. For more information, refer to the *Fabric Watch Administrator's Guide*.

3. Select the **Threshold Configuration** tab.

4. Under **Area Selection**, select the area with the alarms that you want to enable or disable.

5. Select the **Element Configuration** subtab.

6. Select an element from the **Element Selection** menu.

7. In the **Status** area:

   - To disable threshold alarms, click **Disabled** and click **Apply**. The threshold alarms are disabled and you do not need to continue with this procedure.

   - To enable threshold alarms, click **Enabled** and continue with the next step.

8. Select a behavior type for the threshold alarms:

   - Click **Triggered** to receive threshold alarms only when they are triggered by events that you defined.

   - Click **Continuous** to receive threshold alarms at a continuous interval. Select a time interval in which to receive the threshold alarms from the **Time Interval** menu.

9. Click **Apply**.

10. *Optional*: Apply the selections on this panel to multiple elements simultaneously.

    a. Click **Apply More**. The **Multiple Selection** dialog box displays.

    b. Click the boxes next to the indices of all applicable elements.

    c. Click **OK**.

Configuring alarms for FRUs

Configuration for the FRU class is different from configuration for the other classes. Because FRUs are not monitored through a threshold-based system, they have a simpler interface for configuration.

For FRUs, you configure the states for which an event occurs, as described in the following procedure.

1. Open the **Fabric Watch** window.

2. Select the **Threshold Configuration** tab.

3. In **Fabric Watch Explorer**, select a FRU class.

4. Under **Area Selection**, select a FRU type from the list.

5. Select the alarm states for which you want an event to register.

    If a FRU of the selected type is determined that it is one of the selected states, an event will occur.

6. Select the methods by which you want to be notified about the FRU alarms.

    For FRUs, the only options are error log and e-mail alert.

7. Click **Apply** to apply the changes to the switch.

    A confirmation dialog box displays, asking if you want to apply the changes to the switch.

8. Click **OK** to save the changes to the switch.

Fabric Watch alarm information

From Fabric Watch, you can view two types of reports:

- Alarm notifications—Displays the alarms that occurred for a selected class or area.
- Alarm configuration—Displays threshold and alarm configurations for a selected class or area.

Viewing an alarm configuration report

Use the **Threshold Configuration** tab, **Configuration Report** subtab to display a report of the configuration for a selected class or area with the following information:

- Threshold settings (labeled **Threshold Configuration**)
- Notification settings (labeled **Action Configuration**)
- Element settings (not labeled) — You can scroll through this information, but cannot make changes.

Use the following procedure to view an alarm configuration report.

1. Open the **Fabric Watch** window.

2. Select the **Threshold Configuration** tab.

3. Select a previously configured element from **Fabric Watch Explorer** (for instructions, refer to *"Enabling or disabling threshold alarms for individual elements"* on page 176).

4. Under **Area Selection**, select the alarm area report to be viewed.

5. Select the **Configuration Report** subtab.

    This tab displays a report of the configuration for the selected area.

Displaying alarms

Using the **Alarm Notification** tab, you can view a list of all alarms that occurred for a selected class or area (Figure 53 on page 172). Table 15 describes the columns in this report. You can click the header of each column to change the way the information is sorted in your view. You can also right-click the column header and select sort options from a menu.

NOTE
Note that for the FRU class, only the **Name**, **State**, and **Time** columns are displayed. In addition, if the FRU area is Fan, the Name column refers to either a fan or a fan FRU, depending on the switch model. Refer to "Viewing detailed fan hardware status" on page 142 for more information.

TABLE 15   Alarm notification table fields

Field	Description
Name	The string assigned to the element that had an event
State	The current state of the element
Reason	The event type that was triggered
Last Value	The data value of the element when the event was triggered
Current Value	The current data value of the element
Time	Time when the event occurred

Use the following procedure to display the alarms page.

1. Open the **Fabric Watch** window.

2. In **Fabric Watch Explorer**, select the class that you want to check for alarms.

3. Select the **Alarm Notification** tab.

4. In **Area Selection**, select the area that you want to check for alarms from the list. All alarms for that area display.

For troubleshooting responses to alarms, refer to the *Fabric Watch Administrator's Guide*.

E-mail notification

You can be notified of an alarm condition through an e-mail alert. If you have configured alarms to send an e-mail notification, you must also configure the e-mail server and the e-mail recipient, as described in the following sections.

Configuring the e-mail server on a switch

You must set up the e-mail notification recipient's DNS server and domain name on each switch for which e-mail notification is enabled.

Use the following procedure to configure the alert e-mail address on the switch.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Switch** tab.

3. In the **DNS Configuration** area, in the **DNS Server 1** field, enter the primary domain Name Server IP address. You can enter the IP address in IPv4 or IPv6 format.

4. In the **DNS Server 2** field, enter the secondary domain server IP address. You can enter the IP address in IPv4 or IPv6 format.

5. In the **Domain Name** field, enter the domain name (between 4 and 32 characters).

6. Click **Apply**.

Configuring the e-mail alert

You can set a different e-mail alert configuration for each class. For example, you can set one e-mail notification for SFPs and another for E_Ports. Before configuring e-mail alert recipients, you must set up the e-mail notification recipient's DNS server and domain name.

NOTE
You must execute the **fwalfilterset 1** command in CLI to enable e-mail notification.

Use the following procedure to configure an e-mail alert recipient.

1. Open the **Fabric Watch** window.

2. Select the **Email Configuration** tab.

3. Click **Enable** or **Disable** to enable or disable the e-mail alert status.

   NOTE
   When you disable e-mail alerts, Fabric Watch does not send e-mail notification, even if the e-mail notification method is assigned to monitored areas.

   ---

4. Enter the e-mail address of the recipient in the **Recipient Email Address** text box. Messages are sent to this address when e-mail notification is enabled.

   NOTE
   E-mail addresses must not exceed 128 characters.

   ---

5. Click **Apply**.

6. *Optional*: Click **Send Test Email** to receive a test e-mail so you can verify the e-mail notification is working correctly. You can send a test e-mail only after you have applied your settings.

Administering Extended Fabrics

In this chapter

Extended link buffer allocation overview

If the link is used over long distances, use the **Extended Fabric** tab of the **Switch Administration** window to configure the long-distance setting of a port. Because buffer credits are a switch resource, you must own the switch in order to modify extended fabric settings on a port. The EGM license must be enabled on the switch; otherwise, access to configuring long distance is denied and an error message displays.

The **Extended Fabric** tab displays information about the port speed, long-distance settings, and buffer credits, as shown in Figure 54 on page 182. For detailed information on managing extended fabrics, refer to the *Fabric OS Administrator's Guide*.

The **Extended Fabric** tab displays the following information:

• Port Number
• Buffer Limited—Indicates whether the port is buffer limited. A buffer-limited port can come online with fewer buffer credits allocated than its configuration specifies, allowing it to operate at a reduced bandwidth instead of being disabled for lack of buffers.

   Buffer-limited operation is supported for the LS and LD extended ISL modes only and is persistent across reboots, switch disabling and enabling, and port disabling and enabling.

• Port Speed—The port speed is displayed as follows:
  - 1G—1 Gbps
  - 2G—2 Gbps
  - 4G—4 Gbps
  - 8G—8 Gbps
  - N1—Negotiated 1 Gbps
  - N2—Negotiated 2 Gbps
  - N4—Negotiated 4 Gbps
  - Auto-Negotiation

• Buffer Needed/Allocated—The number of buffers needed and the number of buffers that are actually allocated.
• Actual Distance (km)—The actual distance for the link in kilometers.
• Desired Distance (km)—Required for a port configured in LD or LS mode (Table 16 on page 183), the desired distance, in kilometers, for the link.

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

For an LS-mode link, the actual distance is not measured; instead the desired distance is used to calculate the buffers required for the port.

- Long Distance—Table 16 describes the long-distance settings and identifies which settings require a Brocade Extended Fabrics license.

FIGURE 54   **Extended Fabric tab**

For Brocade DCX and DCX-4S enterprise-class platforms, the slots for CPs are not available.

You can view only two FC and two logical ports in the **Extended Fabric** tab for the 7500E Extension switches without upgrade license.

The Brocade Encryption Switch and the FS8-18 Encryption blade support auto-negotiated link speeds of 1, 2, 4, and 8 Gbps. The GE ports are always locked at 1 Gbps.

TABLE 16   Long-distance settings and license requirements

Value	Description	Extended Fabrics License Required?
L0	No long-distance setting is enabled. The maximum supported link distance is 10 km, 5 km, or 2.5 km for ports at speeds of 1 Gbps, 2 Gbps, and 4 Gbps, respectively.	No
LE	Extended normal setting is enabled, 10 km (6 miles) or less.	No
LD	Dynamic setting is enabled. Buffer credits for the given E_Port are dynamically configured based on the actual link distance, as long as this is less than the desired distance. If the actual link distance exceeds the desired distance, the desired distance is used to allocate the buffers. The LD-level link can operate at distances up to 500 km at 1 Gbps, 250 km at 2 Gbps, or 125 km at 4 Gbps, depending on the switch platform and the availability of frame buffers within the port group.	Yes
LS	Static setting is enabled. Buffer credits for the given E_Port are statically configured based on the desired link distance. The LS-level link can operate at distances up to 500 km at 1 Gbps, 250 km at 2 Gbps, or 125 km at 4 Gbit/sec, depending on the switch platform and the availability of frame buffers within the port group	Yes

Configuring a port for long distance

When you configure a long-distance ISL, ensure that the ports on both sides of the ISL have the same configuration in order to avoid fabric segmentation.

Use the following procedure to configure a port for long distance.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Click **Show Advanced Mode**.

3. Select the **Extended Fabric** tab.

4. This step is switch-specific:

   For Brocade DCX and Brocade 48000 enterprise-class platforms, click the slot subtab that corresponds to the correct slot for the logical switch.

   For Brocade 200E, 300, 4012, 4016, 4018, 4020, 4024, 4100, 4900, 5000, 5100, 5300, 7600 switches, the 7500, 7500E, 7800 Extension switches, and the Encryption Switch, proceed directly to the next step.

5. Select a distance that corresponds to the port from the **Long Distance** menu.

   Depending on the distance selected, this might require a license. For information about the various distances, refer to Table 16.

   If you select a long-distance setting of LD or LS, you must also enter a value in the **Desired Distance** column for that port number:

a. Double-click the **Desired Distance** field for the port, as shown in Figure 54.

b. Enter a number in the field to indicate the distance in kilometers. The allowed values depend on the port capability:

- If the port capability is 8 GB, type a number between 10 and 63 inclusive.
- If the port capability is 4 GB, type a number between 10 and 125, inclusive.
- If the port capability is 2 GB, type a number between 10 and 250, inclusive.
- If the port capability is 1 GB, type a number between 10 and 500, inclusive.

This value is the upper limit for calculating buffer availability for other ports in the same port group. If the actual distance is more than the desired distance, the port operates in buffer-limited mode.

c. Press **Enter** or click another port entry for the value to be accepted.

6. Click **Apply**.

The warning message, "DLS should be disabled while enabling Long distance link with Credit Recovery" displays.

7. Click **Yes** to apply the changes, or click **No** to close the confirmation message window.

Administering the iSCSI Target Gateway

In this chapter

iSCSI service overview

Through the iSCSI Target Gateway Admin module, you can perform iSCSI target gateway-related management tasks, such as creating and managing iSCSI virtual targets, managing iSCSI sessions and iSCSI authentications, and editing discovery domain sets, that enforce iSCSI device access control. The iSCSI port configuration is available to both the iSCSI Target Gateway Admin module and the port management module.

Although iSCSI service is fabric-wide, you can manage the iSCSI target gateway service through any iSCSI-capable switch in a fabric. Any applied iSCSI target gateway change is propagated and enforced to the whole fabric.

Web Tools can recognize and manage all Fibre Channel ports and GbE ports on the supported blades. When a GbE port is configured to support iSCSI, it can transport SCSI traffic over an IP network. Each GbE port has a unique IP address called an "iSCSI target portal" and each port supports 64 iSCSI sessions. The TCP/IP stack at the port provides support for multiple TCP connections over a single GbE port.

In Web Tools, ports are addressed using slot number and port number notation (for example, 2,16):

- For Fibre Channel ports on the FC4-16IP blade, the range of ports is 0 through 7.
- For GbE ports on the FC4-16IP blade, the port numbers range from ge0 through ge7. The FC4-16IP blade does not support FCIP functionality.

The iSCSI standard defines several naming conventions to enable location-independent device identification of storage resources. The FC4-16IP blade recognizes the IQN (iSCSI Qualified Name) formatted iSCSI initiator node name. For example, an iSCSI target name of "iSCSI tgt" is presented as follows:

iqn.2002-12.com.brocade:ISCSItgt

After an IQN is defined, you can map a LUN device to the IQN name.

Supported platforms for iSCSI

The iSCSI target gateway service is supported on the Brocade 48000 director with CP blades running Fabric OS v5.2.0 and later releases, and configured with an FC4-16IP blade.

Common iSCSI Target Gateway Admin functions

You can right-click on the table content in the window to access **Export**, **Copy**, and **Search** options:

NOTE
You must accept the Brocade Certificate at the beginning of the login to Web Tools to enable the functionality of Export and Copy.

- Click **Export Row** or **Export Table** to save the contents to a tab-delimited file. For CHAP, the secret is still encrypted.

- Click **Copy Row** or **Copy Table** to copy the contents in tab-delimited text format to a file.

- Click **Search** to search for a specific text string in the table.

    Enter a text string in the box that displays in the table and press **Enter**. This is an incremental search and allows 24 maximum characters including the wildcard characters: question mark (?) and asterisk (*). The first row containing the text string is highlighted. To find the next match, press the down arrow. To find the previous match, press the up arrow.

    If the text is not found in the table, the text displays in red color.

Terminology

iSCSI target gateway services require you to understand some additional terminology. Table 17 shows the terms that are used in this document to explain how the iSCSI target gateway is implemented.

TABLE 17   iSCSI target gateway services terminology

Term	Definition
iSCSI	Internet-SCSI. A transport carrier of the SCSI protocol over IP.
iSCSI target gateway	An intermediate device in the network that allows the iSCSI initiators in an IP SAN to access and utilize storage in a Fibre Channel SAN. the FC4-16IP AP blade in a Brocade 48000 director functions as an iSCSI target gateway.
iSCSI port	A special GbE port used for iSCSI only. A Fibre Channel virtual initiator is created behind each iSCSI port running as a proxy Fibre Channel initiator.
iSCSI virtual target	A unique target device in the IP SAN that contains LUNs from the real Fibre Channel targets and is identified by an IQN.
iSCSI initiator	A device that begins an iSCSI transaction by issuing a command to another device (the iSCSI target), giving it a task to perform. Typically, an iSCSI host adapter is the initiator but targets can also become initiators.
iSCSI session	An iSCSI session is the basic communication "pipe" from an iSCSI initiator to an iSCSI target. A session is a group of TCP/IP connections that link an initiator with a target (loosely equivalent to a SCSI I-T nexus).

TABLE 17   iSCSI target gateway services terminology (Continued)

Term	Definition
LUN mapping	Logical Unit Number mapping. The mapping of the virtual iSCSI target and the physical Fibre Channel target One frontend LUN (VT LUN) maps to a backend LUN (Fibre Channel LUN). The frontend LUN numbers can be different from the backend LUN numbers.
Fibre Channel LUN	The LUN identifier of the Fibre Channel target.
VT LUN	Virtual target LUN. The LUN identifier of the iSCSI virtual target.
Fibre Channel virtual initiator	(FC-VI) The iSCSI port looks like an F_Port to the rest of the system. There is one Fibre Channel virtual initiator per iSCSI. The Fibre Channel proxy initiator solution is used. FC-VI registers to the Name Server with its symbolic port name (PWWN) and node name (NWWN). The FC-VI in the Name Server entry is created irrespective of the host connectivity.
discovery domain (DD)	Created between and iSCSI host and iSCSI targets using their IQN for the purpose of iSCSI device access control.
discovery domain set (DDSet)	Created using DDs. Can be configured to enable or disable the configuration for the purpose of iSCSI device access control.
CHAP authentication	Authenticates the initiators against a list of user names and passwords with CHAP (Challenge Handshake Authentication Protocol) authentication in either one-way or mutual.
IQN	An iSCSI Qualified Name that indicates an iSCSI node name in a form that is of human readable notation using the following syntax: iqn.yyyy-mm.<reverse of DNS>:<optional iSCSI unique string>
GbE port	Gigabit Ethernet port. Uses a copper CAT-5e cable for an IP connection to an RJ-45 copper connector. FC4-16IP has 8 ports of this type that support 1 Gbps speed.
PDU	Protocol Data Unit. A unit of data with a header and an optional data section.

For additional information about iSCSI target gateway, refer to the *Fabric OS Administrator's Guide*.

Saving changes

There are several ways to save changes on the switch and apply them to the fabric (applies to the iSCSI Target Gateway Admin module only):

- Apply—Click **Apply** and your changes are transfered from the Web Tools database to the switches database and distributed throughout the fabric.
- Abort—Click **Abort** to cancel the changes before saving them. The configuration is restored to the last saved data point.
- Clear All—The **Clear All** button, located in the menu bar of the iSCSI module, has the ability to clear all parameters of the iSCSI target gateway databases, including virtual targets, iSCSI initiators, discovery domains, discovery domain sets, and all CHAP users and associated secrets. The IP interface information, however, is not deleted.

> **CAUTION**
>
> The Clear All button deletes the information from the database. Before you use the Clear All function, perform a configUpload and save a backup of the iSCSI target gateway database.

Setting up iSCSI Target Gateway services

The following procedure provides an overview of the basic steps for setting up iSCSI target gateway services. The **iSCSI Setup** wizard guides you through the steps to set up iSCSI connectivity between IP networks and your Fibre Channel SAN.

Click the **Launch Usability** wizard button on the **iSCSI Administration** window to use the **iSCSI Setup** wizard to perform all setup tasks.

ATTENTION
After mapping iSCSI targets, do not move the targets out of AD0 by adding them to other Admin Domains unless you first explicitly add them back to AD0.

Launching the iSCSI Target Gateway Admin Module

When you click **iSCSI** in the **Manage** section of the **Tasks** menu, the **iSCSI Administration** dialog box displays. This option is available on all Brocade 48000 switches with option 5 configured and with an FC4-16IP blade.

NOTE
Because the entire fabric is scanned when you open the **iSCSI Administration** window, larger fabrics may take longer to load.

The Target Group is the first pane that comes up and presents all the iSCSI virtual targets and its mapping to the Fibre Channel targets (physical and virtual) from the fabric. You can create and add LUNs to the existing iSCSI virtual targets from this group.

When you select an IQN, you have the ability to edit or delete virtual targets associated with that IQN. You can view current sessions and discovery domain accessibility.

Use the following procedure to launch the **iSCSI Target Admin Gateway** module.

1. Select a logical switch using the drop-down list under **Fabric Tree** section in the **Switch Explorer** window.

   The selected switch displays in the **Switch View**.

   Make sure that your **Admin Domain Context** is either AD0 or AD255.

   Generally, the default user Admin Domain is AD0. The recommended practice is to perform all iSCSI management from AD0; you can make changes from AD255, but you are not able to make any zoning changes.

2. Click **iSCSI** in the **Manage** section of the **Tasks** menu.

   The **iSCSI Administration** dialog box displays.

NOTE
If **iSCSI Target Gateway Services** are disabled, you must click the **Enable iSCSI** button at the top of the window to enable the services.

Launching the iSCSI setup wizard

Use the following procedure to launch the iSCSI setup wizard.

1. Select a logical switch using the drop-down list under **Fabric Tree** section in the Switch Explorer window.

   The selected switch displays in the Switch View.

   Make sure that your Admin Domain Context is either AD0 or AD255.

   Generally, the default user Admin Domain is AD0. The recommended practice is to perform all iSCSI management from AD0; you can make changes from AD255 but you are not able to make any zoning changes.

2. Click **iSCSI** in the **Manage** section of the **Tasks** menu.

   The **iSCSI Administration** dialog box displays.

3. Click **Launch Usability Wizard**.

   Follow the steps in the wizard to complete all the setup tasks.

Activating the iSCSI feature

By default, a director has iSCSI disabled. If a switch has iSCSI disabled or there is no iSCSI virtual target created, Web Tools assumes that iSCSI is not activated.

Use the following procedure to activate the iSCSI feature.

1. Ensure that the blade is inserted in the director and powered on.

2. Open the **iSCSI Target Gateway Admin** module as described in "Launching the iSCSI Target Gateway Admin Module" on page 188.

3. Click **Enable iSCSI**.

4. Click **Apply**.

Encryption services for the iSCSI Target Gateway

Data-At-Rest encryption services for the iSCSI hosts is supported for the Encryption Switch and the FS8-18 Encryption blade in the case where they are in the same fabric as the FC4-16IP and FA4-18 and the AP7600 platforms.

When the FS8-18 Encryption blade and the FC4-16IP blade are in the same fabric, the Virtual Initiator (VI) exposed by the FC4-16IP blade can be exposed as physical hosts to the FS8-18 encryption blade to provide encryption services. Similarly, the Virtual Target exposed by the FS8-18 encryption blade can be exposed as a Physical Target to the FS8-18 Encryption blade.

Configuring the IP interface

This step configures iSCSI ports (GbE Ports) found on the FC4-16IP. You must have at least one iSCSI port configured to log in to the iSCSI target.

There are two steps in this process:

- Configure the IP interface for iSCSI port.

- Configure the IP route for the iSCSI port.

The **iSCSI Port Group** tab allows you to configure iSCSI ports, display session details on a port, and show the port statistics. It also allows you to view and configure the IP interface and routes that are located on the **IP Interface** tab. You can edit or delete the IP address, but you cannot add any additional IP addresses to this interface.

When you select the switch in the left pane, the right pane lists the tasks you can perform on that switch in relation to one of the GbE ports.

When you select one of the GbE ports, you can perform the same tasks listed previously: view and capture statistics related to the port, add or delete IP addresses, add or delete IP routes, view current sessions, and view the iSCSI statistics in brief.

If an IP address or IP route is already configured on the GbE port, then it is not editable as any edits disrupt any iSCSI traffic.

Configuring the IP route is optional because when an IP address is set up, a route is automatically set up as well.

Use the following procedure to configure the IP interface.

1. Open the **iSCSI Target Gateway Admin** module as described in "Launching the iSCSI Target Gateway Admin Module" on page 188.

2. Select the **iSCSI Port** tab.

3. In the left pane, select the GbE port to use.

4. Select the **IP Interface** subtab and click **Add**.

5. Enter the IP address and subnet mask.

6. Enter the MTU size or accept the default MTU size, and click **Add**.

Editing an IP Address

Use the following procedure to edit an IP address.

1. Launch the **iSCSI Target Gateway Admin** module as described in "Launching the iSCSI Target Gateway Admin Module" on page 188.

2. Select the **iSCSI Port** tab.

3. From the left pane, select the GbE port to use.

4. Select the **IP Interface** subtab and select the item on the tab.

5. Click **Edit**.

6. Click **OK** when you receive the Warning dialog box.

7. Enter the subnet mask.

8. Enter the MTU size or accept the default MTU size, and click **OK**.

NOTE
To change the IP address, delete the current IP address and recreate it. You are not allowed to create an additional IP address *for this interface,* because there can be only one IP address per interface.

Configuring the IP route (optional)

Use the following procedure to configure the IP route.

NOTE
You can create maximum of 32 static IP routes.

1. Launch the **iSCSI Target Gateway Admin** module as described on page 188.

2. Select the **iSCSI Port** tab.

3. From the left pane, select the GbE port to use.

4. Select the **IP Routes** tab.

5. Click **Add**.

6. Enter the IP address, subnet mask, and gateway IP address, and the metric.

7. Click **Add**.

Editing the IP route

Use the following procedure to edit an IP route.

1. Open the **iSCSI Target Gateway Admin** module as described in "Launching the iSCSI Target Gateway Admin Module" on page 188.

2. Select the **iSCSI Port** tab.

3. From the left pane, select the GbE port to use.

4. Select the **IP Routes** tab.

5. Click **Edit**. A warning dialog box displays.

6. Click **OK**.

7. Enter a new value for the metric.

8. Click **OK**.

NOTE
If you want to change a value other than the metric, you must delete this route and create another in its place.

Creating iSCSI virtual targets

SCSI virtual target creation is done from the first pane in the **iSCSI Target Gateway Admin** module. The **iSCSI Virtual Target** wizard provides two ways to create iSCSI targets: **Create** and **Easy Create**. **Create** allows you to double check your work several times before committing the changes.

Use the following procedure to create an iSCSI virtual target.

1. Open **iSCSI Target Gateway Admin** module as described in "Launching the iSCSI Target Gateway Admin Module" on page 188.

2. Select the **Targets** tab.

3. Click **Create**.

The VT Configuration Wizard displays.

4. Enter an IQN. The text field displays the value "iqn" and you must enter the remaining data.

5. Click **Add LUNs**.

6. On the **LUN Addition** dialog box, select the LUNs to add.

   You must expand each unit until you get to the actual LUN.

7. Click **Add LUN(s)**.

   This adds the selected LUNs to your virtual target.

8. Click **Next** and click **Finish**.

Using Easy Create to create iSCSI virtual targets

Easy Create is an alternative method for creating iSCSI virtual targets.

Follow the Easy Create procedure to create iSCSI virtual targets.

1. Open **iSCSI Target Gateway Admin** module as described in "Launching the iSCSI Target Gateway Admin Module" on page 188.

2. Select the **Targets** tab.

3. Click **Easy Create**.

4. Follow the instructions in the wizard to create a virtual target in iSCSI.

   NOTE
   When you click **Add** in the **Easy VT Creation** dialog box, virtual targets are created for all the available physical targets in a one-to-one combination. If you add the virtual target using the **Add** button and click **OK,** the virtual target is created for only the selected physical targets.

Editing an iSCSI target

You can edit a virtual target even when there is an active iSCSI session.

Use the following procedure to edit an existing iSCSI target.

1. Open **iSCSI Target Gateway Admin** module as described in "Launching the iSCSI Target Gateway Admin Module" on page 188.

2. Select the **Targets** tab.

3. Select the IQN in the left pane of where you want to edit the targets.

4. Click **Edit**.

   The **VT Configuration** wizard displays.

5. Follow the instructions in the wizard to edit an iSCSI virtual target.

   NOTE
   The **Remove LUN(s)** button is available only for virtual targets that are fully initialized as a target.

Searching for a specific Fibre Channel target

The creation wizard has a search function to find specific Fibre Channel targets.

Use the following procedure to search for a specific FC target.

1. Click the **Search** link.

2. Input the <domain,port>, partial WWN, or vendor name, or a combination of these values.

3. Click **Next**.

   The search results are shown as selected nodes in the Fibre Channel target tree. No changes are made if search criteria does not match.

Viewing iSCSI initiators

When you set up the iSCSI target gateway on a switch, all initiators may not be online yet, but the initiators automatically pick up and display in the **Initiators** tab. The table size grows automatically to show the initiators.

This view presents all iSCSI initiators (hosts) and its associated mappings. You can view iSCSI initiator sessions here.

Discovery Domain management

You configure discovery domains (DD) and discovery domain sets for managing iSCSI device access control. The Discovery Domains pane displays all discovery domains and discovery domain sets and allows you to manage them.

When you select **DDInfo** from the tree in the left pane, you can create a discovery domain. If you select an object in the discovery domain set listed you can view, create, edit, delete, enable, or disable any of the discovery domain information contained in each object. If you select a discovery domain object, you can edit or delete the data contained in the object.

Discovery domains are placed in a discovery domain set.

Discovery domains

In the **Create DD** wizard you can configure the discovery domains (DD), add DDs to DDSets, and view the confirmation report.

NOTE
When you create new DDs, you specify a DD name, but you cannot edit or change the name when you edit the DDs.

When you launch the **Create DD** or **Create DDSet** wizard, you can add or remove virtual targets to the selected DD or DDSet. The wizard displays all available initiators and targets grouped by initiators and targets on the left side. Depending on how the wizard is launched, the right side is blank or list current members of the DD being used:

- When the wizard is launched using **Create,** the list at the right is blank.

- When the wizard is launched using **Edit**, the list at the right displays current members of the DD being viewed.

Discovery domains can be created with virtual targets, iSCSI initiators, or both.

In the wizard, you can do the following tasks:

- You can configure the DD. You must specify the DD name, and then you can add or remove initiators and targets. You can also add any offline devices by entering the IQN name in the IQN name field and clicking **Add Offline Devices** under the list on the right. The offline device name is added to the **Selected List**.

- You can also filter out initiators and targets from the tree in the **Selection List** by using the **Filter** button. You can enter the full or partial name of an iSCSI member in the IQN Name text box and click the **Filter** button. Based on the filter criteria, the tree displays only those members satisfying the filter criteria. You then add the device by selecting the device and clicking the **Add** button. In order to view the all available initiators and targets, click **Show All**.

Creating a discovery domain

Use the following procedure to create a discovery domain.

1. Open **iSCSI Target Gateway Admin** module as described in "Launching the iSCSI Target Gateway Admin Module" on page 188.

2. Select the **Discovery Domains** tab.

3. Click **Create DD**.

   The **Create DD** wizard displays.

4. Follow the instructions in the wizard to create an iSCSI discovery domain.

Editing a discovery domain

Use the following procedure to edit a discovery domain.

1. Open **iSCSI Target Gateway Admin** module as described in "Launching the iSCSI Target Gateway Admin Module" on page 188.

2. Select the **Discovery Domains** tab.

3. Select a DD in the left pane and click **Edit**.

4. Select virtual targets and use the buttons to add or remove them from the DD.

5. Click **Next**.

   The opening screen with a list of virtual targets that you added to your DDs displays.

6. Click **Next**.

   You can verify the virtual targets that you added to your DDs.

7. Click **Finish**.

   You can confirm the changes that you made before committing them.

Discovery domain sets

The **iSCSI Target Gateway Admin** module provides the flexibility to create discovery domain sets (DDSet) that define the host target access. (This functionality is similar to Fibre Channel zoning.) Use the **Discovery Domains** tab to view and manage access from iSCSI initiators to iSCSI virtual targets.

The DD view displays all DDSets created and allows you to create, edit, enable, or disable a discovery domain set. Select a DDSet from the left pane to view the contents of the set.

Discovery domains can be created but they do not have to be associated with a DDSet. However, a DDSet cannot be created without having at least one discovery domain associated in it. Only floating discovery domains are allowed.

NOTE
You can only rename a discovery domain or discovery domain set from the **iSCSI Usability** wizard.

Creating a discovery domain set

Use the following procedure to create a discovery domain set.

1. Open **iSCSI Target Gateway Admin** module as described in "Launching the iSCSI Target Gateway Admin Module" on page 188.
2. Select the **Discovery Domains** tab.
3. Click **Create DDSet**.

    The **Create DDSet** wizard displays.

4. Follow the instructions in the wizard to create an iSCSI discovery domain set.

Editing a discovery domain set

Use the following procedure to edit a discovery domain set (DDSet).

1. Launch the **iSCSI Target Gateway Admin** module as described in "Launching the iSCSI Target Gateway Admin Module" on page 188.
2. Select the **Discovery Domains** tab.
3. Select a DDSet in the left pane and click **Edit**.
4. Select the discovery domains to add to or remove from the DDSet.
5. Click **Finish** and confirm the changes that you made.

CHAP configuration

This view allows you to define access to log in to that virtual target through the Microsoft iSCSI Initiator. You can create, view, and change CHAP users and their associated secrets. After a CHAP user is created, you can modify only the CHAP secret.

The CHAP module pane lists CHAP secrets in a table with the user name and chap secret in encrypted format (*). You can add, delete, or modify CHAP entries. Each CHAP secret has:

- User name maximum length of 64 characters
- CHAP secret of maximum length of 32 characters

Creating a CHAP user

Use the following procedure to create a CHAP user.

1. Launch the **iSCSI Target Gateway Admin** module as described on .

2. Select the **CHAP** tab.

3. Click **Create**.

4. Enter the CHAP user name.

   Optional: To add more than one user at a time, click **Add**.

5. Enter a CHAP secret and click **Apply**.

Editing a CHAP secret

Use the following procedure to edit a CHAP secret.

1. Launch the **iSCSI Target Gateway Admin** module as described on .

2. Select the **CHAP** tab.

3. Click **Change CHAP Secret**. You can edit the CHAP secret, but not the CHAP user name.

4. Fill in the fields in the dialog box to edit a CHAP secret.

Binding or removing CHAP users

Use the following procedure to bind or remove CHAP users.

1. Launch the **iSCSI Target Gateway Admin** module as described on .

2. Select the **CHAP** tab.

3. Click **Bind/Remove Chap(s)**.

4. Select a virtual target.

5. Enter a new CHAP user, if necessary.

6. Select the CHAP users and click **Add** or **Remove** to move them into the appropriate list (unassociated or associated CHAP users).

7. Click **Apply**.

Connection redirection

Connection redirection allows iSCSI sessions to be evenly distributed across ports on the same blade. Before the maximum number of connections is reached for any given port, logins are redirected to the next available port, resulting in an even distribution of sessions. This distribution occurs only during the first login phase.

Enabling or disabling connection redirection

Use the following procedure to enable or disable connection redirection.

1. Launch the **iSCSI Target Gateway Admin** module as described on page 188.

2. Select the **Connection Redirection** tab.

3. Select the check box in the **Enable/Disable** column to enable connection redirection for the required slot. Clear the check box to disable connection redirection.

   NOTE
   Existing connections are not redistributed when iSCSI ports change state from disabled to enabled or offline to online.

4. Click **Apply.**

iSCSI Fibre Channel zone configuration

After you have finished setting up the iSCSI target gateway and whenever you later modify the iSCSI virtual target mappings, you must create an iSCSI Fibre Channel zone to allow the system's virtual initiators (logical FC devices that represent iSCSI initiators) to communicate in a zoned FC environment with the physical devices to which you have mapped the iSCSI virtual targets.

The procedures in this section show you how to create this zone and add it to the fabric's zone database.

Note the following:

- If you already have zone configurations defined in your fabric, you are also able to add the zone that you create here to some or all of these configurations by selecting them from a list.

- If a defined configuration is currently effective in the fabric and you add your iSCSI FC zone to that configuration, the configuration is automatically re-enabled to include this zone.

  NOTE
  If you do not have a zoning license or any zoning implemented, you do not need to create one for iSCSI target gateway service.

The following default zoning conditions apply:

- If default zoning is set to No Access, then creating an iSCSI Fibre Channel zone is mandatory as there is no way for the devices to talk to each other without one.

- If default zoning is set to All Access and there is no effective zone configuration, then you can create an iSCSI Fibre Channel zone and add it to a defined configuration, but you do not need to enable the defined configuration. Because your default zoning is All Access with no effective zone configuration, all devices can already talk to each other. However, to avoid SAN congestion in the future, you should implement a zoning plan for your devices.

Use the **Zone Admin** module to create zoning or remove or add zone members to reflect your iSCSI devices.

For more information about configuring zones, refer to "Zoning configurations" on page 121.

Creating an iSCSI Fibre Channel zone with no effective zone configuration

Use the following procedure to create an iSCSI FC zone with no effective zone configuration.

1. Open **iSCSI Target Gateway Admin** module as described in "Launching the iSCSI Target Gateway Admin Module" on page 188.

2. Click **Create iSCSI Zone**.

3. Click **Yes**.

   The **Create iSCSI Zone** wizard creates a zone called "ISCSI FC ZONE," which is not placed into a defined configuration or automatically enabled.

4. Add the ISCSI FC ZONE into a configuration.

   For instructions, refer to "Creating zone configurations" on page 131.

Creating an iSCSI Fibre Channel zone with an effective zone configuration

Use the following procedure to create an iSCSI FC zone with an effective zone configuration.

1. Launch the **iSCSI Target Gateway Admin** module as described in "Launching the iSCSI Target Gateway Admin Module" on page 188.

2. Click **Create iSCSI Zone**.

3. Click **Yes**.

4. Select a configuration in the dialog box:

   - If you select a non-effective configuration, the iSCSI Fibre Channel zone is added into that configuration. The configuration is not re-enabled and remains in the defined configuration until you enable it. You must add the iSCSI Fibre Channel zone to the effective configuration at a later date or iSCSI target gateway does not work.

   - If you select an effective configuration, the iSCSI Fibre Channel zone is added into the effective configuration, and then the configuration is re-enabled. This affects the entire SAN; the zoning database must update itself, and then replicate its changes into the fabric.

ATTENTION
Schedule your changes during a maintenance cycle if you decide to add the iSCSI Fibre Channel zoning members to an effective configuration. Re-enabling the effective configuration affects the entire fabric.

5.   Click **OK**.

     The effective configuration is modified and re-enabled.

Managing and troubleshooting accessibility

The Web Tools iSCSI accessibility feature helps you do the following:

- Verify that both host and target are online.
- Verify that the effective discovery domain set has both host and target.
- Allow an initiator or target to access the other.
- Deny an initiator or target to access the other.
- Verify that the iSCSI Fibre Channel zone has been set up and, if appropriate, enable the defined configuration. Refer to "Creating an iSCSI Fibre Channel zone with an effective zone configuration" on page 198.

Routing Traffic

In this chapter

Routing overview

NOTE
To perform routing operations and Dynamic Load Sharing (DSL) configurations, the EGM license must be installed on the switch; otherwise, access to these features is denied and an error message displays.

For Fabric OS v6.4.0, the supported routing policies are:

- Port-based routing — Port-based routing assigns a "static route," in which the path chosen for traffic never changes.

- Exchanged-based routing — Exchange-based routing policy is the default. Exchange-based routing policy always employs "dynamic path selection," in which the software defines a path based on current traffic conditions.

Refer to the *Fabric OS Administrator's Guide* for more information.

To optimize port-based routing, the DSL can be enabled to balance the load across the available output ports within a domain. Exchange-based routing *requires* the use of DLS; when this policy is in effect, you cannot disable the DLS feature.

Use the **Routing** tab of the **Switch Administration** window to view and modify routing information. Figure 55 on page 202 displays the **Routing** tab.
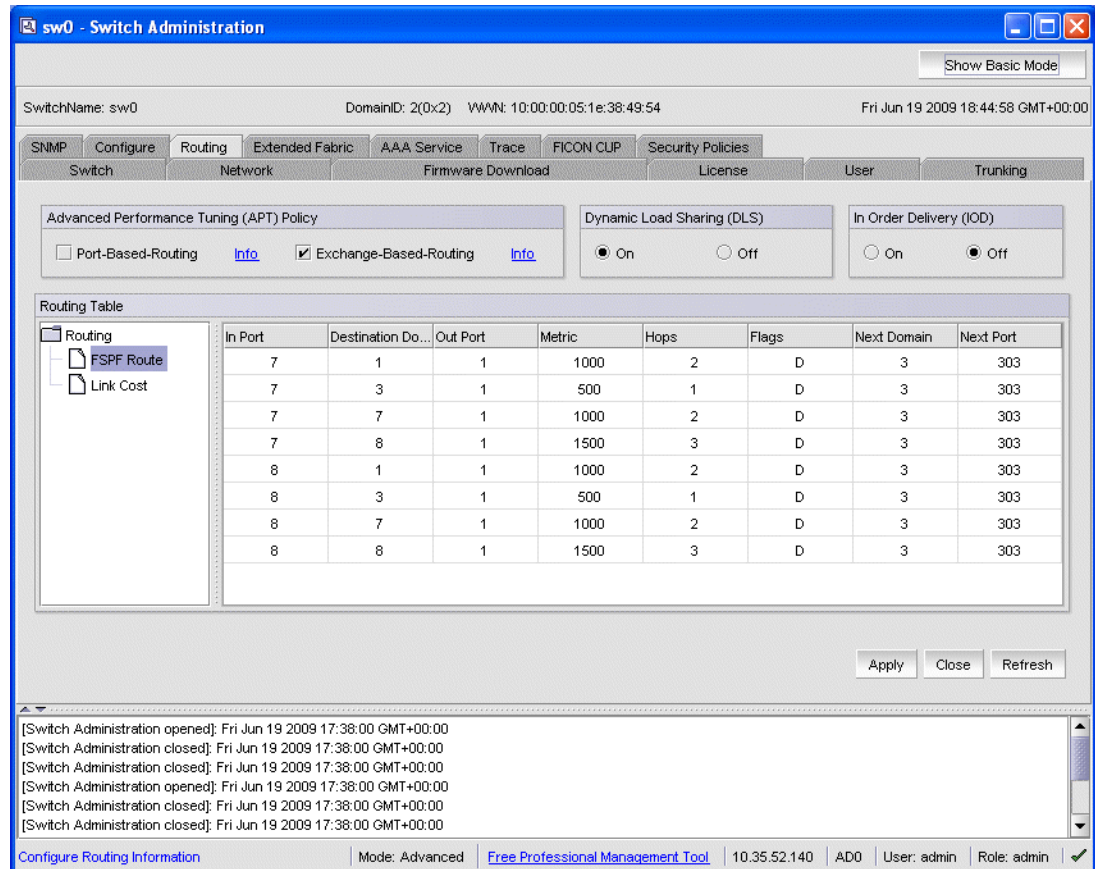
FIGURE 55 Routing tab

Viewing fabric shortest path first routing

The **Routing** tab of the **Switch Administration** window displays information about routing paths.

Use the following procedure to view the fabric shortest path first routing.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Routing** tab.

3. This step is switch-type specific:

- For the Brocade 48000 and Brocade DCX and DCX-4S enterprise-class platforms, click a slot number under the FSPF Route category in the navigation tree.

- For the Brocade 200E, 300, 4012, 4016, 4018, 4020, 4024, 4100, 4900, 5000, 5100, 5300, 7600 switches, the 7500, 7500E and 7800 Extension switches, and the Encryption Switch, click the FSPF Route category in the navigation tree.

Configuring dynamic load sharing

The exchange-based routing policy depends on the Fabric OS dynamic load sharing feature (DLS) for dynamic routing path selection. When this policy is in force, DLS is always enabled and cannot be disabled.

When the port-based policy is in force, you can enable DLS to optimize routing. When DLS is enabled, it shares traffic among multiple equivalent paths between switches. DLS recomputes load sharing either when a switch boots up or each time an E_Port or FX_Port goes online or offline. Enabling this feature allows a path to be discovered automatically by the FSPF path-selection protocol.

For more information regarding DLS, refer to the **dlsset** command in the *Fabric OS Command Reference*.

Use the following procedure to configure dynamic load sharing.

1. Open the **Switch Administration** window as described in “Opening the Switch Administration window” on page 35.

2. Select the **Routing** tab.

3. Select **On** in the **Dynamic Load Sharing (DLS)** area to enable dynamic load sharing or select **Off** to disable dynamic load sharing.

   When the exchange-based routing policy is in effect, the DLS radio buttons display on the **Routing** tab

4. Click **Apply**.

   The warning message, “Credit Recovery for Long distance links should be turned off using CLI while enabling DLS” displays.

5. Click **OK**.

Lossless dynamic load sharing

Lossless dynamic load sharing (DLS) is supported in following platforms:

- DCX with 8G blades
- DCX-4S with 8G blades
- Brocade 300
- Brocade 5100
- Brocade 5300
- Brocade 7800 on FC ports.
- Brocade FX8-24 on FC ports.

You can enable this loss less feature from WT. If you try to enable loss less when DLS is OFF, an error message displays.

Perform the following procedure to enable or disable loss less DLS.

1. Open the **Switch Administration** window as described in “Opening the Switch Administration window” on page 35.

2. Select the **Routing** tab.

3. Select **On** in the **Loss Less (DLS)** area to enable the mode, or select **Off** to disable dynamic load sharing.

   When the exchange-based routing policy is in effect, the Loss Less DLS radio buttons display on the **Routing** tab

4. Click **Apply.**

5. Click **OK.**

Specifying frame order delivery

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 example, if a link goes down), traffic is rerouted around the failure, and some frames could be delivered out of order.

By default, frame delivery is out-of-order across topology changes. However, if the fabric contains destination devices that do not support out-of-order delivery, you can force in-order frame delivery across topology changes.

Enabling in-order delivery (IOD) guarantees that frames are either delivered in order or dropped. For more information regarding IOD, refer to the *Fabric OS Administrator's Guide*.

NOTE
Enabling in-order delivery can cause a delay in the establishment of a new path when a topology change occurs, and therefore should be used with care.

Perform the following procedure to specify frame order delivery.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Routing** tab.

3. Click **On** in the In-Order Delivery (IOD) area to force in-order frame delivery across topology changes or click **Off** to restore out-of-order frame delivery across topology changes.

4. Click **Apply.**

Configuring the link cost for a port

This section describes how to set the cost of an interswitch link (ISL). The cost of a link is a dimensionless positive number. The fabric shortest path first (FSPF) protocol compares the cost of various paths between a source switch and a destination switch by adding the costs of all the ISLs along each path. FSPF defines the path with minimum cost. If multiple paths exist with the same minimum cost, FSPF employs load sharing over these paths.

Every ISL has a default cost that is inversely proportional to its bandwidth. For a 1-Gbps ISL, the default cost is 1000. For a 2-Gbps ISL, the default cost is 500.

Use this procedure to set a non-default, "static" cost for any port.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Routing** tab.

3. This step is switch-specific:

   - For the Brocade 48000 and Brocade DCX, enterprise-class platforms, click the slot number of the logical switch under **Link Cost** in the navigation tree.

   - For Brocade 200E, 300, 4012, 4016, 4018, 4020, 4024, 4100, 4900, 5000, 5100, 5300, 7600 switches, the 7500 and 7500E Extension switches, and the Encryption Switch, click **Link Cost** in the navigation tree.

4. Double-click in the row in the **Cost** column that corresponds to the appropriate port.

5. Enter the link cost. Valid values for link cost are from 1 through 65535. Setting the value to 0 sets the link cost to the default value for that port.

6. Click **Apply**.

Configuring Standard Security Features

In this chapter

User-defined accounts

In addition to the default accounts—root, factory, admin, and user—Fabric OS supports up to 256 user-defined accounts in each logical switch (domain). These accounts expand your ability to track account access and audit administrative activities.

When the Virtual Fabrics capability is enabled, each user-defined account is associated with the following:

- Virtual Fabric ID—Specifies the accessible Virtual Fabrics for a user account.
- Home Virtual Fabric—Specifies the default Virtual Fabric for a user account.
- Role—Determines functional access levels within the Virtual Fabric.

When the Admin Domain capability is enabled, each user-defined account is associated with the following:

- Admin Domain list—Specifies the accessible Admin Domains for a user account.
- Home Admin Domain—Specifies the default Admin Domain for a user account. The home Admin Domain must be a member of the user's Admin Domain list.
- Role—Determines functional access levels within the bounds of the user's current Admin Domain.

NOTE
Virtual Fabrics and Admin Domains are mutually exclusive.

Access rights for any user session are determined by the user's role-based access rights. Refer to Chapter 1, "Introducing Web Tools" for additional information about Role-Based Access Control (RBAC).

The **User** tab of the **Switch Administration** window (Figure 56 on page 209) displays account information. You can create and manage accounts depending on your role. The roles and permissiosn are listed in Table 18.

TABLE 18   User role and permissions

Role	Permissions
admin	Create and manage all predefined and user-defined accounts
operator	Change your own password and cannot create, modify, or view predefined or user-defined accounts
securityadmin	Create and manage all security roles.
switchadmin	Change your own password and cannot create, modify, or view predefined or user-defined accounts
zoneadmin	Change your own password and cannot create, modify, or view predefined or user-defined accounts
fabricadmin	Change your own password and cannot create, modify, or view predefined or user-defined accounts
basicswitchadmin	Change your own password and cannot create, modify, or view predefined or user-defined accounts
user	Change your own password and cannot create, modify, or view predefined or user-defined accounts

Virtual Fabrics considerations

If no home logical fabric ID is specified for a user, the system provides a default home ID. The default home ID is 128.

Admin Domain considerations

For legacy users with no Admin Domain specified, the user has access to AD 0 through 255 (physical fabricadmin) if their current role is Admin. Otherwise, the user has access to AD0 only.

If some Admin Domains were defined for the user and all of them are inactive, the user is not allowed to log in to any switch in the fabric.

If no Home Domain is specified for a user, the system provides a default home domain. The default home domain for predefined account is AD0. User-defined accounts, the default home domain is the Admin Domain in the user's Admin Domain list with the lowest ID.

NOTE
The **User** tab displays and changes information in the switch database. If you have RADIUS configured, note that this tab displays the logged-in RADIUS account information but does not allow the user to modify the RADIUS host server database.
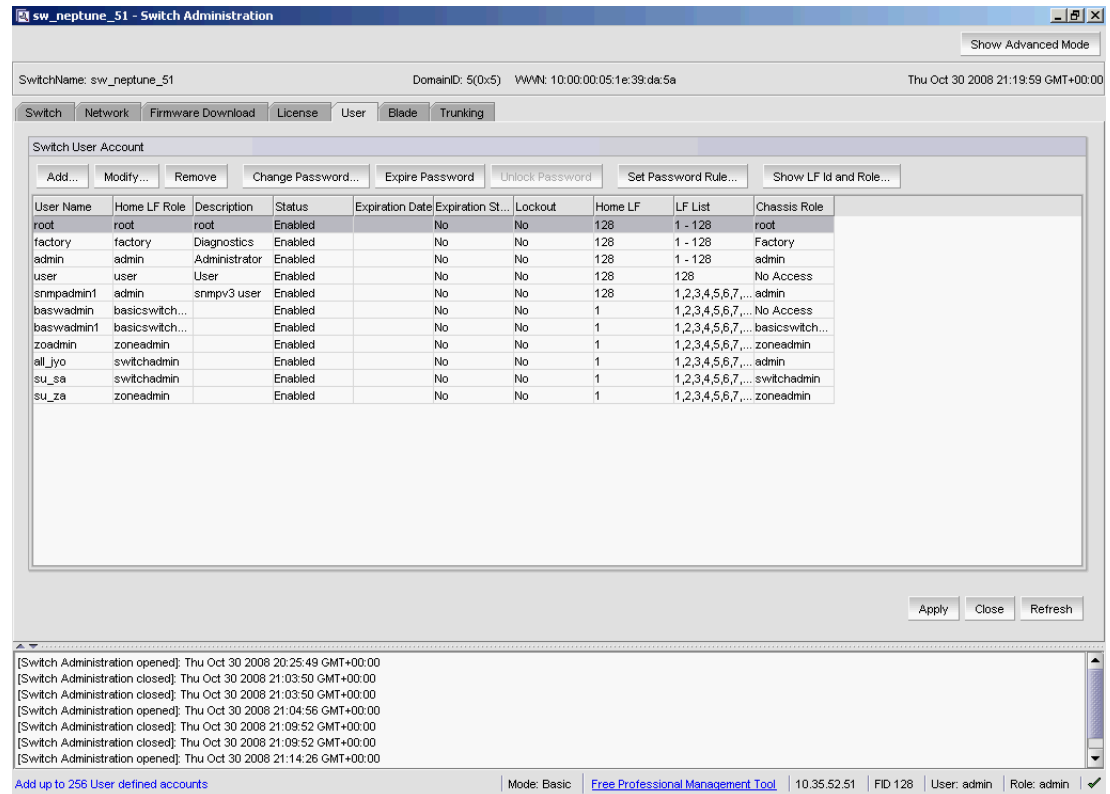
FIGURE 56   **User tab**

Viewing user account information

Use the following procedure to view user account information.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **User** tab.

    A list of the default and user-defined accounts displays. If you are logged in using the switchadmin role, only your account information displays.

Creating user-defined accounts

Use the following procedure to create user-defined accounts.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **User** tab.

3. Click **Add**.

The **Add User Account** dialog box displays. For switches that support Virtual Fabrics, refer to Figure 57. For switches that support Administrative Domains (AD), refer to Figure 58.

FIGURE 57   Add User Account dialog box (VF)

FIGURE 58   Add User Account dialog box (AD)

4. Enter the user name.

   The user name must begin with an alphabetic character.

   The name can be up to 40 characters long. It is case-sensitive and can contain alphabetic and numeric characters, the dot (.) and the underscore (_). It must be different from all other account names on the logical switch.

5. Select a role from the drop-down menu.

   For VF-enabled switches, the selection is done per logical fabric ID. (Refer to "Role-Based Access Control" on page 14 for information about these roles.)

6. *Optional*: Enter a description of the account.

7. Click **Enabled** or **Disabled** to enable or disable the account.

8. Enter the password for the account.

   The password is not displayed when you enter it on the command line.

   Passwords can be from 8 through 40 characters long. They must begin with an alphabetic or numeric character. They can include alphanumeric characters, the dot (.), and the underscore (_). They are case-sensitive.

   Passwords must also meet any additional password rules that were set up. (Refer to the procedure "Setting the rules for passwords" on page 214 for more information.)

9. Retype the password in the **Confirm Password** field for confirmation.

10. Check the available Virtual Fabrics or Admin Domains that you can access.

    For Virtual Fabrics, all logical fabrics IDs (1-128) are displayed, even if they have not all been created. Only Admin Domains that were created and are accessible to you display.

    If all the Admin Domains in the list are inactive, then you cannot log in to the switch.

    The **All** option does not mean all of the listed Admin Domains; it means all Admin Domains from AD0 through AD255, regardless of whether they were already created.

    The **All** button is disabled unless the following conditions are met:

    - The selected role for the target user must be admin or securityadmin.
    - You must be a physical fabric administrator.

    Selecting **All** makes the target user account a physical fabric administrator.

11. Select a home logical fabric ID if Virtual Fabrics are enabled, or select a home domain for the user from the **Home AD** menu if Admin Domains are enabled.

    The default home logical fabric ID is 128.

    NOTE
    If AD0 is deselected in the user's Admin Domain list and no other Admin Domains are selected, the next available Admin Domain becomes the user's default home Admin Domain.

12. For Virtual Fabrics environments, select a **Chassis Role**.

    The chassis role determines the RBAC role and permissions of the user for performing all chassis-level operations in all logical fabrics.

13. Click **OK**.

14. On the **User** tab, click **Apply** to apply your changes.

Deleting user-defined accounts

Use the following procedure to delete user-defined accounts.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **User** tab.

3. Select the account to remove and click **Remove**.

4. Click **Apply** to save your changes.

You cannot delete the default accounts. An account cannot delete itself. All active command line interface (CLI) sessions for the deleted account are logged out.

Changing user account parameters

Use the following procedure to change the role, add or change the description, and enable or disable accounts. Note that you cannot change the user name of the account using this procedure. To change the user name, you must delete the account and create a new account.

Users can select their own accounts in the user account table and change the password. All other buttons are unavailable.

Use the following procedure to change the user account parameters.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **User** tab.

3. Select the account to modify.

NOTE
You cannot modify the default root and factory accounts, even if you are logged in as root.

4. Click **Modify**.

   The **Modify User Account** dialog box displays.

NOTE
If the user account you are modifying does not have a subset of your Admin Domains, a warning message displays to inform you of the permissions conflict.

5. Select a role from the menu.

   You can change the role only on user-level accounts. You cannot change the role on the admin or root accounts. You cannot change the role of your own account.

6. Enter a new description.

   You can change the description only on user-level accounts. You cannot change the description of the default accounts. You cannot change the description of your own account.

7. Click **Enabled** or **Disabled** to enable or disable the account.

   You can enable and disable user- and admin-level accounts, but not your own account. You cannot enable or disable your own account or the factory account. Only the root account can disable itself. If you disable an account, all active CLI sessions for that account are logged out.

8. Check the available Admin Domains that the user can access.

    Only Admin Domains that have already been created and are accessible to you display. If all the Admin Domains in the list are inactive then you cannot log in to the switch.

    NOTE
    The **All** option does not mean all of the listed Admin Domains; it means all Admin Domains from AD0 through AD255, regardless of whether they were already created.

    The **All** button is disabled unless the following conditions are met:

    -   The selected role for the target user must be admin or securityadmin.
    -   You must be a physical fabric administrator.

    Selecting **All** makes the target user account a physical fabric administrator.

9. Select a home domain for the user from the **Home AD** menu.

    If AD0 is deselected in the user's Admin Domain list and no other Admin Domains are selected, the next available Admin Domain becomes the user's default home Admin Domain.

10. Click **OK** and click **Apply** to apply your changes.

Maintaining passwords

When you expire a password, the next time that user logs in, Web Tools requires the user to provide a new password.

NOTE
You have to own the switch in order to modify password rules.

A password becomes locked if a user has exceeded the maximum number of failed login attempts. This number is specified in the **Lockout Threshold** field. To unlock a locked password, refer to the unlock procedure in "Unlocking a password" on page 215.

Changing the password of an account

If you are logged in as admin, you can change the password of your own account, peer admin accounts, switchadmin accounts, and user accounts. You can also change the root or factory account passwords.

If you are changing the password of an admin account, you must also provide the current password. You do not need to provide the current password if you are changing the password of a lower-level user account.

Passwords can be from 8 through 40 characters long. They must begin with an alphabetic or numeric character. They can include alphanumeric characters, the dot (.), and the underscore (_). They are case-sensitive.

Passwords must also meet any additional password rules that were set up. (Refer to "Setting the rules for passwords" on page 214 for more information.)

Use the following procedure to change the password of an account.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **User** tab.

3. Select the account to modify.

   If you are logged in as a switchadmin, you can only change the password of your own account.

4. Click **Change Password**.

   The **Set User Account Password** dialog box displays.

5. Enter the current password of the account.

   This step is required only if you are changing the password of your own or a peer admin account.

6. Enter the new password of the account.

   The new password must have at least one character different from the old password.

7. Retype the new password in the **Confirm Password** field.

8. Click **OK**.

9. Click **Apply** to save your changes.

Setting the rules for passwords

Use the following procedure to set rules for passwords.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **User** tab.

3. Click **Set Password Rule**.

   The **Configure Password Rule** dialog box displays.

4. Fill out the dialog box for the password rules you want to enforce.

   The available options are:

   - Minimum number of days (0–999) before you can change the password again
   - Number of days (0–999) before a password expires
   - Number of password changes before you can reuse a password
   - Minimum password length (8–40 characters)
   - Minimum number of uppercase and lowercase characters required
   - Minimum number of digits and punctuation characters required
   - Number of characters that can be repeated in the password
   - Number of failed login attempts (0–999) before the password is locked from further change attempts, and the amount of time the password is locked (0–99999 minutes)
   - Number of days to warn user before password expiration (0–999)

5. Select whether to enable or disable the lockout administration features.

   If you select to disable the lockout administration, the user is never locked out of the system.

6. Click **OK** to close the dialog box.

7. Click **Apply** to save your changes.

Setting a password as expired

Use the following procedure to set a password as expired.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **User** tab.

3. Select the account.

4. Click **Expire Password**.

   If the button is unavailable, the password is already expired.

5. Click **Apply** to save your changes.

Unlocking a password

Use the following procedure to unlock a password.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **User** tab.

3. Select the account.

4. Click **Unlock Password.**

   If the button is unavailable, the password is already unlocked or was not locked out.

5. Click **Apply** to save your changes.

Displaying roles and assigned logical fabrics

You can display user role assignments for logical fabrics.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **User** tab.

3. Select an account.

4. Select **Show Role and VF**. The role mapping for that user displays.

Access control list policy configuration

Support for the Access Control List (ACL) policies is currently defined in the Switch Connection Control (SCC) and Device Connection Control (DCC) policies. SCC and DCC policy configuration in base Fabric OS is performed on a switch-local basis.

Fabric Configuration Server (FCS) Policy can be created only once. While creating the FCS policy, the local switch WWN is automatically included in the list. In the FCS list, the switch in the first position becomes the primary FCS switch. If the first switch in the FCS list is not reachable, the next switch becomes the primary switch. You can also explicitly specify the primary FCS switch.

If there is no SCC, DCC, or FCS policy, the defined and active list is blank.

Virtual Fabrics considerations

ACL policies can be implemented at the logical switch/logical fabric level.

Admin Domain considerations

ACL management can be done on AD255 and in AD0 only if there are no other user-defined Admin Domains. Both AD0 (when no other user-defined Admin Domains exist) and AD255 provide an unfiltered view of the fabric. If there are user defined Admin Domains, then ACL management can be done on AD255 only.

Creating an SCC, DCC, or FCS policy

You can create the FCS policy only once.

Use the following procedure to create an SCC, DCC, or FCS policy.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Security Policies** tab.

3. Select the **ACL** subtab.

4. Select a policy by clicking on the appropriate tab (**SCC**, **DCC**, or **FCS**).

5. Click **Edit**.

   This launches the **ACL Policy Configuration** wizard.

6. Select the policy type you want to edit.

7. Click **Next** and click **Create**.

8. *SCC Option:* Add switches to an SCC policy by selecting one or more switches and clicking **Add** or **Add All**.

9. *SCC Option:* To add an offline switch, click **Add other Switch** and enter the WWN.

10. *DCC Option:* Select the ports to add to a DCC policy.

    When you launch the **DCC Policy Configuration** wizard, only the launched switch and its ports are listed in the tree. All the devices in the fabric are also listed in the tree.

11. In the **ADD Domain, Port Index** field, enter the value in the Domain, Index format and click **Add**.

12. Click **OK** to confirm the changes to the switch.

13. Activate the policy in order to implement it. Refer to "Activating all SCC, DCC, or FCS policies" on page 217 for instructions.

Editing an SCC, DCC, or FCS policy

Use the following procedure to edit an SCC, DCC, or FCS policy.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Make sure the **Show Advanced Mode** option is selected.

3. Select the **Security Policies** tab.

4. Select a policy by clicking on the appropriate tab.

5. Click **Edit**.

    This launches the **ACL Policy Configuration** wizard.

6. Select the policy type you want to edit.

7. Click **Next** and click **Modify**.

8. Select a switch or highlight multiple switches to add to the policy by clicking **Add** or **Add All**.

9. Select a switch or highlight multiple switches to remove a policy by clicking **Remove**.

10. Click **Next** and click **Finish** to confirm the changes to the switch.

Deleting all SCC, DCC, or FCS policies

You cannot delete the FCS policy from non-primary or non-FCS switches.

The **Delete All** button is enabled only when there is at least one policy activated.

Use the following procedure to delete all SCC, DCC, or FCS policies.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Security Policies** tab.

3. Click **Delete All**.

    A warning message displays.

4. Click **OK** to delete all the policies.

Activating all SCC, DCC, or FCS policies

After a policy is created or modified, you can distribute it to the remaining fabric.

To delete a policy, you must activate a new or empty policy.

Use the following procedure to activate all SCC, DCC, or FCS policies.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Security Policies** tab.

3. Click **Activate All** to activate all the policies.

NOTE
Activating the policy moves it into the **Activate Policy Set** window.

Distributing an SCC, DCC, or FCS policy

Perform this procedure to distribute an SCC, DCC, or FCS policy.

NOTE
SCC and DCC policy can be distributed only for a primary switch.

Use the following procedure to distribute an SCC, DCC, or FCS policy.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Security Policies** tab.

3. Select the appropriate tab (**SCC**, **DCC**, or **FCS**).

4. Click **Distribute Policy**.

5. Select the switches that will receive the policy.

6. Select **OK**.

    If the policy distribution fails, an error dialog box displays.

Moving an FCS policy switch position

You can move the position of a primary switch in the FCS policy list.

Use the following procedure to move an FCS policy switch position.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Security Policies** tab.

3. Select the **FCS** tab.

4. Click **Move FCS Switch**.

5. Select the appropriate from and to positions.

6. Click **Apply**.

7. After you move all the member switches, click **Apply** and **Close**.

Fabric-Wide Consistency Policy configuration

Fabric-Wide Consistency Policy (FWCP) configures the Fabric Wide Consistency behavior of distributable ACL policies. The policy ensures that the switches in the fabric enforce the same policies. Set a strict or tolerant fabric-wide consistency policy for each ACL policy type (SCC, DCC, FCS) to automatically distribute that database when a policy change is activated. If a fabric-wide consistency policy is not set, then the policies are managed on a per switch basis.

Use the following procedure to set the fabric-wide consistency policy for an SCC, DCC or FCS policy

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Security Policies** tab.

3. Under **Security Policies**, click **FWCP**.

4. Select one of the following consistency behavior for the required policy type (**SCC, DCC, FCS**).

    - **Absent**
    - **Tolerant**
    - **Strict**

    NOTE
    You can change the consistency behaviors of SCC, DCC, or FCS policy only for a primary switch.

5. Click **Apply**.

6. Click **Yes** to accept the changes.

    NOTE
    If the switch is not a primary switch, an error message dialog box displays.

7. Click **No** to discard the changes and click **Refresh** in the **FWCP Configuration** window to manually refresh the window.

8. Click **Close.**

Authentication policy configuration

You can configure an authentication protocol policy for E_Port and F_Port authentication, and then distribute the authentication policy to other switches in the fabric. You can also set shared secret keys.

Configuring authentication policies for E_Ports

Use the following procedure to configure authentication policies for E_Ports.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Security Policies** tab.

3. Select **Authentication** on the **Security Policies** menu.

4. In the **Authentication Type** field, select **FCAP** or **DHCHAP**.

5. Select the switch authentication policy mode:

    On   Strict authentication is enforced on all E_Ports.

    Active   The switches can be connected to a switch with any
                    type of policy.

Passive	The switch does not initiate authentication but participates if the connecting switch initiates authentication.
Hash	A hash function (like SHA or MD5) is used for authentication.
Off	The switch does not support authentication. Any authentication negotiation is rejected.

6. Select a **DH-Group** type.

7. *Optional:* Set the **device authentication policy mode** to either **off** or **passive** and click **Apply**.

Configuring authentication policies for F_Ports

Use the following procedure to configure authentication policies for F_Ports.

1. Open the **Switch Administration** window and click **Show Advanced Mode**, if not selected.

2. Select the **Security Policies** tab.

3. Select **Authentication** on the **Security Policies** menu.

4. In the **Authentication Type** field, select **DHCHAP**.

NOTE
You must select **DHCHAP** when you are configuring authentication for an F_Port.

5. Set the **switch authentication mode** to either **off** or **passive** and click **Apply**.

Distributing authentication policies

Authentication policies are distributed only if all the selected switches accept the distribution. Only the policy mode is distributed to the selected switches. The switch initiating the distribution must accept distribution.

Use the following procedure to distribute authentication policies.

NOTE
You cannot distribute authentication policies in AD0 unless it is the only Admin Domain.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Security Policies** tab.

3. Select **Authentication** on the **Security Policies** menu.

4. Click **Distribute Policy**.

5. Select the switches or click the button to distribute to all.

6. Click **OK**.

Re-authenticating policies

A user who has changed authentication policy parameters or a shared secret key pair can re-initialize the authentication.

Use the following procedure to re-authenticate policies.

1. Click a port in the **Switch View** to open the **Port Administration** window.

    The **Port Administration** window displays with the port selected.

2. Click **Re-Authenticate** (active only for F_Ports and E_Ports).

3. Close the window.

Setting a shared secret key pair

DH-CHAP requires a shared secret key pair between two entities to authenticate with each other. A key pair consists of a local secret and a peer secret. The local secret identifies the local switch. The peer secret identifies the entity to which the local switch may authenticate.

Use the following procedure to set a shared secret key pair.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Security Policies** tab.

3. Select **Authentication** on the **Security Policies** menu.

4. Select the **Shared Secret Keys** subtab.

5. Click **Add**.

    The **Add Shared Secret Keys** dialog box displays.

6. Enter the Switch WWN, name, or domain ID, or use the **Browse** button to select a switch.

7. In the **Peer Secret** and **Confirm Peer Secret** fields, enter the peer secret value.

8. In the **Local Secret** and **Confirm Local Secret** fields, enter the local secret value.

9. Click **Add**.

10. When you are finished adding secret key pairs for switches, click **Apply**.

Modifying a shared secret key pair

You can edit and modify the secret key pairs by switch.

Use the following procedure to modify a shared secret pair.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Security Policies** tab.

3. Select **Authentication** on the **Security Policies** menu.

4. Select the **Shared Secret Keys** subtab.

5. Select a secret key pair and click **Edit**.

6. Make the appropriate changes and click **OK**.

Setting the Switch Policy Authentication mode

This setting determines whether or not authentication is required when a switch logs in to a fabric.

Use the following procedure to set the Switch Policy Authentication mode.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **Security Policies** tab.

3. Select **Authentication** on the **Security Policies** menu.

4. Use the **Switch Policy Authentication Mode** option to select the authentication policy.

SNMP configuration

This section describes how to manage the configuration of the SNMP agent in the switch. The configuration includes SNMPv1 and SNMPv3 configuration, accessControl, and systemGroup configuration parameters.

Access is read-only if you do not have admin or security admin authority.

For more information, refer to the **snmpConfig** command in the *Fabric OS Command Reference*.

Setting SNMP trap levels

Use the following procedure to set SNMP trap levels.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **SNMP** tab.

3. Select a trap level for a recipient from the corresponding **Trap Level** menu in the **SNMPv1** and **SNMPv3** sections.

   The level you select identifies the minimum event level that prompts a trap.

   NOTE
   Adding or editing the user name can be done only through the CLI and by selecting a user name from the **User Name** menu in the **SNMPv3** section.

4. Click **Apply**.

Changing the systemGroup configuration parameters

Use the following procedure to change the systemGroup configuration parameters.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **SNMP** tab.

3.   Enter a contact name, description, and location in the **SNMP Information** section.

4.   *Optional*: Select the **Enable Authentication Trap** check box to allow authentication traps to be sent to the reception IP address.

5.   Click **Apply**.

Setting SNMPv1 configuration parameters

Use the following procedure to set SNMPv1 configuration parameters.

1.   Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2.   Select the **SNMP** tab.

3.   Double-click a community string in the **SNMPv1** section and enter a new community string.

4.   Double-click a recipient IP address in the **SNMPv1** section and enter a new IP address.

5.   Click **Apply**.

Setting SNMPv3 configuration parameters

NOTE
The port number is not included.

Use the following procedure to set SNMPv3 configuration parameters.

1.   Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2.   Select the **SNMP** tab.

3.   Select a user name from the **User Name** menu in the **SNMPv3** section.

   NOTE
   The list is scrollable. If you do not see your user name, scroll down using the scroll bar or by clicking the **User Name** heading.

4.   Double-click a recipient IP address in the **SNMPv3** section and enter a new IP address.

5.   Select a trap level from the **Trap Level** menu.

6.   *Optional:* Select the **Enable SNMPv3 Informs for all Trap Recipients** check box to enable or disable inform requests for all trap recipients.

7.   Enabling SNMPv3 informs allows you to enter the **Engine ID**.

   The Engine ID is required to authenticate the inform request. If informs request is disabled, the SNMP manager does not send a response to the sender.

8.   Click **Apply**.

Changing the access control configuration

NOTE
The port number is not included.

Use the following procedure to change the access control configuration.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **SNMP** tab.

3. Double-click an access host IP address in the **Access Control List** section and enter a new host IP address. You can enter an IP address in either IPv4 or IPv6 format. When you use the IPv6 format, you must include a prefix; for example, `fec0::2002/64`.

   NOTE
   The list is scrollable. If you do not see your user name, scroll down using the scroll bar or by clicking the **Access Host** heading.

   ---

4. Select a permission for the host from the **Access Control List** menu.

   Options are **Read Only** and **Read Write**.

5. Click **Apply**.

RADIUS management

Fabric OS supports RADIUS authentication, authorization, and accounting service (AAA). When configured for RADIUS, the switch becomes a Network Access Server (NAS) that acts as a RADIUS client. In this configuration, authentication records are stored in the RADIUS host server database. Login and logout account name, assigned role, and time accounting records are also stored on the RADIUS server.

You should set up RADIUS through a secure connection such as SSH.

The following are the three choices in the drop-down menu when RADIUS is selected as the primary service (Figure 59):

- **Switch Database when RADIUS Authentication Fails**—When selected, the switch user login database is checked whenever RADIUS authentication fails.

- **Switch Database When RADIUS Times Out**—Switch user login database is checked only if the physical connection to the RADIUS server fails.

- **None**—Switch user login database is never checked. Only a RADIUS server can be used for authentication.
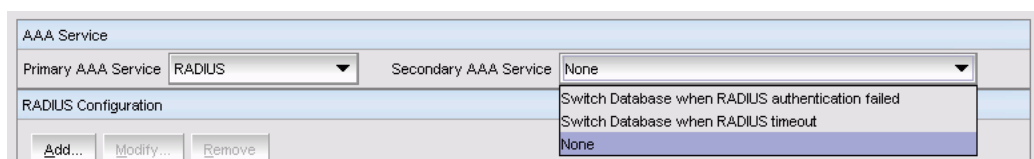
FIGURE 59   **Choices in the Secondary AAA Service drop-down menu**

If the switch database is selected as primary, there is no secondary option. The RADIUS server cannot be configured as a backup for the switch user login database.

When the primary AAA service is RADIUS, you have three secondary service choices:

- **None**
- **Switch Database when RADUIS authorization fails**
- **Switch Database when RADIUS times out**

When RADIUS login fails, even though RADIUS server is available, the additional service allows you the option to use the Switch Database as backup authentication service when the RADIUS server is not available. Alternatively, you can have no secondary AAA service, which means that only the primary service is used for authentication.

Use the **AAA Service** tab of the **Switch Administration** window to manage RADIUS.

Enabling and disabling RADIUS

At least one RADIUS server must be configured before you can enable RADIUS.

Use the following procedure to enable or disable RADIUS.

1. Open the **Switch Administration** window as described in .
2. Select the **AAA Service** tab.
3. To enable RADIUS, select **RADIUS** from the **Primary AAA Service** drop-down menu.
4. Select **None**, **Switch Database when RADIUS Login Failed**, or **Switch Database when RADIUS Login Timeout** from the **Secondary AAA Service** menu.

   NOTE
   To disable RADIUS, select **Switch Database** from the **Primary AAA Service** menu and select **None** from the **Secondary AAA Service** menu.

5. Click **Apply**.

Configuring RADIUS

The configuration is chassis-based, so it applies to all logical switches (domains) on the switch and it is replicated on a standby CP, if one is present. It is saved in a configuration upload, and can be applied to other switches in a configuration download. You should configure at least two RADIUS servers so that if one fails, the other server assumes the service.

You can configure RADIUS even if it is disabled. You can configure up to five RADIUS servers. You must be logged in as admin, switchadmin, or securityadmin to configure RADIUS.

Use the following procedure to configure RADIUS.

1. Open the **Switch Administration** window as described in .
2. Select the **AAA Service** tab.
3. Click **Add**.

> The **RADIUS Configuration** dialog box displays. You can configure up to five RADIUS servers. If five RADIUS servers are already configured, the **Add** button is disabled.

4. Enter the RADIUS server name, as a valid IP address (in either IPv4 or IPv6 format) or Dynamic Name Server (DNS) string.

   Each RADIUS server must have a unique IP address or DNS name for the RADIUS server.

5. Enter the port number.

6. Enter the secret string.

7. Enter the timeout time in minutes.

8. Select either **CHAP** or **PAP** as the authentication protocol.

   The default value is CHAP, and if you do not change it, CHAP becomes the authentication protocol.

9. Click **OK** to return to the **AAA Service** tab.

10. Click **Apply**.

Modifying the RADIUS server

Use the following procedure to change the parameters of a RADIUS server that is already configured.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **AAA Service** tab.

3. Select a RADIUS server from the **RADIUS Configuration** list.

4. Click **Modify**.

   The **RADIUS/ADLDAP Configuration** dialog box displays.

5. Enter new values for the port number, timeout time (in minutes), and secret string.

6. Select either **CHAP** or **PAP** as the authentication protocol.

   The default value is CHAP, and if you do not change it, CHAP becomes the authentication protocol.

7. Click **OK** to return to the **AAA Service** tab.

8. Click **Apply**.

Modifying the RADIUS server order

The RADIUS servers are contacted in the order they are listed, starting from the top of the list and moving to the bottom.

Use the following procedure to modify the RADIUS server order.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2. Select the **AAA Service** tab.

3. Select a RADIUS server from the RADIUS Configuration list.

4.   Click the up and down arrows to rearrange the order of the RADIUS servers.

5.   Click **Apply.**

Removing a RADIUS server

Use the following procedure to remove a RADIUS server.

1.   Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2.   Select the **AAA Service** tab.

3.   Select a RADIUS server from the **RADIUS Configuration** list.

4.   Click **Remove.**

     If there is no RADIUS server configured, the **Remove** button is disabled. You cannot remove the only RADIUS server if RADIUS is the primary AAA service.

     The RADIUS server is not deleted until you apply the changes from the **AAA Services** tab.

5.   Click **Apply** in the **AAA Services** tab.

     A confirmation displays, warning you that you are about to remove the selected RADIUS server.

6.   Click **Yes** in the confirmation.

Active Directory service management

Active Directory is the directory server that holds all the user profiles. Active Directory provides user authentication and authorization using LDAP as authentication protocol. Active Directory provides better security while using remote authentication mechanism.

You can add, remove, and modify settings of Active Directory Server.

Enabling Active Directory service

For adding a new Active Directory server, you must provide the server IP address, port number, secret string, timeout value, and LDAP as the authentication protocol. The server IP address may be in either IPv4 or IPv6 format. Select **Active Directory** as the server type; the dialog box displays LDAP as the only authentication protocol.

Use the following procedure to enable Active Directory service.

1.   Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

2.   Select the **AAA Service** tab.

3.   To enable Active Directory service, select **Active Directory** from the **Primary AAA Service** menu.

4.   Select **None**, **Switch Database when Active Directory authentication failed**, or **Switch Database when Active Directory timeout** from the **Secondary AAA Service** menu.

> **NOTE**
> To disable **Active Directory** service, select **Switch Database** from the **Primary AAA Service**
> drop-down menu and select **None** from the **Secondary AAA Service** drop-down menu.

5. Click **Apply.**

Modifying Active Directory service

Use the following procedure to change the parameters of a Active Directory service that is already
configured.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration
    window" on page 35.

2. Select the **AAA Service** tab.

3. Select a server from the **ADLDAP Configuration** list.

4. Click **Modify.**

    The **RADIUS/ADLDAP Configuration** dialog box displays.

5. Enter new values for the port, timeout, and domain.

6. Click **OK** to return to the **AAA Service** tab.

7. Click **Apply.**

Removing Active Directory service

Use the following procedure to remove a RADIUS server.

1. Open the **Switch Administration** window as described in "Opening the Switch Administration
    window" on page 35.

2. Select the **AAA Service** tab.

3. Select a server from the **ADLDAP Configuration** list.

4. Click **Remove.**

> **NOTE**
> The server is not deleted until you apply the changes from the **AAA Services** tab.

5. Click **Apply** in the **AAA Services** tab.

    A confirmation dialog box displays, warning you that you are about to remove the selected
    server.

6. Click **Yes** in the confirmation dialog box.

IPsec concepts

Internet Security Protocol (IPsec) is a set of open standards that provide cryptographic security
services for IP networks. Several protocols are available for providing authentication and secure
transmission of data.

From Web Tools, you can establish IPsec policies for FCIP implementations on 7800 extension switches with the upgrade license, the 7500 extension switches and FR4-18i blades, and you can establish IPsec policies for IP interfaces that provide management access to switches and control processors.

There are several protocols and algorithms that can be applied. Choosing the protocols and algorithms you want to use may be a matter of adapting to an implementation that is already in place in your LAN, or you may need to do a significant amount of research and planning. The supported protocols and algorithms are defined and described in the RFCs listed in Table 19.

TABLE 19   Relevant RFCs

RFC number	Title
RFC 4301	Security Architecture for the Internet Protocol
RFC 4302	IP Authentication Header
RFC 4303	IP Encapsulating Security Payload
RFC 4304	Extended Sequence Number (ESN) Addendum to IPsec Domain of Interpretation (DOI) for Internet Security Association and Key Management Protocol (ISAKMP)
RFC 4305	Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header
RFC 4869	Suite B Cryptographic Suites for IPsec
RFC 4309	Using Advanced Encryption Standard (AES) CCM Mode with IPsec Encapsulating Security Payload (ESP)
RFC 4306	Internet Key Exchange Version 2 (IKEv2) Protocol
RF C4307	Cryptographic Algorithms for Internet Key Exchange Version 2 (IKEv2)
RFC 3971	Secure Neighbor Discovery
RFC 3972	Cryptographically Generated Addresses
RFC 3041	Privacy Extensions for Stateless Address Auto configuration in IPv6

Transport mode and tunnel mode

Transport mode adds an authentication header (AH) before the IP header. Only a single pair of addresses is used (those in the IP header). When transport mode is used, both endpoints implement IPsec.

Tunnel mode encapsulates an IP datagram in a new datagram, with a new IP header specifying the addresses of the tunnel end points. IPsec is implemented between tunnel endpoints. IPsec is transparent to the actual endpoints within the IP header in the original packet.

Figure 60 provides a basic visual comparison of how transport mode and tunnel mode modify an IP datagram.
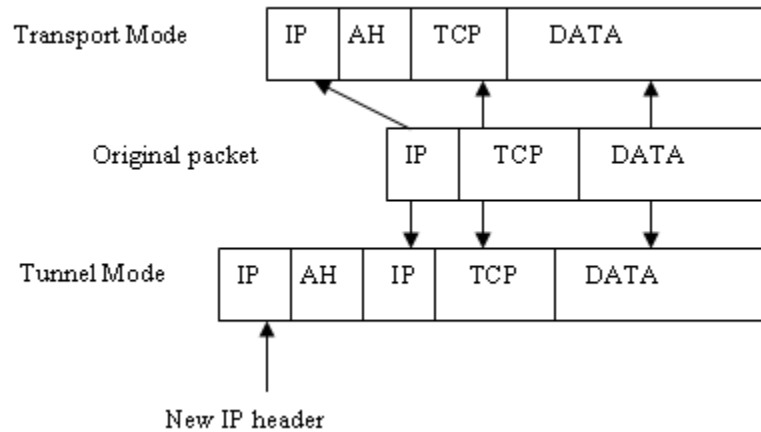
FIGURE 60   Transport mode and tunnel mode comparison

IPsec header options

IPsec adds headers to an IP datagram to enable authentication and privacy. There are two options:

- Authentication Header (AH)
- Encapsulating Security Payload (ESP)

Authentication Header

AH can be used to authenticate a data stream, but does not provide encryption needed for privacy. The AH contains a message authentication code (MAC). The MAC is created by a hash algorithm calculation. The MAC is transmitted in an IP datagram. The same hash algorithm is then used by the receiver to verify the integrity of the packet. AH can be used in either transport mode or tunnel mode, as shown in Figure 61.
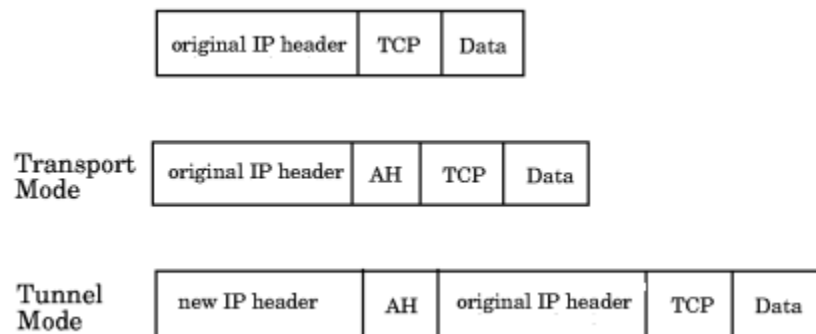
FIGURE 61   AH header in transport mode and tunnel mode

Encapsulating Security Payload

ESP provides authentication, and also provides privacy by encrypting the IP datagram. The use of an ESP header is similar to the use of the AH header. A hash algorithm is used to calculate an authentication value, the authentication value is sent in an IP datagram, and the same hash algorithm is used by the receiver to verify the authentication value. ESP can be used in either transport mode or tunnel mode, as shown in Figure 62.
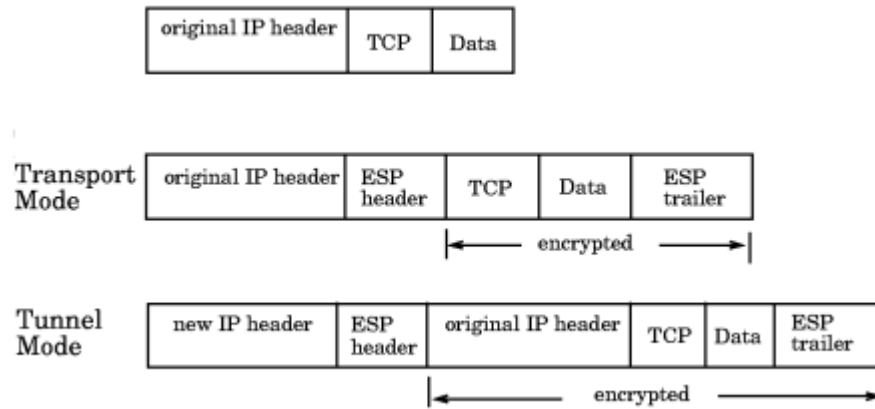
FIGURE 62   ESP header in transport mode and tunnel mode

Basic IPsec configurations

There are three basic configurations for IPsec use:

*   Endpoint to Endpoint
*   Gateway to Gateway
*   Endpoint to Gateway

Endpoint to Endpoint

In an endpoint to endpoint configuration, both endpoints implement IPsec. Transport mode is commonly used in endpoint to endpoint configurations, and only a single pair of addresses is used. Typically, this kind of configuration would be used for direct communication between hosts. There are two drawbacks to consider:

*   If network address translation (NAT) is used on the connection, one or both endpoints may be behind a NAT node. If that is the case, UDP must be used to encapsulate the tunneled packets. Port numbers in the UDP headers can then be used to identify the endpoint behind the NAT node.
*   Packets cannot be inspected or modified in transit. This means that QoS, traffic shaping, and firewall applications cannot access the packets, and does not work.

Gateway to Gateway

In a gateway to gateway configuration, IPsec protection is implemented between network nodes. Tunnel mode is commonly used in a gateway to gateway configuration. A tunnel endpoint represents a set of IP addresses associated with actual endpoints that use the tunnel. IPsec is transparent to the actual endpoints.

Endpoint to Gateway

In an endpoint to gateway configuration, a protected endpoint connects through an IPsec protected tunnel. This can be used as a virtual private network (VPN) for connecting a roaming computer, like a service laptop, to a protected network.

Internet Key Exchange concepts

Internet Key Exchange (IKE) is used to authenticate the end points of an IP connection, and to determine security policies for IP traffic over the connection. The initiating node proposes a policy based on the following:

- An encryption algorithm to protect data.
- A hash algorithm to check the integrity of the authentication data.
- A Pseudo-Random Function (PRF) algorithm that can be used with the hash algorithm for additional cryptographic strength.
- An authentication method requiring a digital signature, and optionally a certificate exchange.
- A Diffie-Hellman exchange that generates prime numbers used in establishing a shared secret key.

Encryption algorithms

An encryption algorithm is used to encrypt messages used in the IKE negotiation. Table 20 lists the available encryption algorithms. A brief description is provided. If you need further information, please refer to the RFC.

TABLE 20   Encryption algorithm options

Encryption algorithm	Description	RFC number
3des_cbc	3DES processes each block three times, using a unique 56-bit key each time.	RFC 2451
null_enc	No encryption is performed.	
aes128_cbc	Advanced Encryption Standard (AES) 128 bit block cipher.	RFC 4869
aes256_cbc	Advanced Encryption Standard (AES) 256 bit block cipher.	RFC 4869

Hash algorithms

Hash message authentication codes (HMAC) check data integrity through a mathematical calculation on a message using a hash algorithm combined with a shared, secret key. Table 21 lists the available encryption algorithms. The sending computer uses the hash function and shared key to compute a checksum or code for the message, and sends it to the receiving computer. The receiving computer must perform the same hash function on the received message and shared key and compare the result. If the hash values are different, it indicates that a third party may have tampered with the message in transit, and the packet is rejected.

TABLE 21   Hash algorithm options

Hash algorithm	Description	RFC/Publication number
aes_xcbc	Uses a cypher block and extended cypher block chaining (CBC).	RFC 3566
hmac_md5	The MD5 computation produces a 128-bit hash.	RFC 1321
hmac_sha1	The SHA1 computation produces a 160-bit hash.	FIPS Pub 180-1

Pseudo-Random Function algorithm

The Pseudo-Random Function (PRF) algorithm generates output that appears to be random data, using the HMAC chosen as the hash algorithm as the seed value. PRF is used to strengthen security.

Public key certificate-based authentication

Industry standard X.500 database servers are available as certificate authority servers to enable certificate-based authentication of computers.

SA lifetime

The SA lifetime may be defined as the number of bytes transmitted before the SA is rekeyed, or as a time value in seconds, or both. When both are used, the SA lifetime is determined by the threshold that is first reached. Whenever an SA lifetime expires, the security association (SA) is renegotiated and the key is refreshed or regenerated.

For example, if a 200 MB file is transferred with a 100 MB lifetime, at least two keys are generated. If a communication takes one hour, and you specify a lifetime of 300 seconds (five minutes), more than 12 keys may be generated to complete the communication.

The SA lifetime limits the length of time a key is used before it is replaced by a new key, thus limiting the amount of time a given key is available to a potential attacker. Part of a message may be protected by an old key, while new keys protect the remainder of the message, so even if an attacker deciphers one key, only a portion of the message is vulnerable.

Diffie-Hellman groups

Diffie-Hellman (DH) groups are used to determine the length of the base prime numbers for the Diffie-Hellman exchange. Diffie-Hellman key exchange is a cryptographic protocol that allows two parties that have no prior knowledge of each other to jointly establish a shared secret key over an insecure communications channel.

DH group choices are 1(modp768), 2(modp1024), 14(modp2048), and 18(modp8192). Each group provides an incrementally more secure key exchange by providing more bits (768, 1024, 2048, 8192).

Authentication methods

The methods used to authenticate the IKE peer are preshared key (psk), DSS digital signature (dss), and RSA digital signature (rsasig):

- A Preshared key (PSK) is a shared secret that is shared between two parties over a secure channel before it is used. Typically, the PSK is a password or pass phrase. PSKs are created in the end systems used by the two parties. There are several tools available to help select a strong key that will work with various operating systems. When choosing a tool and creating a PSK, keep in mind that the cryptographic strength of a key generally increases with length.

- The Digital Signature Standard (DSS) makes use of a private key to generate a digital signature. Each user possesses a private and public key pair. Signature generation can be performed only by the possessor of the user's private key. The digital signature is sent to the intended verifier in a message. The verifier of the message and signature verifies the signature by using the sender's public key.

- The RSA digital signature process uses a private key to encrypt only the message digest. The encrypted message digest becomes the digital signature and is attached to the original data. To verify the contents of digitally signed data, the recipient generates a new message digest from the data that was received, decrypts the original message digest with the originator's public key, and compares the decrypted digest with the newly generated digest. If the two digests match, the integrity of the message is verified. The identity of the originator also is confirmed because the public key can decrypt only data that has been encrypted with the corresponding private key.

IPsec over FCIP

7500 extension switches and FR4-81i blades use FCIP protocol to IP to carry Fibre Channel traffic over IP networks. IPsec can be used to secure the IP flows over an FCIP tunnel.

At a high level, the steps to take are:

- Access the IPsec Policies dialog box.
- Create an IKE policy for authentication.
- Create a security association (SA).
- Create an SA proposal.
- Add an IPsec Transform policy, referencing the IKE policy and the SA proposal.
- Add an IPsec selector that allows you to apply a Transform policy to a specific IP flow.
- Enable the policy.

Accessing the IPsec Policies dialog box

Use the following procedure to access the **IPsec Policies** dialog box.

1. Open the **Switch Administration** window.

2. Select **Show Advanced Mode**.

3. Select the **Security Policies** tab.

4. Under **Security Policies**, select **IPsec Policies**.

   The **IPsec Policies** window displays. The default view shows the **IKE** tab.

Establishing an IKE policy for an FCIP tunnel

Use the following procedure to establish an IKE policy for an FCIP tunnel.

1. From the **IKE** tab of the **IPsec Policies** screen, select **Create**.

   The **Add Policy** dialog box displays.

2. **Policy Type** provides a way to toggle between the IKE and IPsec **Add Policy** dialog box boxes.

   Make sure the **Policy Type** is set to IKE.

3. Assign a policy number.

   The **Policy Number** selector allows you to select a number between 1 and 32.

4. Select the **Encryption Algorithm** used in this policy.

   The choices are 3DES, AES-128, and AES_256.

5. Select an **Authentication Algorithm** for this policy.

   The choices are SHA-1, MD5, and AES-XCBC.

6. Turn **Perfect Forward Secrecy** on or off.

   The default is On. Perfect Forward Secrecy (PFS) provides additional security by means of a Diffie-Hellman shared secret value. With PFS, if one key is compromised, previous and subsequent keys are secure because they are not derived from previous keys.

7. Select a **Diffie-Hellman Group** association.

   The choices are 1 (modp768) and 14 (modp2048).

8. Set a **Security Association Lifetime** (in seconds).

   The Security Association Lifetime is a time value in seconds. When this timer expires, the security association (SA) is rekeyed. This limits the amount of time a given key is available to a potential attacker.

9. Click **OK**.

Establishing an IPsec policy for an FCIP tunnel

Use the following procedure to establish an IPsec policy for an FCIP tunnel.

1. Select the **IPsec** tab**.**

   The **IPsec Policies** window displays.

2. Select **Create**.

   An **Add Policy** dialog box displays.

3. **Policy Type** provides a way to toggle between the IKE and IPsec **Add Policy** dialog boxes.

   Make sure the **Policy Type** is set to IPSEC.

4. Assign a policy number.

   The **Policy Number** selector allows you to select a number between 1 and 32.

5. Select the **Encryption Algorithm** used in this policy.

   The choices are 3DES, AES-128, and AES_256.

6. Select an **Authentication Algorithm** for this policy.

   The choices are SHA-1, MD5, and AES-XCBC. The remaining three fields are grayed out. They apply only to IKE policies.

7. Click **OK**.

IPsec over management ports

IPsec can be applied to the management port on a switch or a CP blade to establish a secure connection between a PC or workstation and Web Tools. The connection can be used as a virtual private network (VPN) interface to Web Tools.

At a high level, the steps to take are:

- Access the Ethernet IPsec Policies dialog box.
- Enable IPsec.
- Create an IKE policy for authentication.
- Create an security association (SA).
- Create an SA proposal.
- Add a IPsec Transform policy, referencing the IKE policy and the SA proposal.
- Add an IPsec selector that allows you to apply a Transform policy to a specific IP flow.

Enabling the Ethernet IPsec policies

Use the following procedure to access the **Ethernet IPsec Policies** dialog box.

1. Open the **Switch Administration** window.

2. Select **Show Advanced Mode**.

3. Select the **Security Policies** tab.

4. Under **Security Policies**, select **Ethernet IPsec**.

   The Ethernet **IPsec Policies** screen displays.

5. Ethernet IPsec policies can be configured only after enabling IPsec by clicking the **Enable** button below the Ethernet **IPsec policies** table.

Establishing an IKE policy

When you establish an IKE policy, you identify a set of algorithms and authentication rules and parameters to use in a key exchange. Refer to the *Fabric OS Administrator's Guide* for details on IKE functionality.

Use the following procedure to establish an IKE policy.

1. Select the **IKE** tab on the **IPsec Policies** window for Ethernet IPsec.

   The **Add IKE Policy** dialog box displays.

2. Enter an **IKE Policy Name**.

3. Enter the IP address of the authentication partner in the **Peer IP Address** field.

4. Enter the switch's local identifier in the **Local Identifier** field.

   This is normally the IP address in IPv4 or IPv6 format, but it may also be a DNS name.

5. Enter the identifier of the remote peer switch in **Peer Identifier**.

   This is normally the IP address in IPv4 or IPv6 format, but it may also be a DNS name.

6. Select the **Encryption Algorithm** option.

7. Select the **Hash Algorithm** option.

8. Select the **PRF Algorithm** option.

9. Select the **DH Group Number** option.

10. Select the **Authentication Method** option.

11. If PSK is chosen as the authentication method, enter the name of the file that holds the pre-shared key in the **Pre-Shared Key filename** field.

12. If you are using an X.509 certificate for authentication, enter the appropriate file names in the **Public Key filename**, **Private Key filename**, and **Peer Public Key filename** fields in PEM format.

13. Use the **PFS** selector to turn Perfect Forward Secrecy (PFS) on or off.

   PFS provides additional security by means of a Diffie-Hellman shared secret value. With PFS, if one key is compromised, previous and subsequent keys are secure because they are not derived from previous keys.

Creating a security association

A security association (SA) describes a set of parameters for providing secure communications between two endpoints.

Use the following procedure to create a security association.

1. Select the **IPsec** tab.

   The **IPsec Policies** screen displays.

2. Select the **SA** tab.

3. Select **Add**.

   The **Add SA** dialog box displays.

4. Enter a name for the SA in the **SA Name** field.

5. Select the **IPsec Protocol**. option.

   The choices are ah (for authentication header) and esp (for encapsulated security protocol).

6. Select the **Authentication Algorithm** option.

7. Select the **Encryption Algorithm** option.

8. Optionally, enter a value in the **SPI number** field.

   A Security Parameter Index (SPI) number is automatically assigned, but may be manually overridden.

9. Click **OK**.

Creating an SA proposal

An SA proposal is sent from one endpoint to another to negotiate IKE and IPsec policies. An SA proposal contains one or more security associations (SA). The endpoints must find a match for each of the following in the SAs sent in the SA proposal:

- The IKE authentication method.
- The IKE encryption algorithm.
- The IKE hash algorithm.
- The Diffie-Hellman group number.
- The IKE SA lifetime.
- The IP addresses of the endpoints.
- The IPsec protocol (AH or ESP).
- The IPsec Transform policy.

Use the following procedure to create an SA proposal.

1. Select the **SA Proposal** tab on the **IPsec Policies** screen.

2. Select **Add**.

   The **Add-SA Proposal** dialog box displays.

3. Enter a name in the **SA Proposal Name** field.

4. Enter the SAs in the **SA(s) to use** field.

5. Optionally, define SA lifetime parameters.

   The SA lifetime may be defined as a time value in seconds (**LifeTime in seconds**), as the number of bytes transmitted before the SA is rekeyed (**LifeTime in bytes**), or both. When both are used, the SA lifetime is determined by the threshold that is first reached.

6. Click **OK**.

Adding an IPsec transform policy

The IPsec transform policy is the combination of protocols and algorithms applied to a flow of IP packets. IPsec unidirectional, and policies need to be applied to both inbound and outbound flows.

Part of adding an IPsec transform policy is to select an IPsec Protection Type. The choices are discard, bypass, and protect:

- Discard causes data packets to be rejected if there is an invalid pair of source and destination addresses or invalid port addresses.
- Bypass allows a data packet to be transmitted or received without IPsec protection.

- Process indicates a data packet is processed using IPsec encryption, IKE authentication, or both, using encapsulation security protocol (ESP) processing, or authentication header (AH) protocol processing.

Use the following procedure to add an IPsec transform policy.

1. Select the **Transforms** tab.

   The **Transforms** screen displays.

2. Select **Add**.

   The **Add Transform** dialog box displays.

3. Enter a name in the **Transform Name** field.

4. Select the **IPsec Mode**.

   The choices are **Transport** or **Tunnel**.

5. Enter the **SA Proposal** name.

6. Select the **IPsec Protection Type** option.

7. Select the **IKE Policy Name** option.

   IKE policies need to be created before adding a transform policy. If there are no names to select from, you must create an IKE policy.

8. *Optional:* Enter a local and peer IP address.

9. Click **OK**.

Adding an IPsec selector

Selectors are used to apply transform policies to an IP flow. Flows are uni-directional. Selectors are associated with a specific source IP address, a specific peer IP address, and a specific transform.

1. Select the **Selectors** tab.

   The **Selectors** screen displays.

2. Select **Add**.

   The **Add Selector** dialog box displays.

3. Enter a name in the **Selector Name** field.

4. Select the **Traffic Flow Direction** (in or out).

   IPsec policies are unidirectional, and must be applied separately to inbound and outbound flows.

5. Enter the IP address of the sender in the **Source IP Address** field.

6. Enter the IP address of the receiver in the **Peer IP Address** field.

7. Enter the **Transform Name** value.

8. The **Protocol Name** selector allows you to select a specific protocol.

9. Click **OK**.

Manually creating an SA

Part of manually creating an SA is to select an IPsec Protection Type. The choices are discard, bypass, and protect:

- Discard causes data packets to be rejected if there is an invalid pair of source and destination addresses or invalid port addresses.

- Bypass allows a data packet to be transmitted or received without IPsec protection.

- Process indicates a data packet is processed using IPsec encryption, IKE authentication, or both, using encapsulation security protocol (ESP) processing, or authentication header (AH) protocol processing.

Use the following procedure to manually create a security association (SA).

1. Select the **SA(Manual)** tab.

2. Select **Add**.

   The **Add Manual-SA** dialog box displays.

3. Enter a security parameter index number in the **SPI (Hexadecimal)** field.

   The SPI must be manually applied when manually adding an SA.

4. Enter the IP address of the endpoint that sends the SA in the **Source IP Address** field.

5. Enter the IP address of the endpoint that receives the SA in the **Peer IP Address** field.

6. Select the protocol used to carry the transmission using the **Protocol Name** selector.

7. Select the **Traffic Flow Direction** (in or out).

   IPsec policies are unidirectional, and must be applied separately to inbound and outbound flows.

   - For the flow from peer to source, select **in**.

   - For the flow from source to peer select **out**.

8. Select the **IPsec Mode**.

   The choices are **Transport** or **Tunnel.** Refer to "Transport mode and tunnel mode" on page 229 if you are unfamiliar with Transport and Tunnel modes.

9. Select the **IPsec Protocol**.

   The choices are **ah** (for authentication header) and **esp** (for encapsulated security protocol).

10. Select the **IPsec Protection Type** option.

11. Select the **Authentication Algorithm** option.

12. Enter or copy a generated encryption key in the **Encryption Key** field.

13. Select the **Encryption Algorithm**.

14. Enter or copy a generated authentication key in the **Authentication Key** field.

15. *Optional:* Enter a local and peer tunnel IP address.

16. Click **OK**.

Editing an IKE or IPsec policy

An existing IKE or IPsec policy can be edited.

Use the following procedure to edit an IKE or IPsec policy.

1. Open the **Switch Administration** window.

2. Click **Show Advanced Mode**.

3. Select the **Security Policies** tab.

4. Under **Security Policies**, select **Ethernet IPsec** or **Ethernet IPsec**.

5. Select the policy you want to edit.

6. Select **Edit**.

    An **Edit Policy** dialog box displays.

7. Edit the policy as needed.

8. Click **OK**.

Deleting an IKE or IPsec policy

You can delete one or more IKE or IPsec policies.

Use the following procedure to delete an IKE or IPsec policy.

1. Open the **Switch Administration** window.

2. Select **Show Advanced Mode**.

3. Select the **Security Policies** tab.

4. Under **Security Policies**, select **Ethernet IPsec** or **Ethernet IPsec**.

5. Select the policy or policies you want to delete.

6. Select **Delete**.

    The policy is deleted from the SA database (SADB), and is removed from the list.

Establishing authentication policies for HBAs

Perform the following procedure to establish and enable authentication policies for HBAs as the log in to a fabric.

1. Open the **Switch Administration** window.

2. Click **Show Advanced Mode**.

3. Select the **Security Policies** tab.

4. Select **Authentication** under Security Policies.

    The **Authentication Policy Settings** screen displays.

5. Under **Configure Authentication Policy**, do the following.

- Select the **Authentication Type**. The choices are FCAP, DHCHAP, or both.
- Select the **Switch Authentication Policy Mode**. The choices are Passive, Active, On, or Off.
- Select the **Hash Type** used. The choices are SHA1, MD5, or both.
- Select the **DH-Group Type**. The choices are 0 (no DH authentication), 1 (1024 bit), 2 (1280 bit), 3 (1536 bit), or 4 (2048 bit).
- Use the **Device Authentication Policy Mode** selector to set the desired mode. The choices are On, Off, or Passive.
- Click **Apply**.

6. If your authentication method uses a shared secret, select the **Shared Secret Keys** tab.

   The **Shared Secret Keys** screen displays.

7. Select **Add**.

   The **Add Shared Secret Keys** dialog box displays.

8. Browse to select the switch WWN or name and domain ID, or enter the switch WWN or name and domain ID in the **Switch WWN: Name/Domain ID** field.

9. Enter the shared secret key for the peer device (an HBA in this case) in the **Peer Shared Secret** and **Confirm Peer Shared Secret** fields.

10. Enter the shared secret for switch in the **Local Shared Secret** and **Confirm Local Shared Secret** fields.

11. Click **Add**.

   An entry is added in the **Switch WWN** box.

12. Click **OK**.

13. Add more shared secrets, if needed.

Administering FICON CUP Fabrics

In this chapter

FICON CUP fabrics overview

Control Unit Port (CUP) is a protocol for managing FICON directors. Host-based management programs manage the switches using CUP protocol by sending commands to the emulated control device implemented by Fabric OS. A Brocade switch or director that supports CUP can be controlled by one or more host-based management programs or director consoles, such as Brocade Web Tools or Brocade Fabric Manager. (Refer to the *Fabric Manager Administrator's Guide* for information about Fabric Manager.) The director allows control to be shared between host-based management programs and director consoles.

NOTE
To perform FICON CUP operations, the EGM license must be enabled on the switches using the CUP protocol. Also, the EGM license must be enabled to set Prohibit Dynamic Connectivity Mask (PDCM) parameters.

To use FICON CUP, you must do the following:

• Install a FICON CUP license on a FICON director.

• Enable FICON Management Server (FMS) mode on the FICON director.

• Configure CUP attributes (FMS parameters) for the FICON director.

You can use Web Tools for all of these tasks. You can also use Web Tools to manage FICON directors (when FMS mode is enabled on those directors) to do the following:

• Display the control device state

• Display a code page

• Manage port connectivity configuration

You do not need to install the FICON CUP license to perform FICON CUP management; you *must* install the FICON CUP license, however, if your switch is to enforce traffic between the FICON director and the host-based management program.

Enabling port-based routing

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

Use the following procedure to enable port-based routing.

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

2. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

3. Click the **FICON CUP** tab.

    If the EGM license is not installed on the switches using CUP protocol, access to this feature is denied and an error message displays. If the EGM license is enabled, the **FICON CUP** tab is available.

4. Click **Enable** in the **FICON Management Server Mode** section to enable the port-based routing policy, or click **Disable** to disable port-based routing.

5. Click **Apply** to save your changes.

Enabling or disabling FICON Management Server mode

FICON Management Server (FMS) is used to support switch management using CUP. To be able to use the CUP functionality, all switches in the fabric must have FICON Management Server mode (FMS mode) enabled. FMS mode is a per-switch setting. After FMS mode is enabled, you can activate a CUP license without restarting the director. You can use Web Tools to install a CUP license. For more information on installing licenses, refer to "Activating a license on a switch" on page 47.

When FMS mode is disabled, mainframe management applications, director consoles, or alternate managers cannot communicate with a director with CUP. In addition, when FMS mode is disabled on a director, you cannot configure CUP attributes.

Use the following procedure to enable or disable FICON Management Server.

1. Select a FICON CUP-capable switch from the **Fabric Tree**.

2. Open the **Switch Administration** window as described in "Opening the Switch Administration window" on page 35.

3. Click **Show Advanced Mode**.

4. Select the **FICON CUP** tab.

    The FICON CUP tabbed page displays the **FICON Management Server** page. All attributes on this tab are disabled until FMS mode is enabled.

5. Click **Enable** in the **FICON Management Server Mode** section to enable FMS mode or click **Disable** to disable FMS mode.

6. Click **Apply** to save your changes.

FMS parameter configuration

FMS parameters control the behavior of the switch with respect to CUP itself, as well as the behavior of other management interfaces (director console, Alternate Managers). You can configure FMS parameters for a switch *only* after FMS mode is enabled on the switch. All FMS parameter settings are persistent across switch power cycles. There are six FMS parameters, as described in Table 22.

TABLE 22   FMS mode parameter descriptions

Parameter	Description
Programmed Offline State Control	Controls whether host programming is allowed to set the switch offline. The parameter is set as enabled by the hardware after system installation, and can be reset by Web Tools.
Active=Saved Mode	Controls the IPL file update. The IPL file saves port connectivity attributes and port names. After a switch restart or power cycle, the switch reads the IPL file and actives its contents as default configuration.
	When this mode is enabled, activating a configuration saves a copy to the IPL configuration file. All changes made to the active connectivity attributes or port names by host programming or alternate managers are saved in this IPL file. It keeps the current active configuration persistent across switch restarts and power cycles.
	You cannot directly modify the IPL file or save a file as an IPL file. When this mode is disabled, the IPL file is not altered for either new configuration activation or any changes made on the current active configuration. This parameter is set as enabled by the hardware after system installation, and can be reset by Web Tools.
	Note: When FMS mode is enabled and the Active=Saved parameter is disabled, you can enable and disable ports, but the setting is not persistent. When the Active=Saved parameter is enabled, you can enable and disable ports and the setting *is* persistent.
Alternate Control Prohibited	Determines whether alternate managers are allowed to modify port connectivity.
	Enabling this mode prohibits alternate manager control of port connectivity; otherwise, alternate managers can manage port connectivity.
	This parameter is set as enabled by the hardware after system installation, and can be reset by Web Tools.
User Alert Mode	Controls director console behavior for alerts.
	Enabling this mode prompts the director consoles to display a warning whenever you attempt an action that changes switch parameters. When you disable this mode, no warning is displayed. In this case, in which Web Tools is the director console, warning messages are displayed by Web Tools regardless of the setting of the parameter, since Web Tools always displays warning messages when you apply a change to a switch that changes parameters.
	This parameter is always read-only in Web Tools. Each time that the switch is powered on, the parameter is reset to disabled.

TABLE 22 FMS mode parameter descriptions (Continued)

Parameter	Description
Director Clock Alert Mode	Controls behavior for attempts to set the switch timestamp clock through the director console. When it is enabled, the director console (Web Tools, in this case) displays warning indications when the switch timestamp is changed by a user application. When it is disabled, you can activate a function to automatically set the timestamp clock. There is no indication for timestamp clock setting.
	This parameter is set as disabled by the hardware after system installation, and can be reset by Web Tools.
Host Control Prohibited	Determines whether host programming allows modifying port connectivity.
	Enabling this mode prohibits host programming control of port connectivity; otherwise, host programming can manage port connectivity.
	This parameter is set as disabled by the hardware after system installation. and can be reset by Web Tools.

Configuring FMS mode parameters

Use the following procedure to configure FMS mode parameters.

1. Select a FICON-enabled switch from the **Fabric Tree**.

2. Open the **Switch Administration** window as described in “Opening the Switch Administration window” on page 35.

3. Select the **FICON CUP** tab.

   The **FICON CUP** page displays the **FICON Management Server** page. All attributes on this page are read-only until FMS mode is enabled.

4. To enable or disable an FMS mode parameter, click the check box next to the parameter.

   A checked check box indicates that the parameter is enabled. You cannot configure the **User Alert Mode** parameter in Web Tools, as it is read-only.

Displaying code page information

The **Code Page** section identifies the language used to exchange information between the FICON director and Host Programming. It is a read-only field in Web Tools, as it is set by Host Programming only. When FMS mode is disabled, the code page is displayed as unavailable.

Use the following procedure to display code page information.

1. Select a FICON-enabled switch from the **Fabric Tree**.

2. Open the **Switch Administration** window as described in “Opening the Switch Administration window” on page 35.

3. Select the **FICON CUP** tab.

   The **FICON CUP** page displays the **FICON Management Server** page in front. All attributes on this tab are read-only until FMS mode is enabled.

   The code page format is displayed in the **Code Page** section as shown in the following example:

   ```
   Language used to exchange information with Host Programming: (EBCDIC)
   USA/Canada -- 00037
   ```

Viewing the control device state

The control device is in either a neutral or a switched state. When it is neutral, the control device accepts commands from any channel that has established a logic path with it and accepts commands from alternate managers. When the control device is switched, it establishes a logical path and accepts commands only from that logical path ("device allegiance"). Commands from other paths cause a FICON CUP Busy Error. Most "write" operations from alternate managers are also rejected.

Device allegiance usually lasts for a very short time. However, under abnormal conditions, device allegiance can get "stuck" and fail to terminate. It might cause the switch to be unmanageable with CUP, and you will continue to receive the FICON CUP Busy Error. In this case, you should check the control device state and the last update time to identify if the device allegiance is stuck. The Web Tools **Switch Administration** window displays the control device state and last update time. You can click **Refresh** to get most recent update.

NOTE
You can manually reset allegiance to bring the control device back to the neutral state by clicking **Reset Allegiance** in the **FICON CUP Busy Error** dialog box.

The following switch parameters being read or modified can cause the FICON CUP Busy error:

- Mode Register
- Port Names (also called Port Address Name)
- PDCM and Port Connectivity Attributes
- Switch enable/disable
- Switch name change

Use the following procedure to access the FICON CUP tab.

1. Select a FICON-enabled switch from the **Fabric Tree**.

2. Open the **Switch Administration** window as described in .

3. Select the **FICON CUP** tab.

The FICON CUP tabbed page displays, with the FICON Management Server subtabbed page in front. All attributes on this tab are read-only until FMS Mode is enabled.

The control device state is displayed as neutral or switched in the Control Device Allegiance field.

NOTE
If FMS mode is enabled, and the control device state is unavailable, the FICON CUP Busy Error is displayed. Click **Reset Allegiance** in the error message to reset the control device state to its correct state.

CUP port connectivity configuration

In the Port Connectivity subpanel, you can manage the configuration files and active configuration. All CUP configuration files and the active configuration are listed in a table. The active configuration is listed as "Active Configuration*" and the description in the table is "Current active configuration on switch." The other special configuration file is the IPL. Any other files displayed are user-defined configurations and are stored on the switch.

You can create, activate, copy, or delete saved CUP port connectivity configurations; however, you can only edit or copy a configuration while it is active.You can also activate, edit, or copy the IPL configuration. You must have FMS mode enabled before you can make any changes to the configurations. Click **Refresh** to get the latest configuration file list from the switch.

When creating a new configuration or editing an existing configuration, the Web Tools port name is restricted to printable ASCII characters. Characters beyond printable ASCII characters are displayed as dots.

When initially installed, a switch allows any port to dynamically communicate with any other port. Two connectivity attributes are defined to restrict this any-to-any capability for external ports: *Block* and *Prohibit*.

Block is a port connectivity attribute that prevents all communication through a port. Prohibit is the port connectivity attribute that prohibits or allows dynamic communication between ports when a port is not blocked. Each port has a vector specifying its Prohibit attribute with respect to each of the other ports in the switch. This attribute is always set symmetrically in that a pair of ports is either prohibited or allowed to communicate dynamically.

The Port Connectivity table (shown in) displays the Port number (in physical-location format), Port Name (port address name), Block attribute, Prohibit attribute, and Area Id (port address, displayed in hexadecimal) in fixed columns. The right side is a port matrix, that lists all ports by Area ID and identifies prohibited ports. Those columns are scrollable and swappable.

Viewing CUP port connectivity configurations

Use the following procedure to display a list of CUP port connectivity configurations.

1. Select a FICON-enabled switch from the **Fabric Tree**.

2. Select **Tasks** > **Manage** > **Switch Admin**.

3. Click **Show Advanced Mode** to see all the available tabs and options.

4. Select the **FICON CUP** tab.

   The **FICON CUP** page displays the **FICON Management Server** page in front. All attributes on this page are read-only until FMS mode is enabled.

5. Click the **CUP Port Connectivity** subtab.

Modifying CUP port connectivity configurations

In the **Edit Port CUP Connectivity Configuration** dialog box, swapped ports are indicated with the "(Swapped)" label (Figure 63).

FIGURE 63   Edit Port CUP Connectivity dialog box swapped label

Use the following procedure to create a new CUP port connectivity configuration or to edit an existing configuration.

1. Display the CUP port connectivity configuration list.

2. You can either create a new configuration or edit an existing configuration:

    - To create a new configuration, click **New**.

      The **Create Port CUP Connectivity Configuration** dialog box displays all ports and port names on the selected switch (similar to the dialog box shown in Figure 64). The **Block** column, **Prohibit** column, and prohibited ports matrix are displayed as empty, for you to configure.

    - To edit an existing configuration, click the configuration, and then click **Edit**.

      The **Edit Port CUP Connectivity Configuration** dialog box displays the content of the selected configuration from the switch in a table format (Figure 64).

3. *Optional*: Select the check box corresponding to a port you want to block on the **Block** column. Repeat this step for all ports you want to block. Select the **Block All** check box to block all ports.

4. *Optional*: Select the check box corresponding to a port you want to prohibit on the **Prohibit** column. Repeat this step for all ports you want to prohibit. Select the **Prohibit All** check box to prohibit all ports.

    The cells in the matrix are updated with "X" icons to identify prohibited ports.

    FE and FF ports are in the disabled state in FMS mode.

5. *Optional*: Click the individual cells corresponding to the combination of ports you want to prohibit. You cannot prohibit a port to itself.

    If you prohibit E_Port, E-E connection, or E-F connection, a warning message is displayed, "You have placed a prohibit on an E-Port. This has no effect for Fabric OS based fabrics".

6. Review your changes. A blue background in a cell indicates that its value has been modified.

7. After you have finished making changes, do any of the following:

- Click **Activate** to save the changes and make the configuration active immediately, as described in *"Activating a CUP port connectivity configuration"* on page 250.

- Click **Save** to save the changes but not make the configuration active.

- Click **Save As** to save the configuration to a new configuration file. When you click **Save As,** a dialog box displays in which you should enter a file name and description for the configuration file.

- Click **Refresh** to refresh the information from the switch.

- Click **Cancel** to cancel all changes without saving.
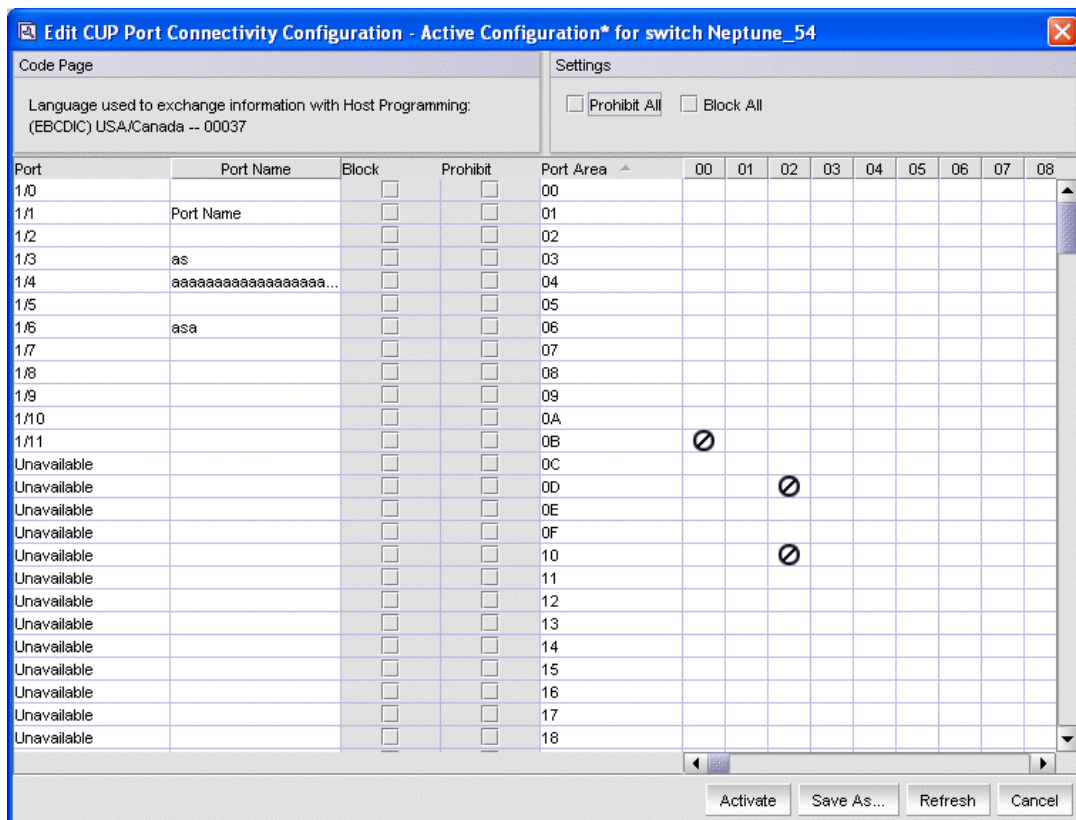
FIGURE 64   Port CUP Connectivity Configuration dialog box

Activating a CUP port connectivity configuration

When you activate a saved CUP port connectivity configuration on the switch, the preceding configuration (currently activated) is overwritten.

Use the following procedure to activate a CUP port connectivity configuration.

1. Open the CUP port connectivity configuration list.

2. Select the saved configuration from the list.

3. Click **Activate**.

   The **Activate CUP Port Connectivity Configuration** confirmation dialog box displays. The message reminds you that the current configuration will be overwritten upon activation.

4. *Optional*: Click **Active=Saved Mode** to enable (selected) or disable (not selected) the **Active=Saved FMS** parameter after the configuration is activated.

5. Click **Yes** to activate the configuration or click **No** to cancel the activation.

Copying a CUP port connectivity configuration

Use the following procedure to copy a CUP port connectivity configuration to a new configuration.

1. Display the CUP port connectivity configuration list.

2. Select a saved configuration or the active configuration from the list.

3. Click **Copy**.

   The **Copy CUP Port Connectivity Configuration** dialog box displays.

4. In the dialog box, enter a name and description for the new configuration and click **OK** to save the configuration to the target file; click **Cancel** to cancel copying the configuration.

   The file name must be in alphanumeric characters and can contain only dashes or underscores as special characters.

Deleting a CUP port connectivity configuration

Use the following procedure to delete a saved CUP port connectivity configuration.

1. Display the CUP port connectivity configuration list.

2. Select the saved configuration from the list.

3. Click **Delete**.

   The **Delete CUP Port Connectivity Configuration** confirmation dialog box displays.

4. Click **Yes** to delete the selected configuration; click **No** to cancel the deletion.

CUP logical path configuration

The logical reporting path is a CUP mechanism for sending FRU-failure type reports to a FICON Logical Path via the FICON Protocol.

Viewing CUP logical path configurations

Use the following procedure to display a list of CUP logical path configurations.

1. Select a FICON-enabled switch from the **Fabric Tree**.

2. Select **Tasks** > **Manage** > **Switch Admin**.

3. Click **Show Advanced Mode** to see all the available tabs and options.

4. Select the **FICON CUP** tab.

   The **FICON CUP** page displays the **FICON Management Server** page in front. All attributes on this page are read-only until FMS mode is enabled.

5. Click the **CUP Logical Paths** subtab.

Configuring CUP logical paths

Use the following procedure to configure a CUP logical path.

1. Select a FICON-enabled switch from the **Fabric Tree**.

2. Select **Tasks** > **Manage** > **Switch Admin**.

3. Click **Show Advanced Mode** to see all the available tabs and options.

4. Select the **FICON CUP** tab.

   The **FICON CUP** page displays the **FICON Management Server** page in front. All attributes on this page are read-only until FMS mode is enabled.

5. Click the **CUP Logical paths** subtab.

6. Select a logical path and click **Set Current**.

Link Incident Registered Recipient configuration

The Link Incident Registered Recipient (LIRR) receives Link Incident Reports (RLIR) on the source N_Port. The LIRR database is stored on the switch.

Viewing Link Incident Registered Recipient configurations

Use the following procedure to display a list of Link Incident Registered Recipient (LIRR) configurations.

1. Select a FICON-enabled switch from the **Fabric Tree**.

2. Select **Tasks** > **Manage** > **Switch Admin**.

3. Click **Show Advanced Mode** to see all the available tabs and options.

4. Select the **FICON CUP** tab.

   The **FICON CUP** page displays the **FICON Management Server** page in front. All attributes on this page are read-only until FMS mode is enabled.

5. Click the **Link Incident Registered Recipient** subtab.

Configuring LIRRs

Use the following procedure to configure the Link Incident Registered Recipients (LIRR).

1. Select a FICON-enabled switch from the **Fabric Tree**.

2. Select **Tasks** > **Manage** > **Switch Admin**.

3. Click **Show Advanced Mode** to see all the available tabs and options.

4. Select the **FICON CUP** tab.

   The **FICON CUP** page displays the **FICON Management Server** page in front. All attributes on this page are read-only until FMS mode is enabled.

5. Click the **Link Incident Registered Recipient** subtab.

6. Click **Set LIRR Port**.

   The **Set LIRR Port** dialog box displays.

7. Enter a port number and click **OK**.

8. Click **Close**.

Displaying Request Node Identification Data

Web Tools displays Request Node Identification Data (RNID) information for the local switch, and for attached FICON devices and FICON channel paths. RNID information for the switch displays in the **Switch Information** tab (Figure 65).
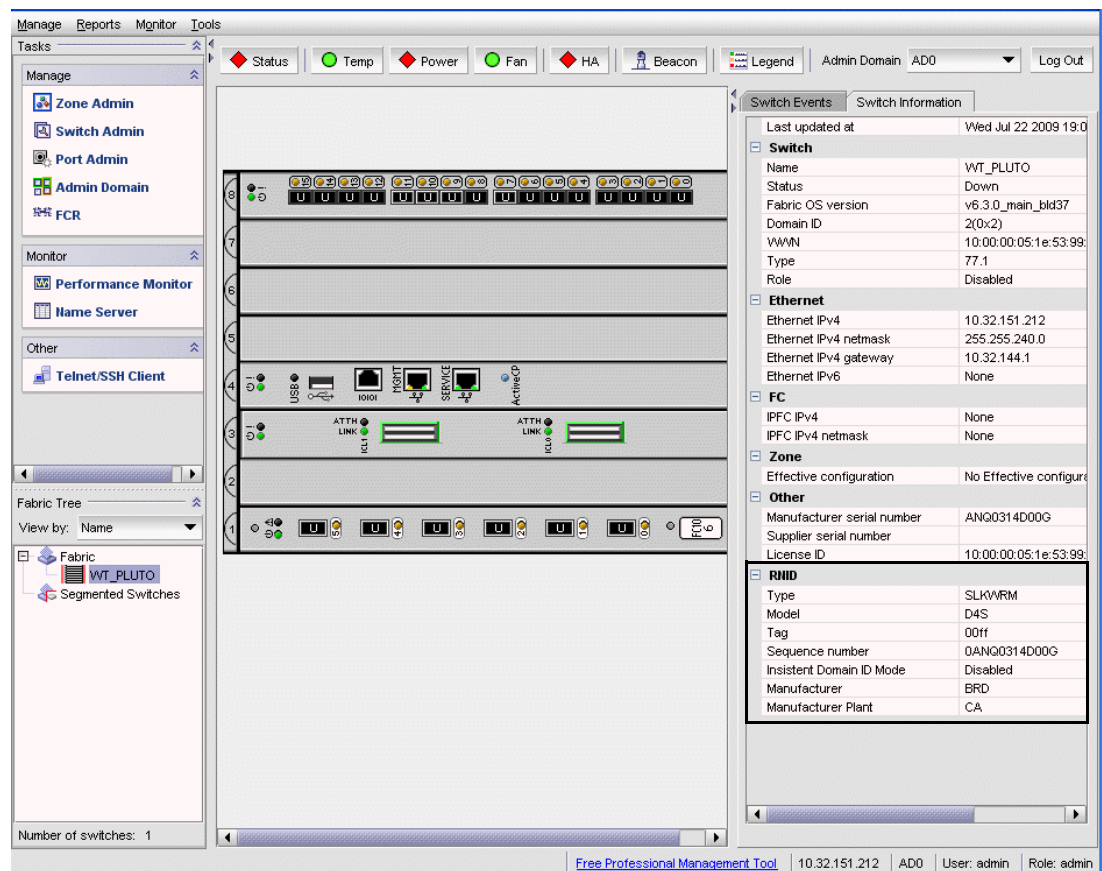

FIGURE 65   Switch RNID information

RNID information for attached FICON devices and channel paths displays on the **Name Server** view. To view this information, Click **Name Server** to display the **Name Server** view. Ports that completed an RNID exchange display **FICON** in the **Capability** column. For those ports, the following information specific to RNID displays in the following columns:

- **Device Type**
- **Model**
- **Manufacturer**
- **Manufacturer Plant**
- **Unit Type**
- **Tag**

Configuring FCoE with Web Tools

In this chapter

Web Tools and FCoE overview

Brocade Web Tools is an embedded graphical user interface (GUI) that enables administrators to monitor and manage single or small fabrics, switches, and ports. Web Tools is launched directly from a web browser, or from the DCFM client.

NOTE
For complete information on Web Tools, refer to the *Web Tools Administrator's Guide*. This chapter only discusses Web Tools and FCoE configuration.

A limited set of features is accessible using Web Tools without a license, and is available free of charge. Additional switch management features are accessible using Web Tools with the Enhanced Group Management (EGM) license. Refer to "Web Tools, the EGM license, and DCFM" for more information.

A new view has been added to Web Tools for the Brocade 8000 switch and FCOE10-24 DCX blade, and new tabs have been added to the **Switch Administration** panel and the **Port Administration** panel to support FCoE interfaces and trunks.

Web Tools, the EGM license, and DCFM

Beginning with Fabric OS version 6.1.1, Web Tools functionality is tiered and integrated with DCFM. If you are migrating from a Web Tools release prior to Fabric OS version 6.1.1, this may impact how you use Web Tools.

A Web Tools license is not required, and a basic version of Web Tools is available for free. Additional functionality may be added by obtaining the Enhanced Group Management (EGM) license. The EGM license is only for 8 Gbps platforms, such as the Brocade DCX and DCX-4S enterprise-class platforms, the Encryption Switch, and the 300, 5100, and 5300 switches. For non-8 Gbps platforms, all functionalities are available without the EGM license.

Port information that is unique to FCoE

The **General** tab of the **Port Administration** panel displays several parameters that are unique to CEE/DCE interfaces:

- **Interface Mode**—The interface mode vales are either **None** or **L2** mode.
- **VLAN ID**—The VLANs that carry traffic on the links are attached to this port.
- **LAG**—The ID of the Link Aggregation Group (LAG), with which this port is associated. If no ID is specified, the port is not associated with any LAG.
- **L2 Mode**—The values are **Access**, **Trunk**, or **Converged**. **Access** mode allows only one VLAN association, and allows only untagged frames. **Trunk** mode allows more than one VLAN association, and allows tagged frames. **Converged** mode interface can be native (untagged or access) in one VLAN and it could be non-native (trunk or tagged) type in more than one VLAN.
- **CEE Map**—The name of a CEE map that was created and associated with the port.
- **Traffic Class Map**—The name of a traffic class map that was created and associated with the port.
- **LLDP Status**—Indicates whether LLDP is active or inactive.
- **LLDP Profile**—The name of an LLDP profile that was created and associated with the port.

- **FCoE Priority Bits**—Each bit represents a user priority that is associated with FCoE traffic.
- **Default CoS**—The default Class of Service.

Switch administration and FCOE

The **CEE** tab on the **Switch Administration** panel is specific to DCE and CEE configuration and management. The CEE tab has five subtabs (Figure 66) that are used for FCoE switch administration:

- **Link Aggregation**
- **VLAN**
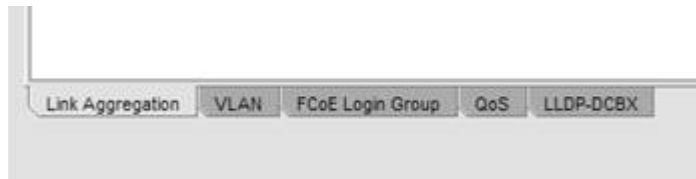- **FCoE Login Group**
- **QoS**
- **LLDP-DCBX**

FIGURE 66   Switch Administration CEE subtabs

FCOE configuration tasks

There are several tasks related to FCOE configuration. The following list describes the high level tasks in a suggested order:

- Quality of Service (QoS) configuration (optional)—If you intend to implement a specific QoS scheme to prioritize data traffic, it is recommended that you finish your QoS configuration before you begin port configuration. QoS values are referenced when you configure ports.

- LLDP-DCBX configuration (optional)—If you intend to implement DCBX, it is recommended that you finish LLDP-DCBX configuration before you configure ports. LLDP-DCBX values are referenced when you configure ports.

- CEE interface configuration (mandatory).

- Link Aggregation Group (LAG) configuration (mandatory)—Ports must be configured before they can be placed into a LAG. The parameters applied to the LAG reflects on each port that is member of the LAG.

- VLAN configuration (optional)—Port and LAG names are referenced in VLAN configuration, and must be defined before you can successfully complete a VLAN configuration.

- Login group configuration (optional)—Login group configuration is not dependent on any of the above configurations. It can be done as a separate task.

Quality of Service configuration

As a general concept, Quality of Service (QoS) is a mechanism for classifying and scheduling data traffic based on priority settings. QoS can be used to control traffic congestion, allocate bandwidth, and carry data traffic with different characteristics over a common interface.

The following two configuration options are available:

- You can create a CEE map. A CEE map defines priority and priority group tables that support Enhanced Transmission Selection (ETS). ETS allows allocation of bandwidth to different traffic classes. CEE maps also allow you to enable Priority Flow Control (PFC).

- You can create a traffic class map. A traffic class map can be used to map a specific class of traffic to a specific Class of Service (CoS).

Adding a CEE map

A CEE map defines priority and priority group tables that support Enhanced Transmission Selection (ETS). ETS allows bandwidth to be allocated based on priority settings through an exchange of priority group tables.

Use the following procedure to add a CEE map.

1. Select the **CEE** tab on the **Switch Administration** panel.

2. Select the **QoS** tab.

3. Select the **CEE Map** tab.

4. Click **Add**.

   The **CEE Map Configuration** dialog box displays.

5. Enter a name for the CEE map in the **Name** field.

6. Enter a precedence value in the **Precedence** field.

   The value is specified as a number. The allowable range is 1 to 100. The default is 1.

   The precedence value controls QoS scheduling policies. If different CEE maps have conflicting policies, the scheduler gives precedence to the CEE map with the highest precedence value (the highest number).

   When the **CEE Map Configuration** dialog box displays, the default values shown in the **Priority Group Map** match the IEEE 802.1Q recommendation for systems supporting eight traffic classes. The **Priority Group Map** displays the Layer 2 **Cos** values mapped to **Priority Group ID** (PGID). PGID values are in the form `<policy>.<priority>`. A policy value of 15 indicates Priority values run from 0 (highest priority) to 7 (lowest priority). Note that this is contrary to the **CoS** values, which run from 7 (highest priority) to 0 (lowest priority).

7. Create a new priority group by clicking **Add** next to the **Priority Group** table.

   An entry is added to the **Priority Group** table.

   NOTE
   When you add an entry, a PGID is automatically assigned. The PGID is an integer from 0 to 7. The first added entry is given a PGID of 0, and the PGID increments by one for each additional added entry until a PGID of 7 is reached.

   ---

8.   Edit the **Bandwidth** entry to indicate the desired percentage of total bandwidth.

9.   Change the **Priority Flow Control Status** to **Enabled** to enable PFC for the entry.

10. Click **OK**.

The new priority group displays in the **Priority Group Map**.

Adding a traffic class map

CoS priorities can be mapped to traffic classes using a traffic class map.

1.   Select the **CEE** tab on the **Switch Administration** panel.

2.   Select the **QoS** tab.

3.   Select the **Traffic Class Map** tab.

4.   Select **Add**.

The **Traffic Class Map Configuration** dialog box displays.

NOTE
This dialog box has the same structure as the **Priority Group Map** in the **CEE Configuration** dialog box. The default CoS-to-traffic class structure is based on IEEE 802.1Q recommendations, as in the default Priority Group Map.

5.   Enter a name for the traffic class map in the **Name** field.

6.   Select the **Traffic Class** that you want to assign to the **CoS** priority.

7.   Click **OK**.

LLDP-DCBX configuration

Link Layer Discovery Protocol (LLDP) is a IEEE standard for collecting and distributing device information. Data Center Bridging Exchange (DCBX) extends LLDP by providing a protocol for discovering, initializing, and managing CEE-compliant devices.

There are two configuration procedures:

- Configuring global LLDP characteristics.
- Configuring an LLDP profile.

Configuring global LLDP characteristics

Configuring at the global level enables you to apply changes to every port.

Use the following procedure to configure the global LLDP characteristics.

1.   Select the **CEE** tab on the **Switch Administration** panel.

2.   Select the **LLDP-DCBX** tab.

3.   Select the **Global** tab.

4.   Select the **LLDP** check box to enable LLDP globally.

You can clear the check box to disable LLDP.

5. Enter a name for the configuration in the **System Name** field.

6. Optionally, add a description in the **System Description** field.

7. Select the **Mode**.

   For **Mode**, the choices are **Tx** (transmit), **Rx**, (receive) or **Both**. The default is **Both**.

8. In the **Hello** field, enter a time value in seconds.

   The Hello value sets the interval between hello bridge protocol data units sent by the root switch configuration messages. The range is 4 to 180 seconds. The default is 30 seconds.

9. In the **Multiplier** field, set the number of consecutive misses allowed before LLDP considers the interface to be down.

   The range is 1 to 10. The default is 4. The multiplier is related to the Hello time interval. Using the defaults, you wait four times (the multiplier value) at 30 second intervals (the hello value) before giving up on the interface.

10. In the **FCOE Priority Bits** field, enter a value that indicates the desired user priority. Each bit represents a user priority associated with FCoE traffic.

    The range is 0-255. The default is 8.

11. Select the parameters you want to exchange.

    Note that the term TLV indicates packaging of parameters into a Brocade-specific Type/Length/Value (TLV):

    - **Advertise Optional-tlv**—Advertises the following optional TLVs:
      - **system-description**—Describes switch or blade characteristics.
      - **port-description**—Describes the configured port.
      - **system-name**—Specifies the system name.
      - **system-capabilities**—Describes the system capabilities.
      - **management-address**—The IP address of the management port on the 8000 switch.
    - **Advertise dot1-tlv**—Select this check box to advertise to any attached device to send IEEE 802.1 LLDP type, length, and values.
    - **Advertise dot3-tlv**—Select this check box to advertise to any attached device to send IEEE 802.3 LLDP type, length, and values.
    - **Advertise dcbx-tlv**—Select this check box to advertise to any attached device the respective LLDP type, length, and values.
    - **Advertise dcbx-fcoe-logical-link**—Select this check box to advertise to any attached device to send DCBX protocol over LLDP to negotiate the logical link type, length, and values.
    - **Advertise dcbx-fcoe-app**—Select this check box to advertise application type, length, and values to ensure interoperability of traffic over DCBX protocol running over LLDP.

12. Click **Apply**.

13. Click **Save Configuration**.

Adding an LLDP profile

The LLDP profile determines LLDP settings per port.

Use the following procedure to add an LLDP profile.

1. Select the **CEE** tab on the **Switch Administration** panel.

2. Select the **LLDP-DCBX** tab.

3. Select the **LLDP Profile** tab.

4. Click **Add**.

   The **LLDP Configuration** dialog box displays.

5. Enter a name for the configuration in the **Name** field.

6. Optionally, add a description in the **Description** field.

7. Select the **Mode**.

   For **Mode**, the choices are **Tx** (transmit), **Rx**, (receive) or **Both**. The default is **Both**.

8. In the **Hello** field, enter a time value in seconds.

   The Hello value sets the interval between hello bridge protocol data units sent by the root switch configuration messages. The range is 4 to 180 seconds. The default is the global configuration range.

9. In the **Multiplier** field, set the number of consecutive misses allowed before LLDP considers the interface to be down.

   The range is 1 to 10. The default is the global configuration range. The multiplier is related to the Hello time interval. Using the defaults, you wait four times (the multiplier value) at 30 second intervals (the hello value) before giving up on the interface.

10. Select the parameters you want to exchange.

    Note that the term TLV indicates packaging of parameters into a Brocade-specific Type/Length/Value (TLV).

    - **Advertise Optional-tlv**—Advertises the following optional TLVs:
      - **system-description**—Describes switch or blade characteristics.
      - **port-description**—Describes the configured port.
      - **system-name**—Specifies the system name.
      - **system capabilities**—Describes the system capabilities.
      - **management-address**—The IP address of the management port on the 8000 switch.
    - **Advertise dot1-tlv**—Advertises to any attached device to send IEEE 802.1 LLDP type, length, and values.
    - **Advertise dot3-tlv**—Advertises to any attached device to send IEEE 802.3 LLDP type, length, and values.
    - **Advertise dcbx-tlv**—Advertises to any attached device the respective LLDP type, length, and values.
    - **Advertise dcbx-fcoe-logical-link**—Advertises to any attached device to send DCBX protocol over LLDP to negotiate the logical link type, length, and values.

- **Advertise dcbx-fcoe-app**—Advertises application type, length, and values to ensure interoperability of traffic over DCBX protocol running over LLDP.

11. Click **Save Configuration**.

Configuring CEE interfaces

The **CEE Interfaces** tab on the **Port Administration** panel is used for configuring the CEE interfaces on a switch.

Use the following procedure to configure the CEE interfaces.

1. Select the **CEE Interfaces** tab on the **Port Administration** panel.

2. Select the port you want to configure under the **CEE Interface Explorer**.

3. Select the **General** tab.

   Normally, this tab is pre-selected.

4. Select **Edit Configuration**.

   The **CEE Edit Configuration** dialog box displays.

5. Select the **Interface Mode**.

   The options are **None** and **L2**. The default is **None**.

   If you intend to use this port in a Link Aggregation Group (LAG), select **None**. L2 mode is applied when you configure the LAG.

6. Select the **L2 Mode**.

   The choices are **Access**, **Trunk,** and **Converged**. The default is **Access**.

   The L2 mode setting determines operation within a VLAN:

   - **Access** mode allows only one VLAN association, and all frames are untagged.
   - **Trunk** mode allows more that one VLAN association, and tagged frames are allowed.
   - **Converged** mode interface can be Native (untagged or access) in one VLAN and it could be non-native (trunk or tagged) type in another VLAN.

7. If you are using a **CEE map** or **Traffic Class Map** to apply QoS traffic priority, select the appropriate button, and enter the name of the map you want to use.

8. Enter the profile name in the **LLDP-DCBX Profile** field for using a specific profile for the interface.

9. In the **FCOE Priority Bits** field, enter a value that indicates the desired user priority. Each bit represents a user priority that is associated with FCoE traffic.

   The range is 0-255. The default is 8.

10. Assign a default class of service in the **Default CoS** field.

    The default CoS range is 0-7. The default is 0.

11. Click **OK**.

12. Click **Enable** for **Status** and **LLDP Status**.

Configuring a link aggregation group

FCoE ports can be grouped to create a link aggregation group (LAG). The LAG is treated as a single interface.

Use the following procedure to configure a LAG.

1. Select the **CEE Interfaces** tab on the **Switch Administration** panel.

2. Select the **Link Aggregation** tab.

3. Click **Add**.

   The **Add LAG Configuration** dialog box displays.

   NOTE
   Only ports that you defined with an **Interface Mode** of **None** can be a **LAG Member**.

   ---

4. Select the **Mode**.

   The choices are **Static** and **Dynamic**. **Static** mode does not use Link Aggregation Control Protocol (LACP) to negotiate and manage link aggregation. Link participation in the LAG is determined by the link's operational status and administrative state. **Dynamic** mode uses LACP. LACP allows partner systems to examine the attributes of the links that connect them and dynamically form a LAG. When you select **Dynamic** mode, the **Active** and **Passive** options are enabled:

   - If you choose **Active**, your switch initiates an exchange of LACP data units.
   - If you choose **Passive**, your switch waits to receive LACP data units from its partner system and then respond. **Passive** is the default behavior.

5. Select the **Type**.

   Type refers to the type of trunking used by the LAG. The choices are **Standard** and **Brocade**.

6. Select the **Interface Mode**.

   The options are **None** and **L2**. The default is **None**.

7. Select the **L2 Mode**.

   The L2 mode setting determines operation within a VLAN:

   - **Access** mode allows only one VLAN association, and all frames are untagged.
   - **Trunk** mode allows more than one VLAN association, and allows tagged frames.

8. Select the operational **Status**.

   The choices are **Administratively Up** and **Administratively Down**.

9. Click **OK**.

Configuring VLANs

The Virtual LAN (VLAN) capability allows multiple virtual LANs within a single physical LAN infrastructure. The physical interface must be configured as L2 prior to configuring a VLAN, either as an individual interface, or as a LAG. Before you start the VLAN configuration procedure, you need to know which interfaces or LAGs you want to associate with each VLAN.

Use the following procedure to configure a VLAN.

1. Select the **CEE** tab on the **Switch Administration** panel.

2. Select the **VLAN** tab.

3. Click **Add**.

   The **VLAN Configuration** dialog box displays.

4. Specify a VLAN ID.

   The format is VLAN<bridge number><ID>. In this Fabric OS release, no bridge instances are supported, so the bridge number is always 0, and the value under **Bridge** is statically defined as VLAN0. The <ID> is an integer from 1 to 3583, that must be entered in the **ID** field.

5. Select the **Native** check box.

6. Under the **Selection List**, click the plus sign (+) next to the **Interface** and **LAG** folders, and select individual interfaces and LAGs you want to associate with the VLAN ID.

7. Click **Add** to move the interfaces or LAGs to the **Selected List**.

   Note the reminder that interfaces must be configured as L2, and that the interfaces or LAGs must be in Trunk mode to be associated with multiple VLANs, Access mode interfaces can be associated with only one VLAN, and the Converged mode interface can be Native in one VLAN and it could be non-native type in more than one VLAN.

8. Click **OK**.

9. Repeat the procedure for additional VLANs.

10. To edit VLAN, select the detail from the table in the **VLAN** tab and click **Edit.**

    The **FCoE** check box is selected by default.

11. Click **OK** to enable FCoE. Clear the check box to disable FCoE.

Configuring FCoE login groups

FCoE login groups control which FCoE devices are allowed to log in to a switch or fabric.

The **FCoE Ports** window is used for configuring the FCoE ports on a switch.

Use the following procedure to configure an FCoE login group.

1. Select the **CEE** tab on the **Switch Administration** panel.

2. Select the **FCoE Login Group** tab.

3. Click **New**.

   The **New Login Group** dialog box displays.

4. Enter a name for the login group in the **Login Group Name** field.

5. Select the switch WWN.

   The choices are:

- **Self** — WWN of your current switch
- **Other Switch WWN**

  If you choose **Other Switch WWN**, you must enter the WWN of that switch in the provided field.

6. Under **Login Member Configuration**, select either **Allow All Members** or **Allow Specific Member.**

   - If you select **Allow All Members**, all devices attached to FCoE ports are allowed to log in to the switch.
   - If you select **Allow Specific Member**, you can control which devices can log in, using **Member Type**, **Member PWWN/MAC**, and the **Add** and **Remove** buttons, as described below.

     a. Select **Model2** as **Member Type**.

     a. Enter the port WWN in hexadecimal format in the **Member PWWN/MAC** field, and click **Add.**

        The WWN displays under **Allowed Login Members**. If you decide a member should not be on the list, highlight the entry and click **Remove**.

7. Click **OK**.

Displaying FCoE port information

There are 24 internal FCoE Ports that bridge FC and Ethernet traffic. You can view FCoE port information from the **Port Administration** panel.

Use the following procedure to display FCoE port information.

1. Select the **FCoE Ports** tab on the **Port Administration** panel.

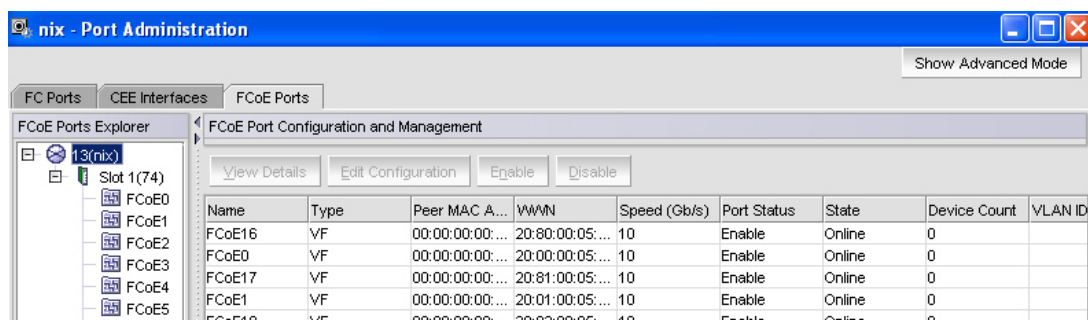   The initial view displays a summary of all FCoE ports on the switch (Figure 67).

FIGURE 67   FCoE Ports tab, Port Administration panel

2. To view information for a specific port, select the trunk in the **FCOE Ports Explorer** or select the port in the **FCoE Port Configuration and Management** table and click **View Details**.

   Port information displays in three tabs. The **General** tab is pre-selected.

   The **Connected Devices** tab displays information about devices connected to the switch (Figure 67). Six columns of information are displayed:

   - **Device WWN** displays the WWN of the connected device.
   - **Device MAC** displays the MAC address of the connected device.

- **Connected Peer Type** displays the port type on the connected device.
- **Is Directly Connected** indicates whether or not the device is directly connected to the trunk.
- **FCoE Port MAC** displays the FCoE port MAC address.
- **Switch Port** displays the switch port WWN.

Displaying LAG information

Use the following procedure to display LAG information.

1. Select the **CEE** tab on the **Switch Administration** panel.

2. Select the **Link Aggregation** tab.

    The LAG information displays.

Displaying VLAN information

Use the following procedure to display VLAN information.

1. Select the **CEE** tab on the **Switch Administration** panel.

2. Select the **VLAN** tab.

    The VLAN information displays.

Displaying FCoE login groups

Use the following procedure to display FCoE login group information.

1. Select the **CEE** tab on the **Switch Administration** panel.

2. Select the **FCoE Login** tab.

    The FCOE login group information displays.

Displaying QoS information

Use the following procedure to display QoS information.

1. Select the **CEE tab** on the **Switch Administration** panel.

2. Select the **QoS** tab.

    From the **QoS** tab, you can select the **CEE Map** tab to display CEE map information, or select the **Traffic Class Map** tab to display traffic class maps information.

Displaying LLDP-DCBX information

Use the following procedure to display LLDP-DCBX information.

1. Select the **CEE** tab on the **Switch Administration** panel.

2. Select the **LLDP-DCBX** tab.

    -   To display global settings, select the **Global** tab.
    -   To display LLDP profile information, select the **LLDP Profile** tab.

Displaying CEE interface statistics

The CEE interface **Port Statistics** tab displays basic and advanced statistics, and allows you to change statistics collection parameters. Use the following procedure to display CEE interface statistics.

The **CEE Interface Statistics Configuration** section allows you to do the following:

-   Toggle between showing **Absolute Values** or **Delta Values** (values that have changed since the last data collection).
-   Use the **Clear Counters** button to clear the counters in port statistics.
-   Change the retrieval interval.

To view additional information, select **Show Advanced Mode**. An **Advanced** tab and an **Error Detail** tab are added next to **Basic Mode**.

The **Advanced** tab displays CEE transmission statistics, and the **Error Details** tab displays transmission error statistics.

Use the following procedure to display CEE interface statistics.

1. Select the **CEE Interfaces** tab on the **Port Administration** panel.

2. Under the **CEE Interface Explorer**, select a port.

3. Select the **Port Statistics** tab.

Configuring a CEE interface from the Switch View

CEE interfaces can be enabled and disabled from a right-click menu on the **Switch View**.

To enable or disable a CEE interface from the **Switch View**, perform the following steps.

1. Right-click the port to display the right-click menu.

2. Select **Configure** to display the **Enable** and **Disable** options.

Configuring a CEE interface from the Port Admin panel

CEE interfaces can be enabled and disabled from the **Port Administration** panel.

To enable or disable a CEE interface from the **Port Administration** panel, perform the following steps.

1. Select the **CEE Interfaces** tab on the **Port Administration** panel.

2. Under the **CEE Interface Explorer**, select the port you want to enable or disable.

3. Select the **General** tab.

   This tab is normally pre-selected. You can follow either of the following options to enable or disable the interface:

   - Click **Enable Interface** or **Disable Interface** to enable or disable the interface, as desired.
   - Click **Edit Configuration** to open the **CEE Edit Configuration** dialog box. Select **Enable** or **Disable** for the **Status** field to enable or disable the interface.

Enabling and disabling a LAG

To enable or disable a LAG, perform the following steps.

1. Select the **CEE** tab on the **Switch Administration** panel.

2. Select the **Link Aggregation** tab.

3. Click **Add**.

   The **LAG Configuration** dialog box displays.

4. Change the **Status** to **Administratively Up** or **Administratively Down**.

Enabling and disabling LLDP

To enable or disable LLDP on a CEE interface, perform the following steps.

1. Select the **CEE Interfaces** tab on the **Port Administration** panel.

2. Under the **CEE Interface Explorer**, select the port.

3. Select the **General** tab.

4. Click **Edit Configuration**.

   The **CEE Edit Configuration** dialog box displays.

5. For the **LLDP Status** option, select **Enable** or **Disable**.

Enabling and disabling QoS priority-based flow control

Priority-based flow control (PFC) can be used to control network congestion. PFC can be used to selectively pause lower priority traffic classes to ensure that high priority and delay-sensitive traffic are not affected by network congestion. For example, if a large storage transfer is monopolizing the network and causing congestion, PFC can be used to pause the storage transfer so other traffic may use the network.

To enable or disable PFC, perform the following steps.

1. Select the **CEE** tab on the **Switch Administration** panel.

2. Select the **QoS** tab.

3.   Select the **CEE Maps** tab.

4.   Under the **Priority Group** area, enable or disable **Priority Flow Control Status** option for each
     Priority Group ID.

Enabling and disabling FCoE ports

You can enable and disable FCoE Ports individually from the **Port Administration** panel.

1.   Select the **FCoE Ports** tab on the **Port Administration** panel.

2.   Select the port you want to enable or disable under the **FCOE Ports Explorer,** or from the list.

3.   Click **Enable** or **Disable** to change the current status of the port.

     You can also enable or disable by clicking **Edit Configuration**, and selecting **Enable** or **Disable**
     on the **FCoE Edit Configuration** dialog box (Figure 68).
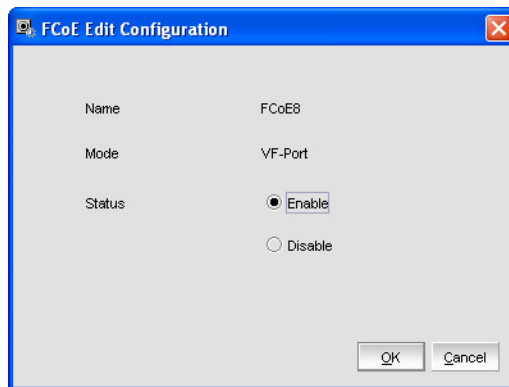
FIGURE 68   **FCoE Edit Configuration dialog box**

Limitations

In this chapter

- General Web Tools limitations 271

General Web Tools limitations

Table 23 lists general Web Tools limitations that apply to all browsers and switch platforms.

TABLE 23   Web Tools limitations

Area	Details
Blade Failure	If a blade fails on the switch, the Web Tools interface can still display slot and ports as healthy. In this case, the failure might not be visible in Web Tools until the Web Tools window is reopened.
Browser	For Internet Explore 7.0, the default setting is to disable Telnet functionality. You must make the appropriate changes in the registry to enable Telnet functionality if you want to use it. Launching the default Telnet is not supported in Windows Vista and Windows 2008 server.
Browser	Fabric Watch, Switch Admin, HA, Name Server, and Zone Admin are separate applets embedded in HTML pages. The successful launch of the applet depends on whether the browser can successfully load the HTML page. Very occasionally, a blank browser window displays with the message "loading pages…" that is stuck. This is likely caused by a sudden loss of switch Web server (either by normal HA failover, restart, or other causes). **Workaround:** If Fabric Watch, Switch Admin, HA, Name Server, or Zone Admin hang, close this window and relaunch the module.
Browser	A Web Tools browser window might stop responding following an HA failover immediately after a zoning configuration was enabled or disabled. It is likely that the Web daemon was terminated by the HA failover before the HTTP request was sent back. **Workaround:** If one of the Web Tools modules is hanging, close the window and relaunch the module. If the module is locked, shut down and relaunch the Web Tools application.
Browser	When you launch Fabric Watch, Switch Admin, Name Server, and Topology from Switch Explorer through Internet Explorer, the applet windows cannot be resized and the Maximize button is disabled.
Chassis not ready for management	If the switch is still in the process of booting and you try to launch the Web Tools by entering the IP address, this message displays in the browser. You should wait for the switch to finish the startup sequence.
Configuration	Web Tools does not support NAT router configurations and does not function correctly with switches behind a NAT router.

TABLE 23   Web Tools limitations (Continued)

Area	Details
Firmware download	There are multiple phases to firmware download and activation. When Web Tools reports that firmware download completed successfully, this indicates that a basic sanity check, package retrieval, package unloading, and verification was successful. Web Tools forces a full package install.
	A restart is required to activate the newly downloaded firmware. This restart is done automatically; however, although Web Tools screens continue to display during the restart, they are not available. Wait approximately 10 minutes to ensure that all of the application windows are restored. If Web Tools fails to respond after 20 minutes, you might need to close all Web Tools applications windows and restart them, or to contact your system administrator for network assistance.
	The Web Tools loss of network connectivity during a failover or restart (initiated through the **firmwareDownload**) varies for different configurations:
	Brocade 48000 and DCX and DCX-4S enterprise-class platforms: loss of network connectivity is up to 5 minutes if the power-on self-test (POST) is disabled. If POST is enabled, the loss of network connectivity can exceed 5 minutes.
	Brocade 200E, 300, 4012, 4016, 4018, 4020, 4024, 4100, 4900, 5000, 5100, 5300, 7500 and 7500E Extension switches, 7600, and the Encryption Switch: Loss of network connectivity is up to 1 minute if POST is disabled. If POST is enabled, the loss of network connectivity can exceed 1 minute.
Firmware downgrade	If you try to run Web Tools on a switch after downgrading the firmware, Web Tools may not open. This is due to the presence of old application cache files in Java. The workaround is to delete the application cache files using the Java Control Panel.
HTTP timeout	Occasionally, you might see the following message when you try to get data from a switch or to send a request to the switch:
	Failed to get switch response. Please verify the status of your last operation and try again if necessary.
	This indicates that an HTTP request did not get a response. The request was sent to the switch, but the connection was down, probably caused by a temporary loss of the Web server on the switch. Due to the nature of an HTTP connection, Web Tools reports this error after a 90-second default timeout.
	In this case, verify the status of your last request, using Telnet to check related status, or click the **Refresh** button from the Web Tools application you were working on to retrieve related data. If your request did not get through to the switch, resubmit it. Executing a refresh from Web Tools retrieves a copy of switch data at that moment; the data you entered can be lost if it had not already committed to the switch.
Java cache	If the Web Tools progress bar stops at 93 percent when initializing switch details, you must clear the Java cache, as described in "Deleting temporary internet files used by Java applications" on page 6.
Java Plug-in	If you remove the certificate (not recommended) from the Java Control Panel, you must close and reopen the browser for the certificate removal to take effect.
Java Plug-in	If you have a Web Tools session open and you open a second session using the **File > New** browser menu, this results in unexpected behavior of the original Web Tools session. For example, you cannot change Admin Domains in the second session.
	Web Tools supports only one browser instance per JRE, and when you open another window using the **File > New** menu, the two windows share the same JRE environment. **Workaround**: Open two independent browser sessions.

TABLE 23   Web Tools limitations (Continued)

Area	Details
Loss of Connection	Occasionally, you might see the following message when you try to retrieve data from the switch or send a request to the switch: *Switch Status Checking* *The switch is not currently accessible.* The dialog box title may vary, because it indicates which module is having the problem. This is caused by the loss of HTTP connection with the switch, due to a variety of possible problems. Web Tools automatically tries to regain the connection. While Web Tools is trying to regain the connection, check if your Ethernet connection is still functioning. If the problem is not with the Ethernet connection, wait for Web Tools to recover the connection and display the following message: "You will have to resubmit your request after closing this message." If the temporary switch connection loss is caused by switch hot code load, or other similar operation, Switch Explorer you are currently running can be downloaded from a different firmware version than the new one. In this case the following message displays: "Switch connection is restored. The firmware version you are running is not in sync with the version currently on switch. Close your browser and re-launch Web Tools." You need to close Switch Explorer and relaunch Web Tools to reopen the connection.
Non-FIPS secure mode HTTPS	HTTPS supports only TLSv1 and SSLv3 protocols with !DH:HIGH:-MD5 cipher in non-fips mode. These options must be enabled in your internet browser.
Out of Memory Errors	If you are managing fabrics with more than ten switches or more than 1000 ports, or if you are using the iSCSI Gateway module extensively, you might encounter out-of-memory errors such as the following: java.lang.OutOfMemoryError: Java heap space To avoid this problem, increase the default heap size in the Java Control Panel. Refer to "Java plug-in configuration" on page 8 for instructions.
Performance Monitor	If the Web browser crashes or the Performance Monitor license is lost while the Performance Monitoring window is running, some of the Performance Monitor resources owned by Web Tools might not be cleaned up correctly. **Workaround:** You might need to use the CLI to manually delete these counters. For example, if you detect Web Tools owned resources (using **perfshoweemonitor**), but you have verified that no Web users are actually using them, use the **perfdeleemonitor** or **perfcleareemonitor** command to free the resources.
Performance Monitor	For SCSI Read, Write, or Read/Write on a LUN per Port graphs, Fabric OS 4.1.0 (and later 4.x versions) allows you to enable only two bytes or less for the LUN value mask setting. Fabric OS 3.1 (and later 3.x versions) allows up to three bytes. Web Tools displays an error message if you exceed this limit. **Workaround:** There is no workaround.
Performance Monitor	For Brocade 48000 directors, while monitoring the performance, if one or all the blades turn Faulty or if they are powered off or on, then the behavior of various monitoring graphs is as follows: The Switch Aggregate and Blade Aggregate graphs freeze without any updates (about the traffic). **Workaround:** Close and relaunch the graphs. The Switch Throughput Utilization, Switch Percent Utilization, and Port Snapshot Error graphs displays the faulty/powered off slot node in the Y-Axis of the graph. **Workaround:** Launch any port selection dialog box and load the graphs accordingly.
Refresh option in browsers	When a pop-up window requesting a user response is pushed into the background and a refresh is requested, a fatal Internet Explorer error might occur. **Workaround:** Restart the browser.

TABLE 23   Web Tools limitations (Continued)

Area	Details
Refresh option in browsers	Web Tools must be restarted when the Ethernet IP address is changed using the **NetworkConfig View** command. Web Tools appears to hang if it is not restarted after this operation is executed. **Workaround:** Restart the browser.
Refresh option in browsers	If you change the switch name or domain ID using the CLI after the Web Tools Switch Administration window has started, the new switch name or domain ID is not updated on the header of the Switch Administration page. Clicking the **Refresh** button does not fix the problem. **Workaround:** Click the **Switch** tab and the Switch Administration header updates.
Refresh option in browsers	If you change the switch name using the Web Tools Switch Administration page or SNMP and then open a Telnet window to verify the name change, the CLI prompt (for example, **switch:admin** >) displays the previous name. The Telnet prompt cannot pick up the new switch name until the switch is fastbooted. **Workaround:** In order to display the correct switch name in the CLI prompt after a switch name update using Web Tools or SNMP, **fastboot** the switch.
Refresh option in browsers	Following a switch enable or disable, you must wait at least 25–30 seconds for the fabric to reconfigure and for FSPF route calculations to complete before requesting routing information. If accessed too early, routing information are not shown. **Workaround:** Following a switch enable or disable, wait at least 25–30 seconds before further action.
Refresh option in browsers	The Web Tools Switch Explorer might continue to display a switch from the Switch View, even when the switch has been removed from the fabric. **Workaround:** If this behavior is seen, relaunch Switch Explorer. If the switch was removed from the fabric, the Fabric View window lists the switch as unavailable.
Refresh option in browsers	In the Switch Administration window, **Switch** tab, if you click the **Refresh** button, you might not be able to click the data entry fields to enter text. This behavior occasionally happens on a notebook or laptop computer; it rarely happens on a desktop computer. **Workaround:** If this happens, you should close the browser window and restart it.
Switch Explorer closure	If a session times out or you log out or close Switch Explorer window, all other windows belonging to the session are invalidated. After a short delay these windows become unusable, but are not closed automatically. You must manually close these windows.
Switch View	Occasionally, switches might display the port icons correctly, but be missing one or more control button icons. **Workaround:** Close the Switch View of the switch and reopen it.
Windows Operating Systems	Occasionally, you will not see the "Lost connection to the switch" message on the Switch View, even though the Ethernet connection has been lost. You might still be able to invoke various features from Switch View, such as Status, Fan Temp, Power, and Beacon. **Workaround:** Verify Ethernet connection to the switch by pinging the logical switch IP address.
Windows Operating Systems	While working on Internet Explorer 6.0, when the user launches Switch Explorer it initially does not activate. Click the window once with the mouse, press the ESC key, the Space Bar, or Enter to activate the window. This is applicable in all applets launched using IE 6.0. **Workaround:** This is not seen while working on Firefox.

Index

Numerics

2 domain/4 domain fabric licenses, *9*
7800 switch, *91*, *92*

A

About Discovery Domains (DD), *193*
Access Control List. *Refer to* ACL
access control. *Refer to* RBAC.
Access Gateway mode
    configuration, *159*
    disable, *163*
    enable, *162*
    F_Port trunk groups, *101*
accessing switch event report, *54*
activating
    CUP port connectivity configuration, *250*
    licenses, *47*
    Ports on Demand, *93*
adding
    performance graphs to a canvas, *116*
    zone alias members, *126*
    zone configuration members, *132*
    zone members, *128*
Admin Domain window, *72*
    closing, *77*
    refreshing, *76*
Admin Domains
    assigning administrators, *212*
    creating, *77*
    deleting, *80*
    direct port membership, *83*
    indirect port membership, *83*
    opening, *72*
    to activate/deactivate, *79*
alarm configuration report for Fabric Watch, *177*
alarms, Fabric Watch
    configuring, *174*, *176*
    displaying, *177*
    enabling and disabling, *176*

aliases, zone. *Refer to* zone aliases
all access zoning, *121*
arbitrated loop parameters, configuring, *45*
automatic trace dump transfers, *140*

B

backbone fabric, *149*
backbone fabric ID, configuring, *156*
backing up configuration file, *61*
basic performance monitoring graphs, *109*
BB credit, *44*
beaconing, enabling, *57*, *58*
best practices for zoning, *138*
blades, enabling and disabling, *38*
browsers
    limitations, *271*, *274*
    refresh frequency, setting, *5*
    supported, *4*
buffer-limited ports, *181*

C

Challenge Handshake Authentication Protocol. *Refer to* CHAP
changing
    domain ID, *41*
    passwords, *213*
    switch name, *41*
CHAP
    authentication, *187*
    secret, editing, *196*
    user, creating, *196*
class F traffic, *44*
clearing temporary internet files, *6*
clearing the zoning database, *138*
closing
    Admin Domain window, *77*
    sessions, *13*
code page, displaying, *246*

configuration
Access Gateway mode, *159*
upload, *162*
configuration file
Admin Domain considerations, *64*
backing up, *61*
restoring, *63*
configuring
arbitrated loop parameters, *45*
backbone fabric ID, *156*
CUP port connectivity, *248*
default heap size, *8*
email notifications, *178*
EX_Ports, *153*
fabric parameters, *43*
Fabric Watch thresholds, *173*
FAN frame notification parameters, *45*
FC ports, *87*
FCR router cost, *155*
FICON Management Server parameters, *245*
FRU alarms, *176*
IOD frames delivery, *204*
IP route for iSCSI Target Gateway, *191*
Java Plug-in, *8*
link cost, *204*
long-distance settings, *183*
port speed, *87*
port type, *87*
ports, *83*
RADIUS server, *225*
routes, *201*
syslog IP address, *37*
system services, *46*
threshold alarms, Fabric Watch, *174*
virtual channel settings, *45*
configuring FCR router port costs, *155*
Control Device state, *247*
Control Unit Port. *Refer to* CUP
copper GigE, *91, 92*
copying CUP port connectivity configuration, *251*
CP failover, initiating, *53*
creating, *194*
Admin Domains, *77*
basic performance graphs, *109*
CHAP user for iSCSI Target Gateway, *196*
DDset, *195*
discovery domains (DD), *194*
iSCSI Fibre Channel zones with an effective zone
  configuration, *198*
iSCSI Fibre Channel zones with no effective zone

configuration, *198*
SCC/DCC policy, *216*
SCSI command graphs, *112*
SCSI vs. IP traffic graphs, *112*
SID-DID performance graphs, *111*
virtual targets for iSCSI Target Gateway, *191, 192*
zone aliases, *126*
zone configurations, *131*
zones, *128*
creating FCS policy, *216*
CUP port connectivity configuration
activating, *250*
copying, *251*
deleting, *251*
displaying, *248, 251, 252*
customizing basic performance graphs, *109*

D

datafield size, *44*
DDSet
creating, *195*
editing, *195*
default zoning, *121*
deleting
Admin Domains, *80*
CUP port connectivity configuration, *251*
user accounts, *212*
zone aliases, *127*
zone configurations, *133*
zones, *130*
device probing, *44*
devices only view, *125*
devices only zoning, *125*
direct port membership in Admin Domains, *83*
disabling
Access Gateway mode, *163*
automatic trace uploads, *141*
blades, *38*
dynamic load sharing, *203*
Fabric Watch threshold alarms, *176*
FICON Management Server mode, *244*
ports, *90, 91*
RADIUS, *225*
RLS probing, *46*
switch, *40*
trunking mode, *99*
zone configurations, *134*
zoning, *134*
Discovery Domain Set. *Refer to* DDSet

Discovery Domains
  iSCSI Target Gateway
    discovery domains (DD), *193*
displaying
  alarms, Fabric Watch, *177*
  Control Device state, *247*
  CUP port connectivity configuration, *248, 251, 252*
  enabled zone configuration, *135*
  fan status, *142*
  FICON code page, *246*
  name server entries, *56*
  power supply status, *143*
  switch events, *54*
  temperature status, *143*
  user account information, *209*
DLS, *203*
domain ID, changing, *41*
downloading
  configuration file, *63*
  firmware, *65*
Dynamic Load Sharing. *Refer to* DLS

E

E_D_TOV, *44*
edge fabrics, *149*
editing
  DDset, *195*
  discovery domains (DD), *194*
  iSCSI Fibre Channel zone members, *198*
EGM licensed features
  FICON CUP, *243*
  Performance Monitoring, *18*
email notifications, *178*
enabled zone configuration, displaying, *135*

enabling
  Access Gateway mode, *162*
  automatic trace dump transfer, *141*
  beaconing, *57, 58*
  blades, *38*
  DLS, *203*
  Fabric Watch threshold alarms, *176*
  FICON Management Server mode, *244*
  insistent domain ID mode, *44*
  iSCSI Target Gateway service, *189*
  ports, *90*
  Ports on Demand, *93*
  RADIUS, *225*
  RLS probing, *46*
  switch, *40*
  trunking mode, *99*
  zone configurations, *134*
ending sessions, *13*
events
  displaying, *54*
  filtering, *55*
  severity levels, *54*
EX_Ports, configuring, *153*
exchange-based routing, *201, 203*
expiring passwords, *215*
extended fabrics, *181*

F

F_Port trunk groups
  Access Gateway mode, *101*
fabric ID, configuring, *156*
fabric information, refreshing, *76, 123*
fabric parameters, configuring, *43*
Fabric Tree, *23*
fabric view, *125*
fabric view zoning, *125*
Fabric Watch
  about, *171*
  alarms, *174*
  thresholds, *173*
failover, initiating, *53*
FAN frame notification parameters, configuring, *45*
fan status, *142*
fast boot, *42*
FC ports, configuring, *87*
FC Routing module, *150*
FC targets, searching for iSCSI Target Gateway, *193*

PDU, *187*
performance graphs
    adding to a canvas, *116*
    modifying, *117*
    printing, *116*
    types of, *105*
Performance Monitoring window, *108*
per-frame routing priority, *44*
persistently disable a port, *91*
platforms, supported, *5*
polling rates, *30*
port membership in Admin Domains, *83*
port menu, *29*
port names, assigning, *89*
port speed, configuring, *87*
port swapping, *97*
port type, configuring, *87*
port-based routing, *201*
ports
    buffer-limited, *181*
    configuring, *83*
    disabling, *90, 91*
    enabling, *90*
    LEDs, *146*
    long distance parameter, *183*
    naming, *89*
Ports on Demand, enabling, *93*
power supply status, *143, 144*
preferences
    persist, *22*
printing
    effective zone configuration, *136*
    performance graphs, *116*
protocol options, *273*

R

R_A_TOV, *44*
RADIUS server
    about, *224, 227*
    configuring, *225*
    enabling and disabling, *225*
    modifying, *226*
    modifying server order, *226*
    removing, *227*
RAM requirements, *5*
RBAC, pre-defined roles, *14*
rebooting the switch, *42*

recommendations
    configuration tasks, *31*
    for Web Tools, *31*
    for zoning, *138*
refresh frequency, setting, *5*
refresh rates, *30*
refreshing
    Admin Domain window, *76*
    fabric information, *76, 123*
    Zone Administration window, *124*
removing
    licenses, *48*
    RADIUS server, *227*
    zone alias members, *126*
    zone configuration members, *132*
    zone members, *128*
renaming
    zone aliases, *127*
    zone configurations, *133*
    zones, *129*
replacing a WWN in zoning database, *137*
requirements, Web Tools, *4*
restoring configuration file, *63*
right-click menu, *29*
RLS probing, enabling and disabling, *46*
Role-Based Access Control. *Refer to* RBAC
router cost path, *155*
routes, configuring, *201*

S

saving
    performance graphs, *114*
    zoning changes, *76, 124*
SCC/DCC policy
    activate, *217*
    create, *216*
    deactivate, *217*
    delete, *217*
    edit, *216*
SCSI command graph, *112*
SCSI vs. IP traffic graph, *112*
searching zone member selection lists, *137*
sequence level switching, *44*
session management, *14*
sessions, ending, *13*
setting
    refresh frequency, *5*
    SNMP trap levels, *222*

severity levels, *54*
SID-DID performance graph, *111*
SNMP trap levels, *222*
Solaris patches, installing, *7*
SSLv3, *273*
starting Web Tools, *9*
swapping port index IDs, *97*
switch
    7800, *91*, *92*
    changing the name of, *41*
    enabling and disabling, *40*
    mouse over information, *28*
    rebooting, *42*
Switch Administration window, *33*
    opening, *35*
Switch Events and Switch Information, *27*
switch events, displaying, *54*
Switch Explorer, Admin Domains, *23*
switch name, changing, *41*
switch report, *41*
switch status report, *144*
Switch View, *25*
Switch View buttons, *25*
syslog IP address
    configuring, *37*
    removing, *37*
system services, configuring, *46*

T

Telnet, *32*
temperature status, *143*
temporary internet files, *6*
threshold alarms, Fabric Watch
    configuring threshold alarms, *174*
    enabling and disabling, *176*
timeout, session, *14*
TLS, *273*
trace dumps, *139*
transition, partial Web Tools functions to DCFM, *3*
troubleshooting
    iSCSI Target Gateway, *199*
    Web Tools, *31*
trunking mode, enabling and disabling, *99*

U

unlocking passwords, *215*
user accounts, managing, *207*

V

value line licenses, *9*
VC Priority, *45*
viewing
    EX_Ports, *153*
    LSAN devices, *156*
    LSAN fabrics, *152*
    LSAN zones, *156*
    swapped ports, *97*
    Switch Explorer, *17*
    switch status, *144*
viewing FCR router cost, *155*
virtual channel settings, configuring, *45*
virtual targets, creating for iSCSI Target Gateway, *191*,
    192

W

Web Tools
    Access Gateway mode, enable, *162*
    GUI preferences, *22*
    launching, *9*
    partial function transition to DCFM, *3*
WWN
    adding to zones, *136*
    removing from zones, *136*
    replacing in zones, *137*

Z

Zone Admin module, saving changes, *76*
Zone Administration window
    about, *121*
    refreshing, *124*
    saving changes, *124*
zone aliases
    creating, *126*
    deleting, *127*
    description, *126*
    modifying, *126*
    renaming, *127*