

Brocade Fabric OS Web Tools Administration Guide, 8.2.0

Supporting Fabric OS 8.2.0

Copyright © 2017 Brocade Communications Systems LLC. All Rights Reserved.

Brocade and the stylized B logo are among the trademarks of Brocade Communications Systems LLC. Broadcom, the pulse logo, and Connecting everything are among the trademarks of Broadcom. The term "Broadcom" refers to Broadcom Limited and/or its subsidiaries

Brocade, a Broadcom Limited Company, reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Brocade is believed to be accurate and reliable. However, Brocade does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

The product described by this document may contain open source software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <http://www.brocade.com/support/oscd>.

Contents

About This Document	11
Supported hardware and software.....	11
Brocade Gen 5 (16-Gbps) fixed-port switches.....	11
Brocade Gen 5 (16-Gbps) directors.....	11
Brocade Gen 6 (32-Gbps) fixed-port switches.....	11
Brocade Gen 6 (32-Gbps) directors.....	12
What's new in this document for Fabric OS 8.2.0.....	12
Document conventions.....	12
Notes, cautions, and warnings.....	12
Text formatting conventions.....	13
Command syntax conventions.....	13
Document feedback.....	14
Introducing Web Tools.....	15
Web Tools overview.....	15
Web Tools and Brocade Network Advisor.....	15
Web Tools functionality moved to Brocade Network Advisor.....	15
System requirements.....	16
Setting the refresh frequency for Internet Explorer.....	18
Deleting temporary Internet files used by Java applications.....	19
Java installation on the workstation.....	20
Installing the JRE on your Solaris or Linux client workstation.....	20
Installing patches on Solaris.....	20
Installing the Java plug-in on Windows.....	21
Java plug-in configuration.....	21
Enabling Java content in the browser.....	21
Configuring the Java plug-in for Windows.....	21
Configuring the Java plug-in for Mozilla family browsers.....	22
Opening Web Tools.....	22
Server authentication for HTTPS.....	23
Logging in.....	25
Logging out.....	27
Role-Based Access Control.....	27
Session management.....	28
Ending a Web Tools session.....	28
Web Tools system logs	28
Application SupportSave logs.....	29
Requirements for IPv6 support.....	30
Using the Web Tools Interface.....	31
Viewing Switch Explorer.....	31
Persisting GUI preferences.....	33
Tabs.....	33
Fabric Tree.....	33
Switch View buttons.....	35
Switch View.....	35
Switch Events and Switch Information.....	37

Free Professional Management tool.....	38
Displaying tool tips.....	38
Right-click options.....	39
Refresh rates.....	39
Displaying switches in the fabric.....	40
Recommendations for working with Web Tools.....	40
Opening an SSH client window.....	41
Collecting logs for troubleshooting.....	41
Managing Fabrics and Switches.....	43
Fabric and switch management overview.....	43
Opening the Switch Administration window.....	45
Configuring IP and subnet mask information.....	47
Configuring Netstat Auto Refresh.....	47
Configuring a syslog IP address.....	48
Removing a syslog IP address.....	48
Configuring IP filtering.....	48
Blade management.....	49
Enabling or disabling a blade.....	49
Setting a slot-level IP address.....	50
Viewing IP addresses.....	51
Switch configuration.....	51
Enabling and disabling a switch.....	51
Enabling and disabling switch persistence.....	52
Changing the switch name.....	52
Changing the switch domain ID.....	52
Viewing and printing a switch report.....	53
Setting a principal switch.....	53
Switch restart.....	54
Performing a fast boot.....	54
Performing a reboot.....	54
System configuration parameters.....	55
WWN-based persistent PID assignment.....	55
Configuring fabric settings.....	56
Enabling insistent domain ID mode.....	56
Configuring virtual channel settings.....	57
Configuring arbitrated loop parameters.....	57
Configuring system services.....	58
Configuring CSCTL QoS mode.....	58
Dynamic Port Name configuration.....	58
Slow-Drain Device Quarantine configuration.....	59
Licensed feature management.....	60
Activating a license on a switch.....	60
Assigning slots for a license key.....	61
Removing a license from a switch.....	61
Universal time-based licensing.....	61
High Availability overview.....	62
Launching the High Availability window.....	62
Synchronizing services on the CP.....	63
Initiating a CP failover.....	64
Event monitoring.....	64

Displaying switch events.....	65
Filtering switch events.....	65
Filtering events by event severity levels.....	66
Filtering events by message ID.....	66
Filtering events by service component.....	66
System Monitor.....	67
Monitoring memory usage.....	68
Monitoring CPU usage.....	69
Displaying Name Server entries.....	70
Printing Name Server entries.....	71
Displaying Name Server information for a particular device.....	71
Displaying zone members for a particular device.....	71
Physically locating a switch using beaconing.....	71
Locating logical switches using chassis beaconing.....	72
Virtual Fabrics overview.....	72
Selecting a logical switch from the Switch View.....	72
Viewing logical ports.....	73
MAPS limited monitoring support.....	74
Maintaining Configurations and Firmware.....	75
Creating a configuration backup file.....	75
Restoring a configuration.....	76
Uploading and downloading from USB storage.....	77
Performing a firmware download.....	77
Managing Ports.....	79
Port management overview.....	79
Opening the Port Admin tab.....	79
Port Admin tab components.....	80
Controllable ports.....	85
Configuring FC ports.....	85
Allowed port types.....	86
Speed.....	87
Long-distance mode.....	88
Available buffer credit calculation.....	89
Assigning a name to a port.....	89
Port beaconing.....	90
Port peer beaconing.....	90
Enabling and disabling a port.....	92
Considerations for enabling or disabling a port.....	92
Persistently enabling and disabling ports.....	92
Configuring NPIV ports.....	95
Enabling Target Driven Zoning Mode.....	95
Port activation.....	96
Enabling Ports on Demand.....	96
Diagnostic ports.....	97
Reserving and releasing licenses on a port basis.....	97
Port swapping index.....	98
Port swapping.....	98
Determining if a port index was swapped with another switch port.....	99
Configuring port binding.....	100

Unbinding a port.....	101
Configuring BB credits on an F_Port.....	102
Configuring ALPA	102
Configuring the port octet speed combination	103
Configuring CSCTL.....	104
Enabling CSCTL mode.....	105
Disabling CSCTL mode.....	105
Configuring compression and encryption.....	106
Enabling or disabling encryption.....	106
Enabling or disabling compression.....	106
Displaying the compression ratio.....	107
Forward Error Correction.....	107
In-band management.....	108
GigE port modes.....	109
Enabling ISL Trunking.....	111
ISL trunking overview.....	111
Disabling or enabling ISL Trunking	111
Viewing trunk group information.....	111
F_Port trunk groups.....	112
Creating and maintaining F_Port trunk groups.....	113
CEE switch level configuration.....	115
CEE tab configuration overview.....	115
Link Aggregation.....	115
Adding Port Channel.....	116
Editing Port Channel.....	116
FCoE Management.....	117
LLDP.....	118
Monitoring Performance.....	121
Performance Monitor overview.....	121
Basic monitoring	121
Performance graphs.....	121
Predefined performance graphs.....	122
Opening the Performance Monitor window.....	123
Creating basic performance monitor graphs.....	123
Customizing basic monitoring graphs.....	123
Tunnel and TCP performance monitoring graphs.....	126
Tunnel and TCP graph chart properties.....	127
Printing graphs.....	127
Administering Zoning.....	129
Zoning overview.....	129
Basic zones.....	129
Traffic Isolation zones.....	129
Peer zones.....	129
Target Driven Zoning.....	131
LSAN zone requirements.....	131
QoS zone requirements.....	131
Zoning configurations	131
Opening the Zone Administration window.....	132

Setting the default zoning mode.....	132
Zoning management.....	132
Refreshing fabric information.....	134
Refreshing Zone Administration window information.....	135
Saving local zoning changes.....	135
Selecting a zoning view.....	136
Creating and populating zone aliases.....	136
Adding and removing members of a zone alias.....	137
Renaming zone aliases.....	138
Deleting zone aliases.....	138
Creating and populating zones.....	138
Adding and removing members of a zone.....	139
Renaming zones.....	140
Cloning zones.....	140
Deleting zones.....	140
Creating and populating enhanced Traffic Isolation zones.....	141
Zone configuration and zoning database management.....	141
Creating zone configurations.....	142
Adding or removing zone configuration members.....	143
Renaming zone configurations.....	143
Cloning zone configurations.....	144
Deleting zone configurations.....	144
Enabling zone configurations.....	144
Disabling zone configurations.....	145
Displaying enabled zone configurations.....	145
Viewing the enabled zone configuration name without opening the Zone Administration window.....	145
Viewing detailed information about the enabled zone configuration.....	146
Adding a WWN to multiple aliases and zones.....	146
Removing a WWN from multiple aliases and zones.....	147
Replacing a WWN in multiple aliases and zones.....	147
Searching for zone members.....	148
Clearing the zoning database.....	148
Zone configuration analysis.....	148
Best practices for zoning.....	149
Working with Diagnostic Features.....	151
Trace dumps.....	151
How a trace dump is used.....	151
Setting up automatic trace dump transfers.....	152
Specifying a remote server.....	152
Enabling automatic transfer of trace dumps.....	152
Disabling automatic trace uploads.....	152
Displaying switch information.....	153
Viewing detailed fan hardware status.....	153
Viewing the temperature status.....	154
Viewing the power supply status.....	155
Port LED interpretation.....	156
Port icon colors.....	156
Using the FC-FC Routing Service.....	159
Fibre Channel Routing overview.....	159

Supported switches for Fibre Channel Routing.....	160
Setting up FC-FC routing.....	160
FC-FC routing management.....	160
Opening the FC Routing module.....	161
Viewing and managing LSAN fabrics.....	161
Viewing EX_Ports.....	162
Configuring an EX_Port.....	162
Editing the configuration of an EX_Port.....	163
Configuring FCR router port cost.....	163
Viewing LSAN zones.....	163
Viewing LSAN devices.....	163
Configuring the backbone fabric ID.....	164
Using the Access Gateway.....	165
Access Gateway overview.....	165
Viewing Switch Explorer for Access Gateway mode.....	165
Access Gateway mode	167
Restricted access in the Port Admin tab.....	167
Enabling Access Gateway mode.....	167
Disabling Access Gateway mode.....	168
Viewing the Access Gateway settings.....	168
Port configuration.....	168
Editing a port.....	168
Creating port groups.....	169
Editing or viewing port groups.....	170
Deleting port groups.....	171
Defining custom primary F-N port mappings.....	171
Defining custom static F-N port mappings.....	171
Defining custom WWN-N port mappings.....	172
Access Gateway policy modification.....	172
Path Failover and Failback policies.....	172
Modifying Path Failover and Failback policies.....	172
Enabling the Automatic Port Configuration policy.....	173
Administering Extended Fabrics.....	175
Extended link buffer allocation overview.....	175
Configuring a port for long distance.....	178
Routing Traffic.....	181
Routing overview.....	181
Viewing Fabric Shortest Path First routing.....	182
Configuring dynamic load sharing.....	182
Lossless dynamic load sharing.....	183
Specifying frame order delivery.....	183
Configuring the link cost for a port.....	184
E_Port balance priority.....	184
Configuring Standard Security Features.....	187
User-defined accounts.....	187
Virtual Fabrics considerations.....	188
Viewing user account information.....	188
Creating user-defined accounts.....	188

Deleting user-defined accounts.....	190
Changing user account parameters	190
Maintaining passwords.....	191
User-defined roles.....	193
Guidelines and restrictions.....	194
Creating a user-defined role.....	194
Editing a user-defined role.....	195
Access control list policy configuration.....	197
Virtual Fabrics considerations.....	197
Creating an SCC, DCC, or FCS policy.....	197
Editing an SCC, DCC, or FCS policy.....	197
Deleting all SCC, DCC, or FCS policies.....	198
Activating all SCC, DCC, or FCS policies.....	198
Distributing an SCC, DCC, or FCS policy.....	198
Moving an FCS policy switch position.....	199
Configuring the Advanced Device Security policy	199
Fabric-Wide Consistency Policy configuration.....	200
Authentication policy configuration.....	200
Configuring authentication policies for E_Ports.....	201
Configuring authentication policies for F_Ports.....	201
Distributing authentication policies.....	201
Re-authenticating policies.....	202
Setting a shared secret key pair.....	202
Modifying a shared secret key pair.....	202
Setting the Switch Policy Authentication mode.....	203
SNMP configuration.....	203
Setting SNMP trap levels.....	203
Changing the systemGroup configuration parameters.....	203
Setting SNMPv1 configuration parameters.....	204
Setting SNMPv3 configuration parameters.....	204
Changing the access control configuration.....	204
RADIUS management.....	205
Enabling and disabling RADIUS.....	205
Configuring RADIUS.....	206
Modifying the RADIUS server.....	206
Modifying the RADIUS server order.....	207
Removing a RADIUS server.....	207
Active Directory service management.....	207
Enabling an Active Directory service.....	207
Modifying an Active Directory service.....	208
Removing an Active Directory service.....	208
TACACS+ management.....	209
Enabling and disabling TACACS+.....	209
Configuring TACACS+.....	209
Modifying TACACS+.....	209
Removing TACACS+.....	210
IPsec concepts.....	210
Transport mode and tunnel mode.....	211
IPsec header options.....	211
Basic IPsec configurations.....	212

Internet Key Exchange concepts.....	213
IPsec over management ports.....	215
Enabling the Ethernet IPsec policies	215
Establishing an IKE policy.....	216
Creating a security association.....	216
Creating an SA proposal.....	217
Adding an IPsec transform policy.....	217
Adding an IPsec selector.....	218
Manually creating an SA.....	218
Editing an IKE or IPsec policy.....	219
Deleting an IKE or IPsec policy.....	220
Establishing authentication policies for HBAs.....	220
Administering FICON CUP Fabrics.....	223
FICON CUP fabrics overview.....	223
Enabling port-based routing.....	224
Enabling or disabling FICON Management Server mode.....	224
FMS parameter overview.....	225
Configuring FMS mode parameters.....	226
Displaying code page information.....	226
Viewing the control device state.....	226
Allow / Prohibit Matrix configuration.....	227
Viewing Allow / Prohibit Matrix configurations.....	228
Modifying Allow / Prohibit Matrix configurations.....	228
Activating an Allow / Prohibit Matrix configuration.....	230
Copying an Allow / Prohibit Matrix configuration.....	230
Deleting an Allow / Prohibit Matrix configuration.....	230
CUP logical path configuration.....	230
Viewing CUP logical path configurations.....	231
Configuring CUP logical paths.....	231
Link Incident Registered Recipient configuration.....	231
Viewing Link Incident Registered Recipient configurations.....	231
Configuring LIRRs.....	232
Displaying Request Node Identification Data	232
Web Tools limitations.....	235
General Web Tools limitations.....	235

About This Document

• Supported hardware and software.....	11
• What's new in this document for Fabric OS 8.2.0.....	12
• Document conventions.....	12
• Document feedback.....	14

Supported hardware and software

The following hardware platforms are supported by Fabric OS 8.2.0.

Although many different software and hardware configurations are tested and supported by Brocade for Fabric OS 8.2.0, documenting all possible configurations and scenarios is beyond the scope of this document.

Fabric OS support for the Brocade Analytics Monitoring Platform (AMP) device depends on the specific version of the software running on that platform. For more information, refer to the Brocade Analytics Monitoring Platform documentation and release notes.

Brocade Gen 5 (16-Gbps) fixed-port switches

- Brocade 6505 Switch
- Brocade 6510 Switch
- Brocade 6520 Switch
- Brocade M6505 blade server SAN I/O module
- Brocade 6542 blade server SAN I/O module
- Brocade 6543 blade server SAN I/O module
- Brocade 6545 blade server SAN I/O module
- Brocade 6546 blade server SAN I/O module
- Brocade 6547 blade server SAN I/O module
- Brocade 6548 blade server SAN I/O module
- Brocade 6558 blade server SAN I/O module
- Brocade 7840 Extension Switch

Brocade Gen 5 (16-Gbps) directors

For ease of reference, Brocade chassis-based storage systems are standardizing on the term “director.” The legacy term “backbone” can be used interchangeably with the term “director.”

- Brocade DCX 8510-4 Director
- Brocade DCX 8510-8 Director

Brocade Gen 6 (32-Gbps) fixed-port switches

- Brocade G610 Switch
- Brocade G620 Switch

- Brocade G630 Switch

Brocade Gen 6 (32-Gbps) directors

- Brocade X6-4 Director
- Brocade X6-8 Director

What's new in this document for Fabric OS 8.2.0

This document includes new and modified information for the Fabric OS 8.2.0 release.

The following changes have been made for this release:

- Updated the JRE version in [System requirements](#) on page 16.
- Added support for the Brocade G630 Switch throughout the document.
- Added a note regarding Flex port configuration in [Port Admin tab components](#) on page 80.
- Added support for the CEE Interfaces, FCoE Ports, LLDP Neighbors subtabs in [Port Admin tab components](#) on page 80.
- Added a note regarding current and voltage information to the basic mode QSFP tab in [Subtabs](#) on page 84.
- Added a note for speed configuration on ICL ports in [Speed](#) on page 87.
- Added support for the link aggregation, FCoE management, and LLDP in [CEE switch level configuration](#) on page 115.
- Added a note regarding aliases as a members in [Target Driven Zoning](#) on page 131.
- Added support for Zoning Checksum in [Viewing detailed information about the enabled zone configuration](#) on page 146.
- Added support for the Brocade G630 Switch for firmware download in the [General Web Tools limitations](#) on page 235.
- Removed support for the FC8-32E, FC8-48E, FC8-64 blades throughout the document, as these blades are no longer supported.
- Made editorial revisions throughout for clarity and accuracy.

Document conventions

The document conventions describe text formatting conventions, command syntax conventions, and important notice formats used in Brocade technical documentation.

Notes, cautions, and warnings

Notes, cautions, and warning statements may be used in this document. They are listed in the 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 a stronger note, for example, to alert you when traffic might be interrupted or the device might reboot.

**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.

Text formatting conventions

Text formatting conventions such as boldface, italic, or Courier font may be used to highlight specific words or phrases.

Format	Description
bold text	Identifies command names. Identifies keywords and operands. Identifies the names of GUI elements. Identifies text to enter in the GUI.
<i>italic text</i>	Identifies emphasis. Identifies variables.
Courier font	Identifies document titles. Identifies CLI output. Identifies command syntax examples.

Command syntax conventions

Bold and italic text identify command syntax components. Delimiters and operators define groupings of parameters and their logical relationships.

Convention	Description
bold text	Identifies command names, keywords, and command options.
<i>italic text</i>	Identifies a variable.
value	In Fibre Channel products, a fixed value provided as input to a command option is printed in plain text, for example, --show WWN .
[]	Syntax components displayed within square brackets are optional.
{ x y z }	Default responses to system prompts are enclosed in square brackets. A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x y	In Fibre Channel products, square brackets may be used instead for this purpose. A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, for example, passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <i>member[member...]</i> .
\	Indicates a "soft" line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

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.

Send your feedback to documentation.pdl@broadcom.com

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

Introducing Web Tools

• Web Tools overview.....	15
• Web Tools and Brocade Network Advisor.....	15
• System requirements.....	16
• Java installation on the workstation.....	20
• Java plug-in configuration.....	21
• Opening Web Tools.....	22
• Role-Based Access Control.....	27
• Session management.....	28
• Web Tools system logs	28
• Application SupportSave logs.....	29
• Requirements for IPv6 support.....	30

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 Brocade Network Advisor.

A limited set of features is accessible using Web Tools without a license and is available free of charge.

Refer to [Web Tools and Brocade Network Advisor](#) on page 15 for more information.

Web Tools and Brocade Network Advisor

Beginning with Fabric OS 6.1.1, Web Tools functionality is tiered and integrated with Brocade Network Advisor. If you are migrating from a Web Tools release before Fabric OS 6.1.1, how you use Web Tools may be impacted.

A Web Tools license is not required; a basic version of Web Tools is available for free.

Beginning with Fabric OS 6.1.1, some Web Tools capabilities are moved from Web Tools to Brocade Network Advisor.

Web Tools functionality moved to Brocade Network Advisor

The functionality that was moved from Web Tools into Brocade Network Advisor is detailed in the following table.

TABLE 1 Web Tools functionality moved to Brocade Network Advisor

Function	Web Tools 6.1.0	Brocade Network Advisor
Add Un-Zoned Devices	Zone Admin	Configure > Zoning Reverse Find in the Zoning dialog box provides a 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 a view of the zoned and unzoned devices in the fabric if all zone members are selected for Find.

TABLE 1 Web Tools functionality moved to Brocade Network Advisor (continued)

Function	Web Tools 6.1.0	Brocade Network Advisor
		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 zoned devices if all zones are selected in the active zone configuration.
Define Device Alias	Zone Admin	Configure > Zoning
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 for Fabric OS 6.1.1, but New , Edit Config , and Delete are available only in Brocade Network Advisor.
GigE Ports Interface	Port Admin Module > GigE tab	Configure > FCIP Tunnel
GigE Ports Route	Port Admin Module > GigE tab	Configure > FCIP Tunnel
Nonlocal Switch Ports Display in Zoning Tree	Zone Admin Admin Domain Switch Admin > DCC policies Performance Monitoring	Configure > Zoning In Web Tools, a nonlocal switch port ID/WWN can be added using a 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 databases.

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 4.0, JavaScript 1.0, and JRE 1.8.0_152 update or later.

NOTE

If there are multiple JRE versions installed, go to the Java Control Panel and uncheck the lower JRE versions to launch Web Tools using the latest JRE version.

Brocade has certified and tested Web Tools on the platforms shown in the “Certified and tested platforms” table.

TABLE 2 Certified and tested platforms

Operating System	Browser
Oracle Enterprise Linux 6.5	Firefox 56
Oracle Enterprise Linux 6.6 Adv - 64 Bit	Firefox 56
Oracle Enterprise Linux 6.7 Adv - 64 Bit	Firefox 56
Oracle Enterprise Linux 7.0 - 64 Bit	Firefox 56
Oracle Enterprise Linux 7.1 - 64 Bit	Firefox 56
Red Hat Enterprise Linux 6.1 Adv (32-bit)	Firefox 56
Red Hat Enterprise Linux 6.3 Adv (64-Bit)	Firefox 56
Red Hat Enterprise Linux 6.5 Adv	Firefox 56
Red Hat Enterprise Linux 6.6 Adv (Beta)	Firefox 56
Red Hat Enterprise Linux 6.6 Adv - 64 Bit	Firefox 56
Red Hat Enterprise Linux 6.7 Adv - 64 Bit	Firefox 56
Red Hat Enterprise Linux 7.0 Adv - 64 Bit	Firefox 56
Red Hat Enterprise Linux 7.1 Adv - 32 Bit	Firefox 56
Red Hat Enterprise Linux 7.1 Adv - 64 Bit	Firefox 56
SUSE Linux Enterprise Server 11.3 - 64 Bit	Firefox 56
SUSE Linux Enterprise Server 12.0 - 64 Bit	Firefox 56
Windows Server 2008 R2 Standard - 64 Bit	Firefox 56, Internet Explorer 11
Windows Server 2008 R2 (SP1) Enterprise - 64 Bit	Firefox 56, Internet Explorer 9/10/11
Windows Server 2008 SP2 Standard - 64 Bit	Firefox 56, Internet Explorer 9/10/11
Windows Server 2008 R2 Datacenter - 64 Bit	Firefox 56, Internet Explorer 11
Windows Server 2012 Standard - 64 Bit	Firefox 56, Internet Explorer 10/11
Windows Server 2012 R2 Datacenter - 64 Bit	Firefox 56, Internet Explorer 11
Windows 8 Enterprise - 64 Bit	Firefox 56, Internet Explorer 10/11
Windows 8.1 Enterprise - 32 Bit	Firefox 56, Internet Explorer 11
Windows 8.1 Enterprise - 64 Bit	Firefox 56, Internet Explorer 11
Windows 10 Enterprise - 64 Bit	Firefox 56, Internet Explorer 11

Brocade supports the platforms shown in the following table.

TABLE 3 Supported platforms

Operating System	Browser
Oracle Enterprise Linux 6.5	Firefox 56
Oracle Enterprise Linux 6.8	Firefox 56
Oracle Enterprise Linux 6.9	Firefox 56
Oracle Enterprise Linux 7.0	Firefox 56
Oracle Enterprise Linux 7.3	Firefox 56
Red Hat Enterprise Linux 6.1 Adv (32-bit)	Firefox 56
Red Hat Enterprise Linux 6.3 Adv (64-Bit)	Firefox 56
Red Hat Enterprise Linux 6.5 Adv	Firefox 56
Red Hat Enterprise Linux 6.9	Firefox 56

TABLE 3 Supported platforms (continued)

Operating System	Browser
Red Hat Enterprise Linux 7.3	Firefox 56
SUSE Linux Enterprise Server 10 (32-bit)	Firefox 56
SUSE Linux Enterprise Server 11 (SP2) (32-Bit)	Firefox 56
SUSE Linux Enterprise Server 11.3	Firefox 56
SUSE Linux Enterprise Server 12.0 (Beta)	Firefox 56
Windows 7 Professional (x86)	Firefox 56, Internet Explorer 8/9/10
Windows 7 SP1	Firefox 56, Internet Explorer 8/9/11
Windows 8 Enterprise(64-Bit)	Firefox 56, Internet Explorer 10
Windows 8.1 Enterprise	Firefox 56, Internet Explorer 11
Windows Server 2003 Standard SP2 (x86 32-bit)	Firefox 56, Internet Explorer 8/9
Windows Server 2008 R2 (SP1) Enterprise (64-Bit)	Firefox 56, Internet Explorer 9/10
Windows Server 2008 SP2 Standard - 64 Bit	Firefox 56, Internet Explorer 9
Windows Server 2012 Standard (64-Bit)	Firefox 56, Internet Explorer 10
Windows Server 2012 R2	Firefox 56, Internet Explorer 11
Windows Server 2016 R2 Standard(64-bit)	Firefox 56, Internet Explorer 10

For Windows systems, a minimum of 1 GB of RAM and a minimum of 8 MB of video RAM are recommended.

Setting the refresh frequency for Internet Explorer

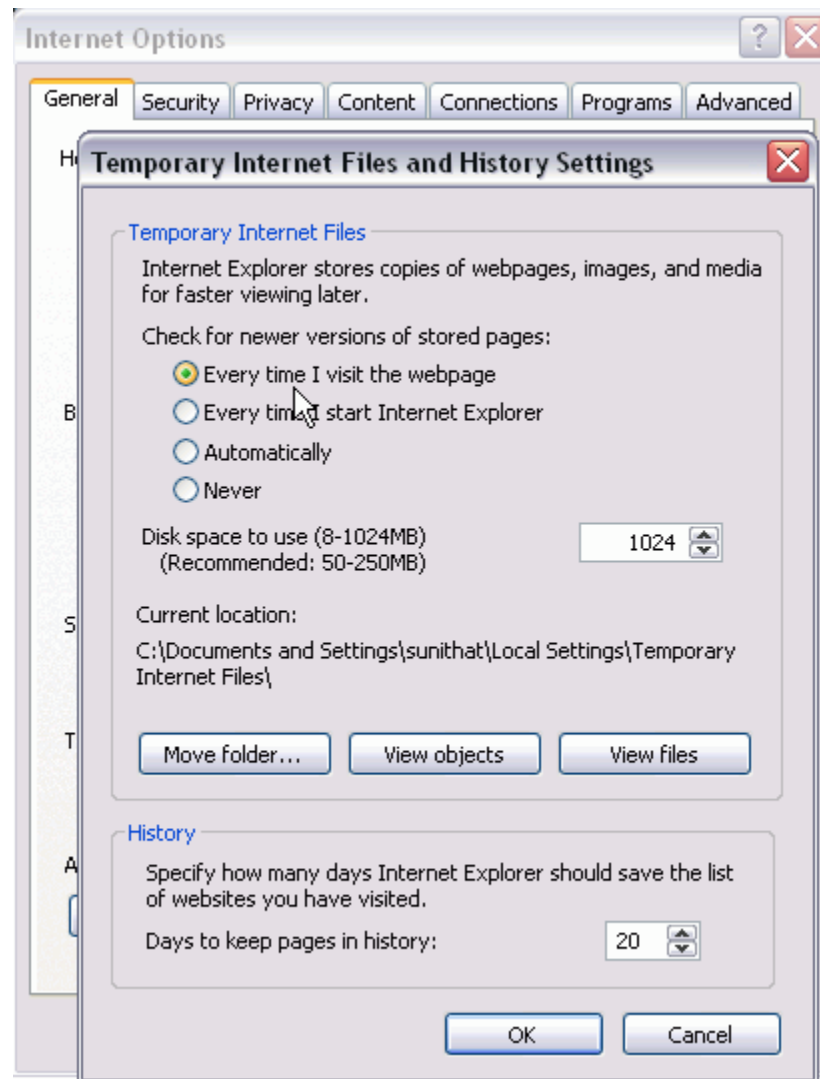
Correct operation of Web Tools with Internet Explorer requires specifying the appropriate settings for browser refresh frequency and process model. Browser pages should be refreshed frequently to ensure the correct operation of Web Tools.

To set the Internet Explorer options, perform the following steps.

1. Open your web browser and select **Tools** > **Internet Options**.
2. Select **General** > **Browsing History** > **Settings**.

3. Choose **Every time I visit the webpage** under **Check for newer versions of stored pages** as shown in the following figure.

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.

To delete these files, perform the following steps.

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 temporary files, the list is empty.

Java installation on the workstation

The 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

To install the JRE on your Solaris or Linux client workstation, perform the following steps.

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

<http://www.oracle.com/technetwork/java/archive-139210.html>

NOTE

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

2. On locating the JRE link, follow the instructions to install the JRE.
3. 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

To install patches on Solaris, perform the following steps.

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

<http://www.oracle.com/technetwork/java/javase/downloads/jre8-downloads-2133155.html>

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

To install the Java plug-in on Windows, perform the following steps.

1. From the **Start** menu, select **Control Panel** and select the **Java 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 8.1.0 or later to install JRE 1.8.0. 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 8.1.0 or later to install JRE 1.8.0. 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, 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.

Enabling Java content in the browser

Launching Web Tools from a browser or Brocade Network Advisor is done using Java Web Start technology. This requires the local system's web browser to be able to run Java web start applications. This setting may have been turned off in the wake of recent Java zero-day vulnerabilities.

To turn on Java content in the browser, perform the following steps.

1. Launch **Java Control Panel**.
2. Click the **Security** tab and select the **Enable Java content in the browser** check box.
3. Click **Apply**.

When the **Windows User Account Control** (UAC) dialog box displays, allow permissions to make the changes.

4. Click **OK** in the Java plug-in confirmation window.

You can now enter the IP address of the switch and launch Web Tools from a browser.

Configuring the Java plug-in for Windows

To configure the Java plug-in for Windows, perform the following steps.

1. From the **Start** menu, select **Control Panel** > **Java**.
2. Click the **Java** tab.
3. In the **Java Applet Runtime Settings** section, click **View**.

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

- Double-click the **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.

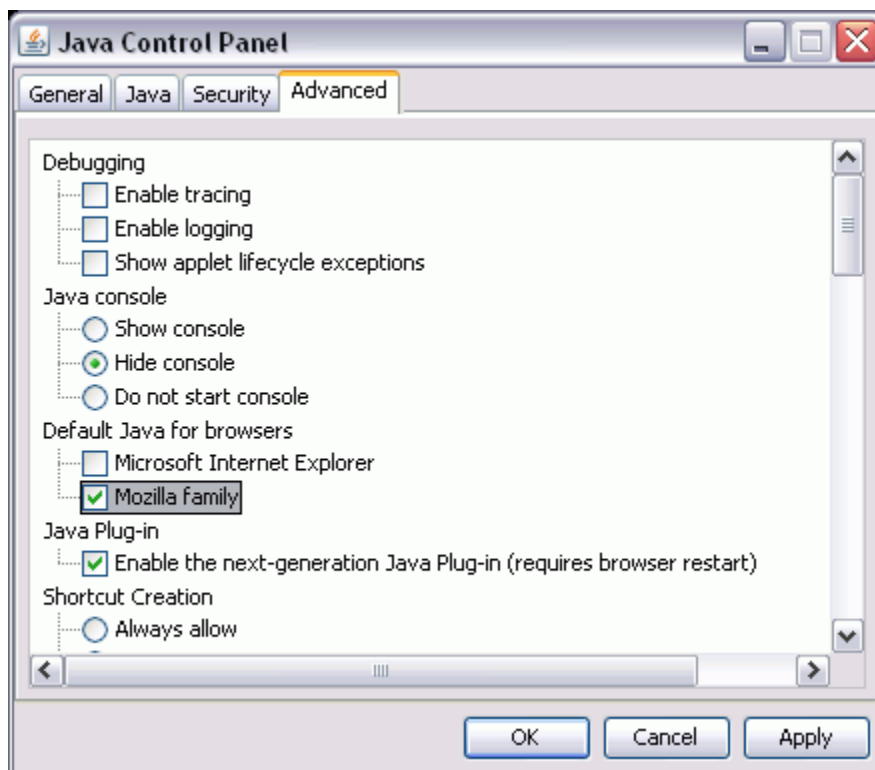
- Click **OK** to apply your settings and close the **Java Control Panel**.

Configuring the Java plug-in for Mozilla family browsers

To configure the Java plug-in for Mozilla family browsers, perform the following steps.

- From the **Start** menu, select **Control Panel**.
- Click the **Advanced** tab and expand the **Default Java for browsers** option, as shown in the following figure.

FIGURE 2 Default Java for browsers option



- Select **Mozilla family** and click **Apply**.
- Click **OK** to apply your settings and close the **Java Control Panel**.

Opening Web Tools

You can open Web Tools on any workstation with a compatible web browser installed. For a list of web browsers that are compatible with the Fabric OS 8.1.0, refer to [System requirements](#) on page 16. Web Tools supports both HTTP and HTTPS.

To open Web Tools, perform the following steps.

1. Open the web browser and enter the IP address of the device in the **Address** field, for example:

```
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 25 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 the *Brocade Fabric OS EZSwitchSetup Administration 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 guide describes only the Web Tools interface.

NOTE

To avoid a potential POODLE attack and establish a secure connection, disable the SSL 3.0 protocol option from your web browser settings.

Server authentication for HTTPS

When you use HTTPS to open Web Tools, the certificate presented by the switch will be validated by Web Tools based on the following parameters:

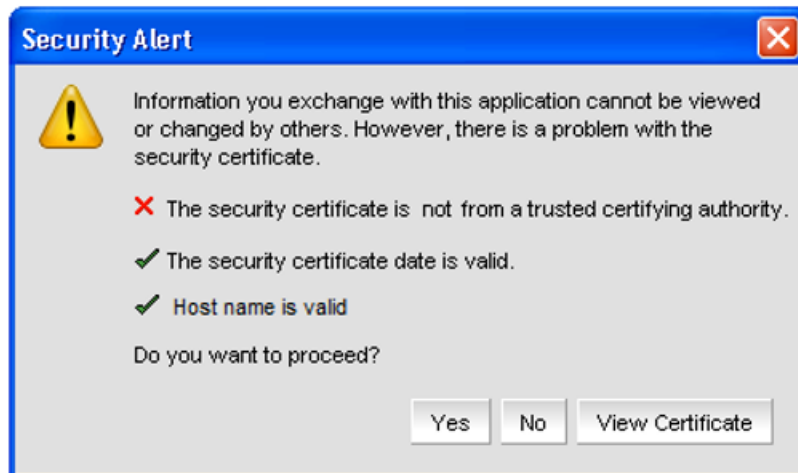
- Trusted certifying authority
- Date
- Host name

The following procedure provides the basic validation required in the security certificate.

Certificate is not issued by a trusted certifying authority

If the certificate issued is not from a trusted CA, you will get a pop-up message that the certificate on the switch is not signed by a trusted certifying authority as shown in the following figure.

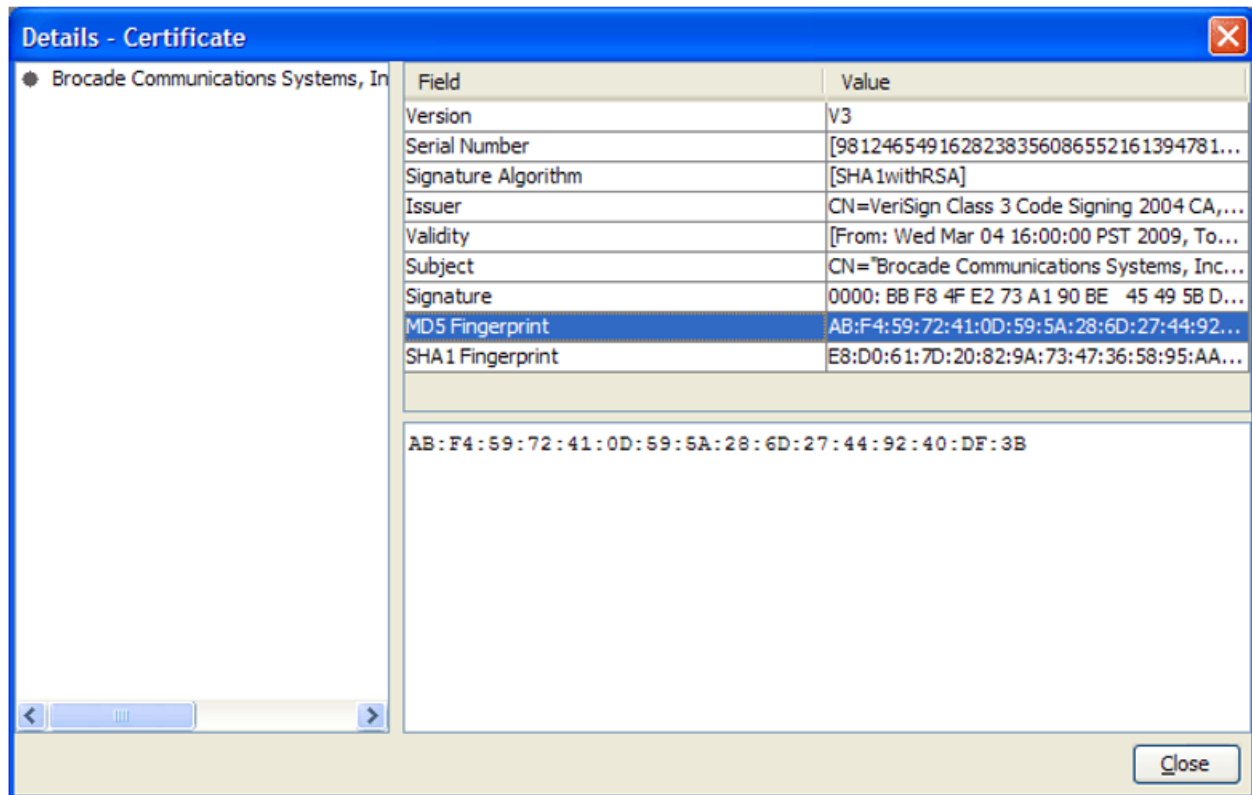
FIGURE 3 Certifying authority invalid dialog box



The figure indicates that the server certificate is not issued by a trusted certifying authority and will give you the following options:

- Yes: Accepts the certificate, and the client launches as usual.
- No: Rejects the certificate and closes the certificate validation error dialog box.
- View Certificate: Launches the certificate details dialog box indicated in the following figure.

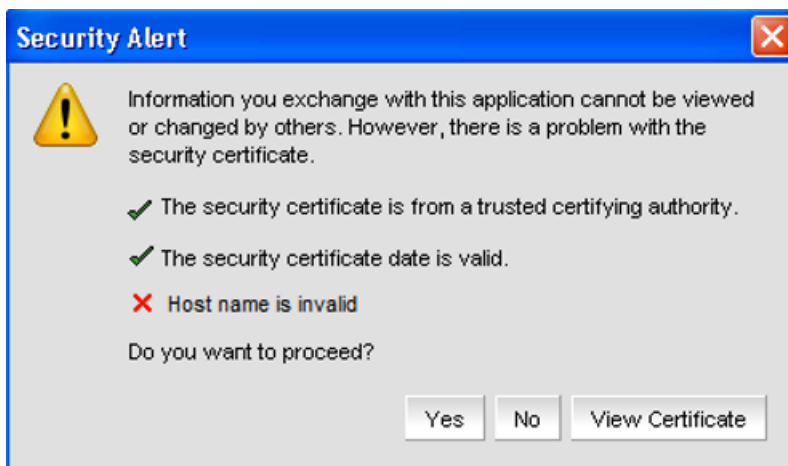
FIGURE 4 Certificate details



Host name in the certificate not matching the switch host name

When you launch Web Tools using a switch IP address or host name, Web Tools will verify the common name of the certificate against the fully qualified domain name obtained from the switch IP address or host name. If they do not match, you will continue to receive the pop-up message because the host name of the switch is not valid as shown in the following figure.

FIGURE 5 Host name invalid dialog box



Certificate date is expired or not valid

Web Tools will verify the “valid from” and “valid till” date of the certificate present on the switch. If either certificate date is not valid or if the certificate has expired, the switch will not be allowed to launch.

NOTE

Web Tools does not display a warning when the security certificate is from a trusted CA with a valid date and host name.

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.

Before 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 log in to Web Tools as a default user with the default password, Web Tools displays a message requesting you to change the password. You will be logged out of Web Tools once the password is changed.

NOTE

You must log in before you can view Switch Explorer (shown in [Figure 7](#) on page 32).

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 for this release:

- Brocade 6510

- Brocade 6520
- Brocade 6548
- Brocade 7840
- Brocade DCX 8510-4 and Brocade DCX 8510-8
- Brocade G620
- Brocade G630
- Brocade X6-4 and Brocade X6-8

To log in to a Virtual Fabric, perform the following steps.

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

You are given a choice between **Home Logical Fabric** and **User Specified Logical Fabric** as shown in the following figure. **Home Logical Fabric** is the default.

FIGURE 6 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 or the context name, and click **OK**.

On providing the context name, a dialog box displays with the available list of VF ID-Context Name (role of the context). You can select the role from the list and log in.

NOTE

The **Options** button is not available for non-VF platforms.

Switching between Virtual Fabrics

To switch between one Virtual Fabric to another, perform the following steps.

1. Log in to Web Tools using the **User Specified Logical Fabric** option.
You can enter the context name to log in.
2. Select the context name that you want to access from the **Logical Switch** list.
The base switch has **Base** appended to the context name.
3. Click **Yes** on the **Switch Virtual Fabric Context** confirmation dialog box.
4. Repeat step 2 and step 3 to switch to another Virtual Fabric.

Logging out

You can end a Web Tools session either by selecting **Manage > Exit** or by closing the **Switch Explorer** window.

You might be logged out of a session involuntarily, without explicitly selecting **Manage > Exit**, under the following conditions:

- You initiate a firmware download from the Web Tools **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. The following table describes these roles.

For information about creating unique user account roles, refer to [User-defined accounts](#) on page 187.

TABLE 4 Predefined Web Tools roles

Role	Description
admin	You have full access to all 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: <ul style="list-style-type: none"> • 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.

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 exit.
- You close the **Switch Explorer** window.
- The session ends due to inactivity (timeout).

A session automatically ends if no information is sent to the switch for more than two hours. Because user keystrokes 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. 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:

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

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 under the Web Tools directory if they are not already present:

- 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
- useractions.txt

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

The maximum size of the webtools.log file is 5 MB. It is rolled into a new file when the 5 MB 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.

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

- Switch name
- Fabric OS version
- Switch type
- Ethernet IPv4 address
- Ethernet IPv4 subnet mask
- Ethernet IPv4 gateway address

The useractions.txt file logs the user actions performed in Switch admin and Zone admin.

Application SupportSave logs

Web Tools allows admin users to collect the following Web Tools-related SupportSave data:

NOTE

Web Tools SupportSave collects logs for troubleshooting Web Tools. To collect switch logs, use the **SupportSave** command through the Fabric OS CLI or the Technical Support feature in Brocade Network Advisor.

- HTML files
- CAL files
- Web Tools logs folder

The admin users are:

- admin
- switchadmin
- zoneadmin
- fabricadmin
- basicswitchadmin
- securityadmin

For non-admin users, only the Web Tools logs folder is collected in SupportSave data.

NOTE

SnmpConfig.html file is no longer collected as a part of SupportSave in Fabric OS 8.1.0 or later.

To collect SupportSave logs, perform the following steps.

1. Click **Tools** > **Application SupportSave**.

The **Application SupportSave** dialog box displays.

2. Click **Browse** to select the location where the Application SupportSave output must be saved.

The default path for Windows is C:\Documents and Settings\<user>\Webtools\<log file>.

The default path for Linux is /root/WebTools/<switch log file>.

3. Click **Capture** to start collecting the Web Tools SupportSave logs.

A zipped Application SupportSave folder is saved in the defined location. The Application SupportSave zip file name format is "WT-SS-switchIP-FID-dd-mm-yy-hh-mm-ss". The Application SupportSave file name will show the VF ID if VF is enabled on the switch.

NOTE

SupportSave collection is terminated if the time exceeds 20 minutes.

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.

Using the Web Tools Interface

• Viewing Switch Explorer.....	31
• Displaying tool tips.....	38
• Right-click options.....	39
• Refresh rates.....	39
• Displaying switches in the fabric.....	40
• Recommendations for working with Web Tools.....	40
• Opening an SSH client window.....	41
• Collecting logs for troubleshooting.....	41

Viewing Switch Explorer

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

- Three tabs: **Switch View**, **Port Admin**, and **Name Server**.

If Access Gateway is enabled, the **Access Gateway Devices** tab displays instead of **Name Server**. For more information, refer to [Using the Access Gateway](#) on page 165.

Fabric Tree under **Switch View** displays a list of all switches in the fabric.

- The menu bar at the top of the window provides access to the following tasks:
 - Viewing tasks, such as the switch view, port administration, and name server.
 - Configuration tasks, such as switch administration and zone administration.

NOTE

You can manage basic zoning and Traffic Isolation zoning using Web Tools. You must use Brocade Network Advisor 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 141.

- Monitoring tasks, such as performance monitoring, system monitoring, and viewing the power, temperature, or fan status.
 - Reporting tasks, such as viewing the status of a switch.
 - Tools tasks, such as opening a Telnet or SSH client window, beaconing a switch or chassis, and accessing SupportSave.
- The buttons below the menu bar provide access to switch information: temperature, power, and fan data.

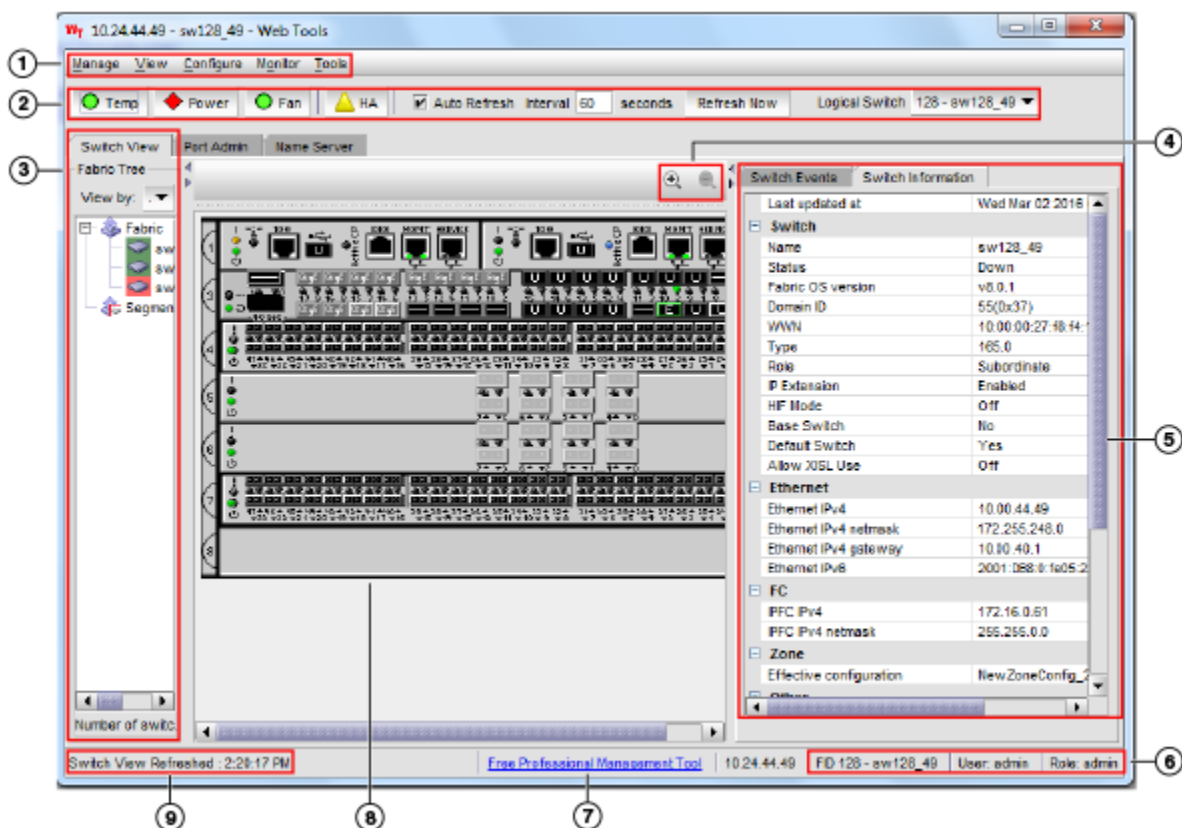
Although clicking a button can open a separate dialog box or window displaying the management tasks, all access control is established when you first log in to the switch.

Buttons in Switch Explorer are unavailable for two reasons: because your account does not have sufficient privileges to access this feature or because your user account does not meet some condition to access the feature.

- 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 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 exit 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, the **Switch View** displays.
- Web Tools displays host time and switch time. The timestamp displayed in the **Last polling refresh time** field in the lower-left corner of the window is the host time and Web Tools is launched. All other timestamps displayed in the application are the switch time. The following are the features that display the host time.
 - Port Admin tab
 - Switch View tab
 - Name Server tab
 - FCR dialog box
 - System Monitoring dialog box
 - Netstat Performance dialog in the Switch Admin tab

FIGURE 7 Switch Explorer



- Menu bar
- Switch View buttons
- Tabs and Fabric Tree
- Zoom in and Zoom out buttons
- Switch Events and Switch Information
- Indicator bar
- Professional Management Tool offering
- Switch View
- Last polling refresh time - The host time is displayed and Web Tools is launched

Persisting GUI preferences

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

If you launch Web Tools from Brocade Network Advisor, all Web Tools GUI persistence data for each user name is stored in the Brocade Network Advisor database.

The **Port Admin** GUI preferences that persist are:

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

The **Switch Admin** GUI preferences that persist are:

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

The Switch Explorer GUI preferences that persist are:

- Last selected tab

The **Name Server** GUI preferences that persist are:

- Table column sorting
- Table column positions

The **Zone Admin** GUI preferences that persist are:

- Basic Zones
- Traffic Isolation Zones
- Last selected tab
- Table column sorting
- Table column positions

Tabs

Switch Explorer provides access to the following three tabs:

- **Switch View**
- **Port Admin**
- **Name Server**-- Name Server information is collected from the selected switch. Refer to ["Displaying the Name Server entries" on page 53](#) for more information.

Fabric Tree

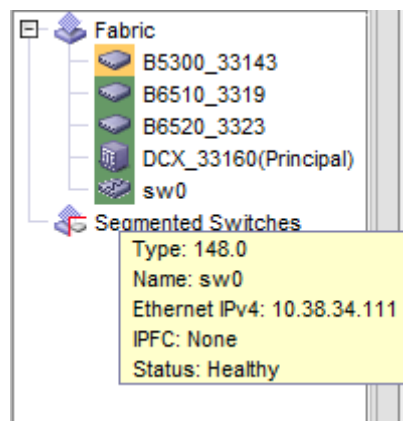
Fabric Tree under **Switch View** displays all switches in the fabric, even those that do not have a Web Tools license. **Fabric Tree** does not display switches segmented before you opened Web Tools.

Three types of switch icons display in **Fabric Tree**: one for a fixed-port switch, one for a chassis, and one for analytics platform. The background color of the switch icon indicates the switch status as follows:

- Green – Switch status is healthy.
- Yellow – Switch status is marginal.
- Red – Switch status is down.
- Gray – Switch status is unmonitored.

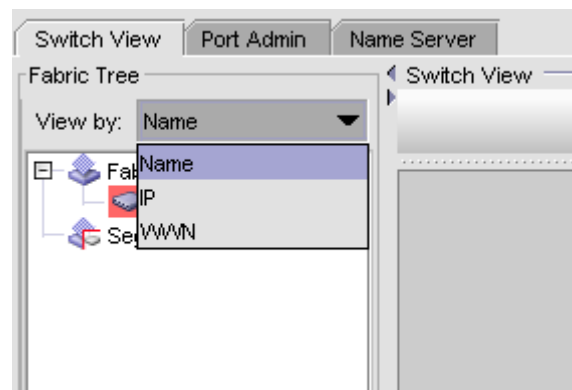
The switch (either principal or subordinate) that is taking the principal role is appended with the term "Principal", as shown in the following figure. Refer to [Setting a principal switch](#) on page 53 for more information on setting the principal switch.

FIGURE 8 Principal switch



Use the menu at the top of the **Fabric Tree** area to view switches in the **Fabric Tree**. Selecting the **View by** drop-down under the **Fabric Tree** section, displays the options like Name, IP, or WWN as shown in the following figure.

FIGURE 9 View by drop-down



By default, **Name** is chosen under **View by** drop-down.

Selecting **IP** > displays the **Fabric Tree** components identified by their IP addresses. If components have IPv6 addresses, they are displayed using IPv6 addresses. If there are IPv4 addresses, then the IPv4 addresses are displayed.

Selecting **WWN** > displays the **Fabric Tree** components identified by their WWN values.

You can rest the pointer over a switch to display the IP address, role, current status, and other details. 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 switches in the fabric, you can manage switches that support Fabric OS 6.1 and later versions because it does not require a Web Tools license. If a switch is launched from **Fabric Tree**, preference will be given to IPv4, even though both IPv4 and IPv6 are configured for that particular switch.

NOTE

Web Tools does not support analytics platforms. When you launch an analytics platform from the **Fabric Tree**, an error message "Web Tools is not supported on Analytical Devices" displays.

Switch View buttons

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

- **Temp**--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.

NOTE

From Fabric OS 8.1.0 or later, the Switch View buttons are updated irrespective of the tab selection on the Switch Explorer.

Switch View

You can click the small right arrow toward 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 [Figure 7](#) on page 32.

NOTE

Blades are graphically represented in the Web Tools GUI. They are vertical in the X6-8 chassis and horizontal in the X6-4 chassis.

The default **Switch View** display refresh rate is 60 seconds. However, the initial display of Switch Explorer may take from 30 to 60 seconds after the switch is booted. Refresh rates are fabric-size dependent. The auto-refresh interval may not be less than 45 seconds. However, the refresh rate varies depending on the activity in the fabric and on the host system you are using. 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.

Port representations

The ports in the **Switch View** show the port type. A colored border indicates the status of the port; for example, a green border indicates that the port is connected and traffic is flowing. A white border indicates that the port has SPF/QSFP inserted but is not online. A violet border indicates that the port is disabled. Ports that are not accessible do not display the port type and do not have borders.

When you pause the pointer over a port, a yellow color toggling effect is available to identify the selected port.

NOTE

A violet border (port disabled) takes precedence over the other colored border indications.

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 156 for more information.

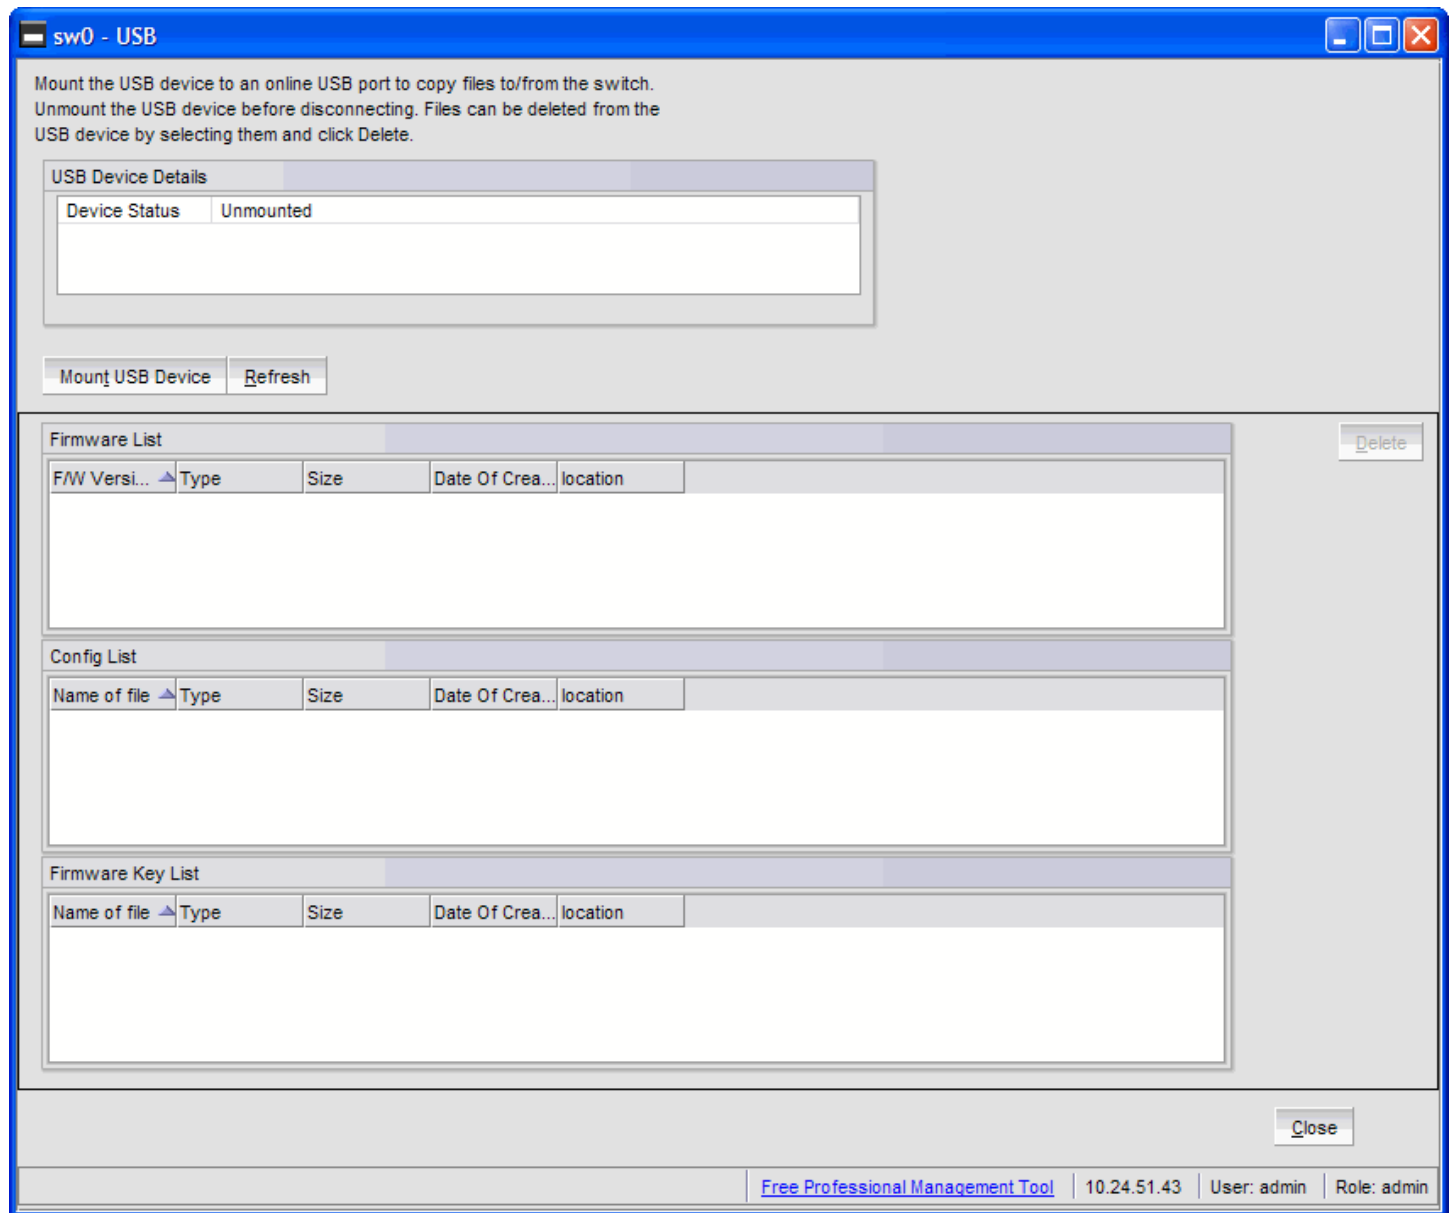
Right-click a port in the **Switch View** to get a menu that opens the **Port Admin** tab, allowing you to view detailed information about the port. From **Port Admin**, you can access information on all other ports. Refer to [Managing Ports](#) on page 79 for more information.

E_Ports are visible in all domains.

USB port representation

For switches with USB ports, the **USB Storage Management** view is launched for USB ports as shown in the following figure.

FIGURE 10 USB port storage management




NOTE

Click the USB port on the Switch View to launch the **USB Storage Management** window.

Zoom in and zoom out



Use the zoom buttons () above the graphical Switch View to magnify the hardware image. Click the zoom in button (+) to see an enlarged view of the switch, and click the zoom out button (-) to see the default view of the switch. Pausing the mouse cursor over the buttons displays the tool tip.

Switch Events and Switch Information

Switch Events and **Switch Information** display as tabs under **Switch View**. The **Switch Events** tab is polled every 15 seconds, and the **Switch Information** tab is polled every 60 seconds. Polling for both tabs is based on the switch time.

NOTE

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

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

- **Last updated at:** Displays the switch time.
- **Switch**
 1. **Name:** Name of the switch.
 2. **Status:** Status of the switch.
 3. **Fabric OS Version:** Fabric OS version of the switch.
 4. **Domain ID:** Domain ID of the switch.
 5. **WWN:** World Wide Name of the switch.
 6. **Type:** Type of the switch.
 7. **Role:** Role of the switch.
 8. **IP Extension:** IP Extension status of extension blades (Enable/Disable).
 9. **HIF:** High Integrity Fabric mode of the switch (Enable/Disable).
 10. The following information is specific to 8G and 16G directors:
 - **Model Type:** Type of the director model.
For the Brocade DCX 8510-4 and Brocade DCX 8510-8, the value is 1. This field is not displayed for any other platform.
 11. The following information is specific to Virtual Fabrics:
 - **Base Switch:** Indicates whether the logical switch can act as a base switch.
 - **Default Switch:** Indicates whether the logical switch is the default logical switch.
 - **Allow XISL Use:** Indicates whether the logical switch is allowed to connect to other logical switches using an extended inter-switch link (XISL).
 - **Inherited FID:** Indicates whether the FID is inherited from the virtual fabric.
- **Ethernet**
 - **Ethernet IPv4:** Ethernet IPv4 address of the switch.
 - **Ethernet IPv4 netmask:** Ethernet IPv4 subnet mask 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 net mask**: Fiber Channel IPv4 subnet mask.
- **Zone**
 - **Effective Configuration**: Indicates whether zone configuration is enabled.
- **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 Mode**: Current status of the Insistent Domain ID mode 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 153.

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. 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** window.

Displaying tool tips

When you pause the pointer over a Web Tools button, the system displays a brief description of the button. If you pause the pointer over most components, the system displays tool tips about the component.

In the **Fabric Tree**, you can pause the pointer over a switch to view its type, subtype, name, Ethernet IPv4 and IPv6 addresses, IPFC, and status.

In **Switch View**, you can pause the pointer over a blade to view the blade ID and 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. You can select a port and pause the pointer over the **Zoom In** and **Zoom Out** buttons to see an enlarged view or the default view of the switch. Firmware versions and IP addressing are displayed for CP blades.

When you pause the pointer over a port, you can view the:

- Port name
- Port ID
- Port beacon
- Port peer beacon
- Port number

- Port index
- Port type (E_Port, F_Port, L_Port, D_Port, EX_Port, GigE port, U_Port, or AE_Port)
- Port status (online or offline)
- Port state (in-sync, no_sync, no light, or no module)

NOTE

The **Port Peer Beacon** status displays only if it is enabled on a port and the **Port beacon** status is not displayed for the same port, as both are mutually exclusive.


NOTE

The port connected to Brocade Analytics Monitoring Platform displays as an AE_Port.

Right-click options

You can right-click a port to quickly perform some basic port administration tasks, as shown in the following figure.

FIGURE 11 Right-click menu for ports (from Switch Explorer)



The tasks are:

- The **Port Admin** option displays the **Port Admin** tab.
- The **Port Details** option displays read-only information about a port, without opening the **Port Admin** tab. You can right-click on the table content to export or copy the information from the **Port Details** window.
- The **Configure** option provides a submenu of options to allow you to rename, enable, and disable ports, and to set persistent enable or disable without opening the **Port Admin** tab.

Refresh 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 60 seconds does not ensure that a refresh occurs every 60 seconds. It ensures that the time between each refresh activity is no more than 60 seconds.

Auto-refresh intervals might not be exactly 60 seconds. The refresh rate varies depending on the activity in the fabric and on the host system that you are using. Following are some variables that you should consider when refreshing the fabric:

- The **Refresh Now** button is disabled for 6 to 8 seconds on every click.
- Retrieval time increases when you are in a large fabric because there is more data to retrieve from the switches.
- The processor speed of the system that you are using may slow down the refresh rate.

- OS-Job Scheduling impacts the refresh rate if you are using a host system in the data center.
- JVM-Performance can cause interval differences between what is on-screen and how long time 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. The following table lists polling rates by module. Although these rates are sample rates, they correctly illustrate variance in the refresh rates throughout Web Tools.

TABLE 5 Polling rates

Module	Polling rate
Name Server	User-defined; 45 sec minimum
Zoning Database	60 sec
Performance Monitor	30 sec
Port Management	User-defined; 45 sec minimum
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 Brocade Network Advisor client as Element Manager. This lets you manage Web Tool requests when the fabric is in a private network.

Launch Web Tools from Brocade Network Advisor 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 4.4.x or later from a **Fabric Tree** displayed for a version earlier than a 4.4.x switch, some of the features may be disabled.

To display switches in the fabric, perform the following steps.

1. Open Web Tools as described in ["Opening Web Tools" on page 10](#) and log in to the switch.
2. If the **Fabric Tree** is not expanded, click the plus sign (+) in the **Fabric Tree** to view all switches in the fabric.
3. Click a switch in the **Fabric Tree**.

A separate browser dialog box displays the selected switch.

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

Recommendations for working with Web Tools

Brocade recommends the following for working with Web Tools:

- If you receive an error when saving changes in the **Switch Administration** window, note the error message, 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 that contains 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 performing fabric-level configuration tasks, wait until you have received confirmation that the changes have been implemented before performing subsequent tasks. For a large fabric, this configuration task 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 to 40 seconds) to collect and display data before taking any action or assuming that the application is "hanging."
- A maximum of five simultaneous HTTP sessions to any one switch is recommended. An HTTP session is considered a Brocade Network Advisor or Web Tools connection to the switch.

Opening an SSH client window

Opening an SSH client window connects to the switch IP interface. You cannot connect to a command CP blade on a director switch through an SSH client window opened from Web Tools, even when the blade has an IP address. Refer to the *Brocade Fabric OS Command Reference* for information about Telnet commands.

To open an SSH client window, perform the following steps.

1. Select a switch in **Fabric Tree**.
You are prompted to log in. Once you do, the selected switch displays in **Switch View**.
2. Select **SSH Client** under the **Tools** menu.
The **Preference** dialog box displays.
3. Enter the SSH path defined for your implementation.
To avoid the need to remember and enter 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.
4. Click **OK**.
The SSH window displays.
5. Enter your user credentials at the login prompt.
6. 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 the **Control Panel** and select **Java**.
2. Click the **Advanced** tab.
3. Expand the **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

• Fabric and switch management overview.....	43
• Configuring IP and subnet mask information.....	47
• Configuring Netstat Auto Refresh.....	47
• Configuring a syslog IP address.....	48
• Removing a syslog IP address.....	48
• Configuring IP filtering.....	48
• Blade management.....	49
• Switch configuration.....	51
• Switch restart.....	54
• System configuration parameters.....	55
• Licensed feature management.....	60
• High Availability overview.....	62
• Event monitoring.....	64
• System Monitor.....	67
• Displaying Name Server entries.....	70
• Physically locating a switch using beaconing.....	71
• Locating logical switches using chassis beaconing.....	72
• Virtual Fabrics overview.....	72
• MAPS limited monitoring support.....	74

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 the following figure.

The **User** tab is editable because most of the information does not require switch membership in the current domain.

FIGURE 12 Switch Administration window, Switch tab, Basic mode

10.24.33.28 - sw0 - Switch Administration

Show Advanced Mode

SwitchName: sw0 DomainID: 10(0xA) VVWN: 10:00:00:27:f8:a7:a1:25 Mon Oct 08 2012 03:17:37 GMT+00:00

Switch Network Firmware Download License User Trunking

Switch Name and Domain ID

Name: sw0 Manufacturer Serial #: ALM0628E00C

Domain ID: 10 Supplier Serial #: none

Switch Status: ☒ Enable ☐ Disable

Switch Persistent: ☒ Enable ☐ Disable switch immediately ☐ Disable when the switch reboots

DNS Configuration

DNS Server 1:

DNS Server 2:

Domain Name:

Remove All

Principal Switch

☐ Set as preferred Principal Switch ☐ Rebuild Fabric after setting preferred principal switch

Priority: 0 (1 - FF)

Reboot/Fastboot

Reboot Fastboot

Report

View Report

Access Gateway Mode

☐ Enable ☒ Disable

[Switch Administration opened]: Mon Oct 08 2012 02:42:20 GMT+00:00

[Switch Administration closed]: Mon Oct 08 2012 02:42:20 GMT+00:00

[Switch Administration closed]: Mon Oct 08 2012 02:42:20 GMT+00:00

[Switch Administration opened]: Mon Oct 08 2012 02:45:22 GMT+00:00

Change current switch settings Mode: Basic Free Professional Management Tool 10.24.33.28 FID 128 - sw0 User: admin Role: admin

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

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. Click **Configure** > **Switch Admin**.

The **Switch Administration** window displays in basic mode, as shown in [Figure 12](#) on page 44. The basic mode displays the "basic" tabs and options.

- Click **Show Advanced Mode** to see all available tabs and options, as shown in the following figure.

FIGURE 13 Switch Administration window, Switch tab, Advanced mode

The screenshot shows the 'Switch Administration' window for switch 'Stinger_6'. The interface is in 'Advanced Mode' and displays various configuration tabs and settings.

Switch Name and Domain ID:

- Name:
- Domain ID:
- Manufacturer Serial #:
- Supplier Serial #:

Switch Status:

- ☒ Enable ☐ Disable

Switch Persistent:

- ☒ Enable ☐ Disable switch immediately ☐ Disable when the switch reboots

DNS Configuration:

- DNS Server 1:
- DNS Server 2:
- Domain Name:
-

Principal Switch:

- ☐ Set as preferred Principal Switch ☐ Rebuild Fabric after setting preferred Principal Switch
- Priority: (1 - FF)

Reboot/Fastboot:

-

Report:

-

Access Gateway Mode:

- ☐ Enable ☒ Disable

Footer:

- Change current switch settings
- Mode: Advanced
- [Free Professional Management Tool](#)
- 10.24.33.6
- FID 128 - Stinger_6
- User: admin
- Role: admin

NOTE

Allow the switch time to respond to your request. Web Tools displays a message that the switch is busy when you overload the switch with many requests at the same time.

Configuring IP and subnet mask information

Before proceeding, collect all the information that you need to configure the Ethernet IP interface. This includes the subnet mask, gateway IP address, IPFC, and subnet mask for your system. When you configure or change the Ethernet IP, subnet mask, gateway IP, 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 FCO for switches that do not support Virtual Fabrics.

To configure the IP and subnet mask information, perform the following steps.

1. Select the **Network** tab.
2. In the appropriate IP address section, enter the IP address that 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 **IPFC Net IP** field, enter the IPFC net IP address.
 - In the **Ethernet Mask** field, enter the Ethernet mask address.
 - In the **IPFC Net Mask** field, enter the IPFC net mask address.
 - In the **Gateway IP** 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 addresses are created for a DCX or X6 chassis (eight per chassis and eight per each installed CP).

Configuring Netstat Auto Refresh

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

To configure Auto Refresh, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Network** tab.
3. Click **Netstat Performance**.
4. Select the **Auto Refresh** check box to automatically refresh the port 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 port details are automatically refreshed, based on the configured time interval. The minimum value is 15 seconds.

Configuring a syslog IP address

The syslog IP address 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 the *Brocade Fabric OS Administration Guide* for more information on configuring the syslog daemon.

To configure a syslog IP address, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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

To remove a syslog IP address, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Network** tab.
3. Select a syslog IP address in the table and click **Remove**.

You can click **Clear All** to remove all 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.

To set up IP filtering, perform the following steps.

1. Open the **Switch Administration** window (in Basic mode) as described in [Opening the Switch Administration window](#) on page 45.
2. Click **Show Advanced Mode**.
3. Select the **Security Policies** tab.
4. Select **IP Filter**.
5. Click **Create Policy**.

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

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

7. Enter the rule order, rule type, source and destination IP addresses, and then modify the service or destination port, protocol, and action as necessary.

Both the source and destination IP addresses are needed for the FWD rule type.

Only the source IP address is needed for the INPUT rule type, since the destination IP address field is disabled.

NOTE

Once the policy is activated, the rules are processed in a top-down sequence.

8. Click **OK**.

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

- 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 DCX 8510-4 and Brocade DCX 8510-8 platforms.

Enabling or disabling a blade

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

NOTE

The blade state is always shown as enabled, even if you perform a blade disable operation. When a blade is set to a disable state, only the ports on the blade are disabled. The blade remains active.

To enable or disable a blade, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.

2. Select the **Blade** tab.

FIGURE 14 Blade tab

172.26.24.135 - Sw135_128 - Switch Administration

SwitchName: Sw135_128 DomainID: 1(0x1) WWN: 10:00:00:05:33:c1:ef:00 Thu Oct 19 2017 08:52:31 GMT+00:00

CEE SNMP Configure Routing Extended Fabric AAA Service Trace FICON CUP Security Policies

Switch Network Firmware Download License User **Blade** Trunking

Set IP Address... Show IP Address... Blade Action ▼

Blade Administration

Slot No	Blade Id	Blade Type	Firmware Version		Blade State	Power Consumption
			Primary	Secondary		
1	97	SW_BLADE	N/A	N/A	ENABLED	140
2	75(FX8-24)	AP_BLADE	7.4.0_bld10	7.4.0_bld10	FAULTY (51)	250
3	99	CORE_BLADE	N/A	N/A	ENABLED	135
4	50	CP_BLADE	7.4.0_bld10	7.4.0_bld10	STANDBY	40

Close Refresh

[Switch Administration opened]: Thu Oct 19 2017 08:51:22 GMT

Blade Administration - Enable/Disable Blade Mode: Advanced Free Professional Management Tool 172.26.24.135 FID 128 - Sw135_128 User: root Role: root

3. Select **Blade Action** > **Enable Blade** for each blade that you want to enable, or select **Blade Action** > **Disable Blade** to disable a blade, and click **Yes** in the confirmation dialog box.

Disabling a blade does not turn the blade off; it disables the ports on the blade. You cannot enable or disable CP blades.

Setting a slot-level IP address

To configure an IP address, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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.

- 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**.

- 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.
- 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.

- Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
- Select the **Blade** tab.
- Click **Show IP Address**.
- Scroll through the list to view all the information.
- When you are finished, click **Close**.

Switch configuration

Use the **Switch** tab of the **Switch Administration** window to perform basic switch configuration. [Figure 12](#) on page 44 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 pause the pointer over the icon, the system displays text that indicates the status of the switch.

Use this procedure to enable or disable a switch.

- Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
- Select the **Switch** tab.
- In the **Switch Status** section, click **Enable** to enable the switch or **Disable** to disable the switch.
- 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.

Enabling and disabling switch persistence

Use this procedure to enable or disable switch persistence. By default, switch persistence is enabled.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Switch** tab.
3. In the **Switch Persistence** section, do one of the following:
 - Click **Enable** to enable the switch persistence.
 - Click **Disable switch immediately** to disable the switch persistent immediately.
 - Click **Disable when the switch reboots** to set the switch persistence in the disabled state and disable switch persistence on reboot.

The switch remains in the enabled or temporarily online state until it reboots. After reboot, the switch goes to the disabled state.

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 name. Switch names must begin with an alphabetic [a-z] or numeric [0-9] character. The remainder of the name can include alphanumeric characters, the hyphen, and the underscore (_). 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 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 45.
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.

To change the switch domain ID, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Disable the switch, as described in [Enabling and disabling a switch](#) on page 51.
3. Select the **Switch** tab.
4. Enter a new domain ID in the **Domain ID** field.

For IMO (Brocade Native mode), the range of valid values is from 1 through 239.
5. Click **Apply**.
6. Enable the switch, as described in [Enabling and disabling a switch](#) on page 51.

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/QSFP serial ID information

To view or print a report, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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.

Setting a principal switch

To set the preference of a switch to become the next principal switch in the fabric, perform the following steps.

NOTE

Principal switch selection is not supported in Access Gateway mode.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Switch** tab.

3. In the **Principal Switch** section, perform one of the following actions to set the switch as a principal switch:
 - Select **Set as preferred Principal Switch**. The switch is set as the principal switch in the next fabric rebuild.
 - Select **Set as preferred Principal Switch** and then select **Rebuild Fabric after setting preferred principal switch**. The switch is set as the preferred principal switch with a rebuild triggered forcefully.
 - Select **Set as preferred Principal Switch** and enter the priority value in a range from 1 through FF in the **Priority** field.

You can change the priority value and select the **Rebuild Fabric after setting preferred principal switch** option from a principal or a subordinate switch. The principal switch selection is based on the factors in the following table.

TABLE 6 Principal switch selection factors

	Priority value with force option	Expected behavior
Subordinate switch	Lesser than principal	Fabric rebuild occurs, and the switch comes up as a principal switch.
	Greater than principal	Fabric rebuild occurs, and the switch remains as a subordinate switch.
	Equal to principal	Fabric rebuild occurs, and the principal switch is selected based on the WWN check (lower becomes principal switch).
Principal switch	Lesser than subordinate	No fabric rebuild occurs, and the switch remains principal.
	Greater than subordinate	Fabric rebuild occurs, and the switch becomes subordinate.
	Equal to subordinate	Fabric rebuild occurs, and the principal switch is selected based on the WWN check (lower becomes principal switch).

4. Click **Apply**.

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).

To perform a fast boot, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Click **Fastboot**.
3. On the **Fastboot Confirmation** window, click **Yes** to continue.
4. Click **Apply**.

Performing a reboot

To reboot the CP and issue the normal power-on booting sequence, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.

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
- CSCTL QoS mode settings

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 in to the switch, the PID is bound to the device WWN. If the device is moved to another port in the same switch, or if 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 56.

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 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 the following table 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 8510-4	Enabled	Yes, if dynamic area addressing is enabled in the default switch.	Yes	0	If 8-bit dynamic mode is enabled, FMS is not supported
DCX 8510-8	Enabled			1	Can be set.
X6-4				2	Not supported.
X6-8					
Brocade 6510	Enabled	Yes	Yes	Configurable Default-8 bit dynamic	
Brocade 6520	Disabled	N/A	N/A		
Brocade 7840					

Configuring fabric settings

To configure the fabric settings, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Click **Show Advanced Mode**.
3. Select the **Configure** tab.
4. Select the **Fabric** subtab.
5. Make the fabric parameter configuration changes.
6. Click **Apply**.
7. Enable the switch as described in [Enabling and disabling a switch](#) on page 51.

Fabric settings

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

TABLE 8 Fabric settings

Setting	Description
Address mode	Displays the addressing mode present in the switch.
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 of valid values is from 1 through 27.
Datafield size	The largest possible data field size (in bytes). The range of valid values is from 256 through 2112.
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.
Dynamic Port Name	Displays the switch name, port type, port index, and alias name as part of port name for all port types.
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).
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.
Per-Frame Routing Priority	Select whether to enable 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.
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.
Sequence Level Switching	Select this check 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.
Suppress Class F Traffic	Applies only if VC-encoded address mode is also set. When selected, translatable addressing (which allows private devices to communicate with public devices) is disabled.
WWN-based Persistent PID	Set this mode to configure a PID persistently using a device's WWN. When the device logs in to the switch, the PID is bound to the device WWN. Refer to WWN-based persistent PID assignment on page 55.

Enabling insistent domain ID mode

To enable insistent domain ID mode, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Disable the switch as described in [Enabling and disabling a switch](#) on page 51.

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 51.

Configuring virtual channel settings

You can configure parameters for eight virtual channels (VCs) 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.

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

To configure the virtual channel settings, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Disable the switch as described in [Enabling and disabling a switch](#) on page 51.
3. Select the **Configure** tab.
4. Select the **Virtual Channel** subtab.
5. Enter a value in the **VC Priority** field that you want to change.
The only valid numeric values for all fields are "2" and "3".
6. Click **Apply**.
7. Enable the switch as described in [Enabling and disabling a switch](#) on page 51.

Configuring arbitrated loop parameters

To configure arbitrated loop parameters, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Disable the switch as described in [Enabling and disabling a switch](#) on page 51.
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 51.

Arbitrated loop parameters

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

TABLE 9 Arbitrated loop parameters

Parameter	Description
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.

To configure system services, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Disable the switch as described in [Enabling and disabling a switch](#) on page 51.
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 51.

Configuring CSCTL QoS mode

You can configure switch-level Class-Specific Control (CSCTL) Quality of Service (QoS) mode.

To configure CSCTL QoS mode, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Configure** tab.
3. Select the **CSCTL QoS Mode** subtab.
4. Select either of the following options:
 - **Default** - This is the default option. The **Default** option clears any previously configured CSCTL to VC mapping and sets one-to-one mapping between the CSCTL value and the VC number.
 - **Auto** - This option maps the CSCTL value to more than one VC.
5. Click **Apply**.

Dynamic Port Name configuration

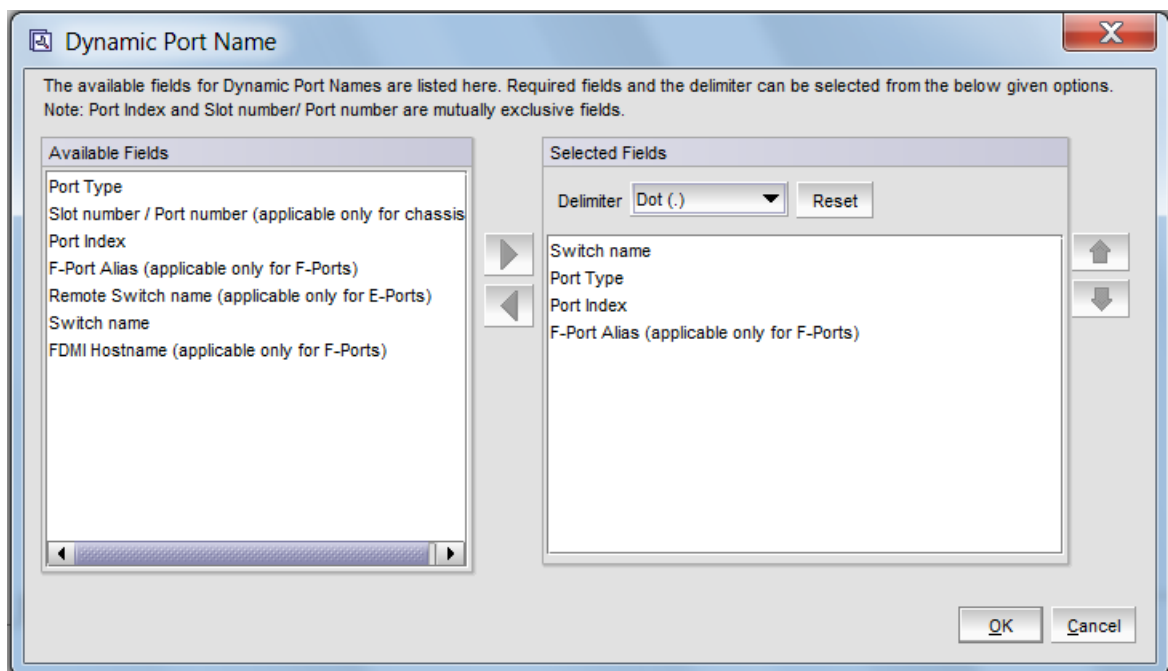
You can configure the Dynamic Port Name feature to display any available field as part of the port name.

By default, the supported list has switch name, port type, port index, and alias name. You must select at least one field in the supported list to set the dynamic port name format. The available port name fields are Switch name, Port Type, Port Index, slot number/port number, F-Port Alias, FDMI Hostname, and Remote Switch name. The supported delimiters are period (.), dash (-), and underscore (_).

To enable Dynamic Port Name configuration, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Configure** tab.
3. Select the **Fabric** subtab.
4. Enable select the **Dynamic Port Name** check box.
5. You can do any of the following to configure the dynamic port name:
 - If you want to use the default fields as part of the port name, go to step 6.
 - If you want to modify the fields in the dynamic port name format, click **Configure**. The **Dynamic Port Name** window opens.

FIGURE 15 Dynamic Port Name window



- Move the fields between the **Available Fields** and **Selected Fields** list as required using the left and right arrows. Rearrange the fields in the **Selected Fields** pane using the up and down arrows.
 - Select a **Delimiter** from the list in the **Selected Fields** pane. By default, a dot (.) is the delimiter.
 - Click **OK**.
6. Click **Apply**.
Dynamic Port Name is enabled on the switch.

Slow-Drain Device Quarantine configuration

The Slow-Drain Device Quarantine (SDDQ) feature is used to reduce unnecessary side effects (back pressure) caused by slow-draining devices. The SDDQ-supported switches that quarantine the slow-drain devices affect the switch CPU performance. The number of slow-drain devices to be quarantined is limited and is a configurable parameter.

Beginning with Fabric OS 7.4.0, Web Tools allows slow-drain device quarantine configuration. SDDQ is not supported in Access Gateway mode.

To configure the SDDQ limit, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Configure** tab.
3. Select the **Slow-Drain Device Quarantine** subtab.
4. Set the **Chassis SDDQ Limit (0-32)** value.
5. Click **Apply**.
SDDQ is configured on the switch.

NOTE

The default slow-drain limit is 10. You can configure up to 32 ports.

Licensed feature management

The licensed features currently installed on the switch are listed in the **License** tab of the **Switch Administration** window. If the feature is listed, 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 expiration date is included. Right-click a license key to export data, copy data, or search the table.

NOTE

If there are no licenses installed on a switch, a message displays, "No License Installed" in the License Administration table.

Activating a license on a switch

Before you can unlock a licensed feature, you must obtain a license key. Use the license key provided in the paperpack document supplied with switch software, or refer to the *Brocade Fabric OS Administration Guide* for instructions on how to obtain a license key at the Brocade website (mybrocade.com).

To activate a license, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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 Brocade 7800upgrade license, do not take effect until the switch is restarted.

NOTE

For the Ports on Demand license, **Capacity** displays the maximum number of ports supported by license.

Assigning slots for a license key

The slot-based licensing feature allows you to increase the capacity without disrupting the slots that already have licensed features running.

NOTE

You can enable slot-based licenses such as FTR_UPG1 and FTR_UPG2 only on the 10 Gigabit Ethernet (FTR_10G), Advanced Extension (FTR_AE), and Advanced FICON Acceleration (FTR_AFA) features.

The Brocade 7840 Extension Switch requires the WAN Rate Upgrade 1 and WAN Rate Upgrade 2 licenses with the capacity value set to 1.

To assign slots for a license key, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **License** tab.
3. Select the license key for which you want to assign slots from the **License Administration** table.
4. Click **Assign Slot(s)**.

The **Assign Slots** window displays.

5. Select the slots that you want to assign.
6. Click **OK**.

NOTE

The **Assign Slot(s)** option is not supported in fixed-port switches.

Removing a license from a switch

To remove a license from a switch in the **Switch Administration** window, perform the following steps.

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 45.
2. Select the **License** tab.
3. Select the license key that you want to remove.
4. Click **Remove**.

Universal time-based licensing

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

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

- Fabric
- Fabric Vision
- Extended Fabric
- Performance Monitor

- Trunking
- High-Performance Extension over FCIP/FC
- Advanced Extension
- Advanced FICON Acceleration
- FICON Management Server (CUP)
- 10 GbE
- Integrated Routing
- Integrated Routing Ports on Demand
- Enterprise Inter-Chassis Link (EICL) license

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 DCX 8510-4, Brocade DCX 8510-8, Brocade X6-4, and Brocade X6-8 platforms.

Refer to the *Brocade Fabric OS Administration Guide* for additional information about High Availability.

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

The background color of the HA button indicates the overall status of High Availability on the switch.

- 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.

Launching the High Availability window

To launch the **High Availability** window, perform the following steps.

1. Select a Brocade director from the **Fabric Tree**.

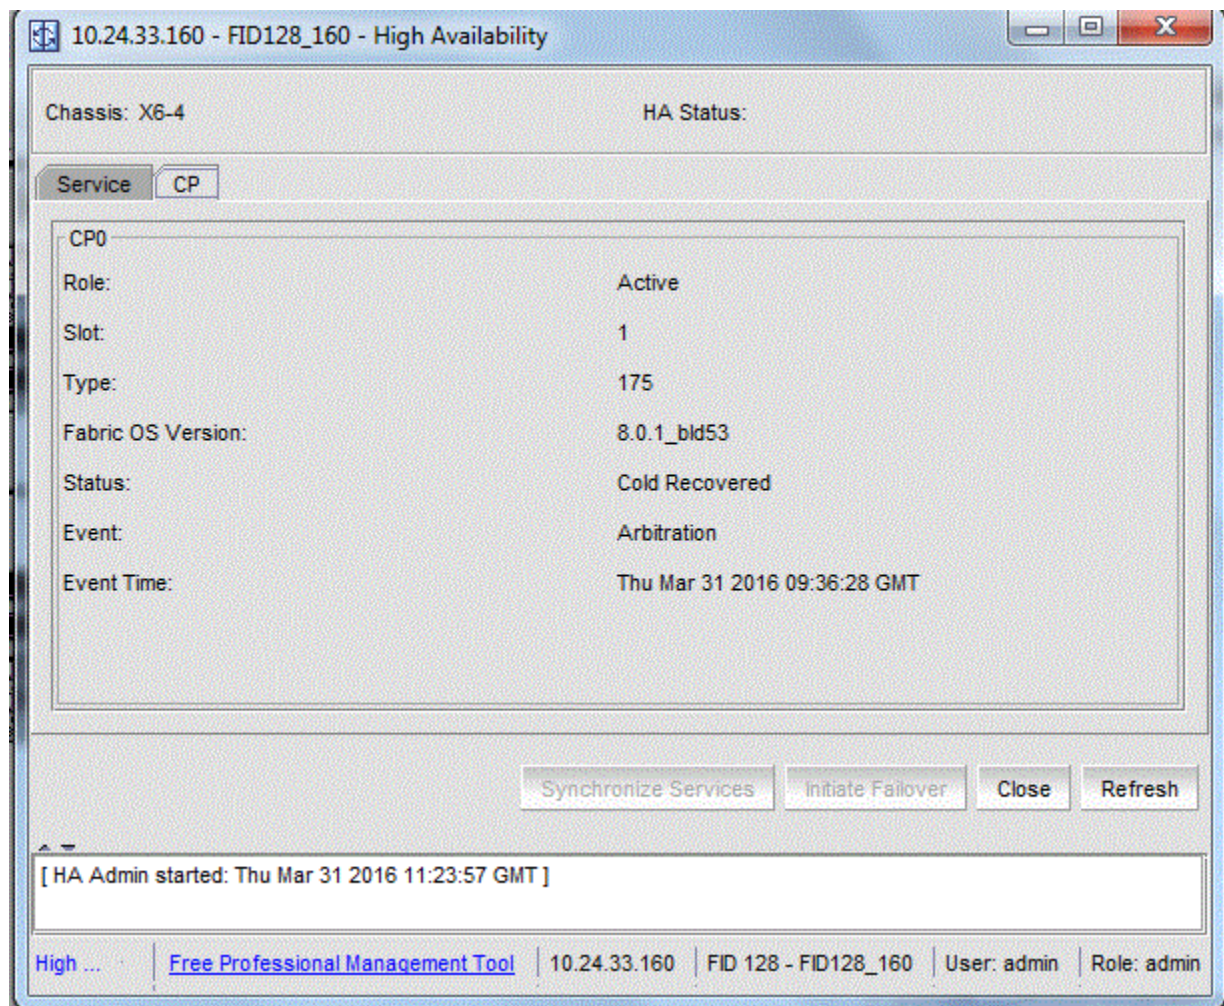
The **Switch View** displays.

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

The **High Availability** window displays. 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 the Brocade DCX 8510-4, CP blades are placed in slot 4 and slot 5. For the Brocade DCX 8510-8, CP blades are placed in slot 6 and slot 7.

FIGURE 16 High Availability window, CP tab



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

Synchronizing services on the CP

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

To synchronize services on the CP, perform the following steps.

1. Open the **High Availability** window.
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, then the services are in sync.
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 synchronize the 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 Web Tools windows and all associated child windows are invalidated. You must close all Web Tools windows and open Web Tools again.

To initiate a nondisruptive failover, perform the following steps.

1. Open the **High Availability** window.
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
- 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 six message severity levels:

- Critical
- Alert

- Error
- Warning
- Information
- Debug

The following table lists the event message severity levels displayed on the **Switch Events** tab and explains what qualifies event messages to be at certain levels.

On the **Switch Events** tab, you can click **Filter** to launch the **Event Filter** dialog box. The **Event Filter** dialog box allows you to define which events should be displayed on the **Switch Events** tab. For more information on filtering events, refer to [Filtering switch events](#) on page 65.

TABLE 10 Event severity levels

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 they are not corrected immediately. For example, a power supply failure or a rise in temperature must receive immediate attention.
Alert	This event does not compromise data or prevent the use of the system; however, the event warrants your 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.
Information	Information-level messages report the current nonerror status of the system components, such as the online and offline status of a fabric port.
Debug	Debug messages deliver status messages relating to debugging systems.

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 for convenience in the event list of each logical switch .

To display switch events, perform the following steps.

1. Select the switch from the **Fabric Tree**.
The **Switch View** displays.
2. Select the **Switch Events** tab, if necessary.

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 unfiltered information.

To filter switch events, perform the following the procedure.

1. Open the **Switch Events** tab as described in [Displaying switch events](#) on page 65.

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

To filter events by event severity levels, perform the following steps.

1. Open the **Switch Events** tab as described in [Displaying switch events](#) on page 65.
 2. Click **Filter**.
- The **Event Filter** dialog box displays.
3. Select **Level**.
 4. Select the event levels that 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

To filter events by message ID, perform the following steps.

1. Open the **Switch Events** tab as described in [Displaying switch events](#) on page 65.
 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). 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

To filter events by service component, perform the following steps.

1. Open the **Switch Events** tab as described in [Displaying switch events](#) on page 65.

2. Click **Filter**.

The **Event Filter** dialog box displays.

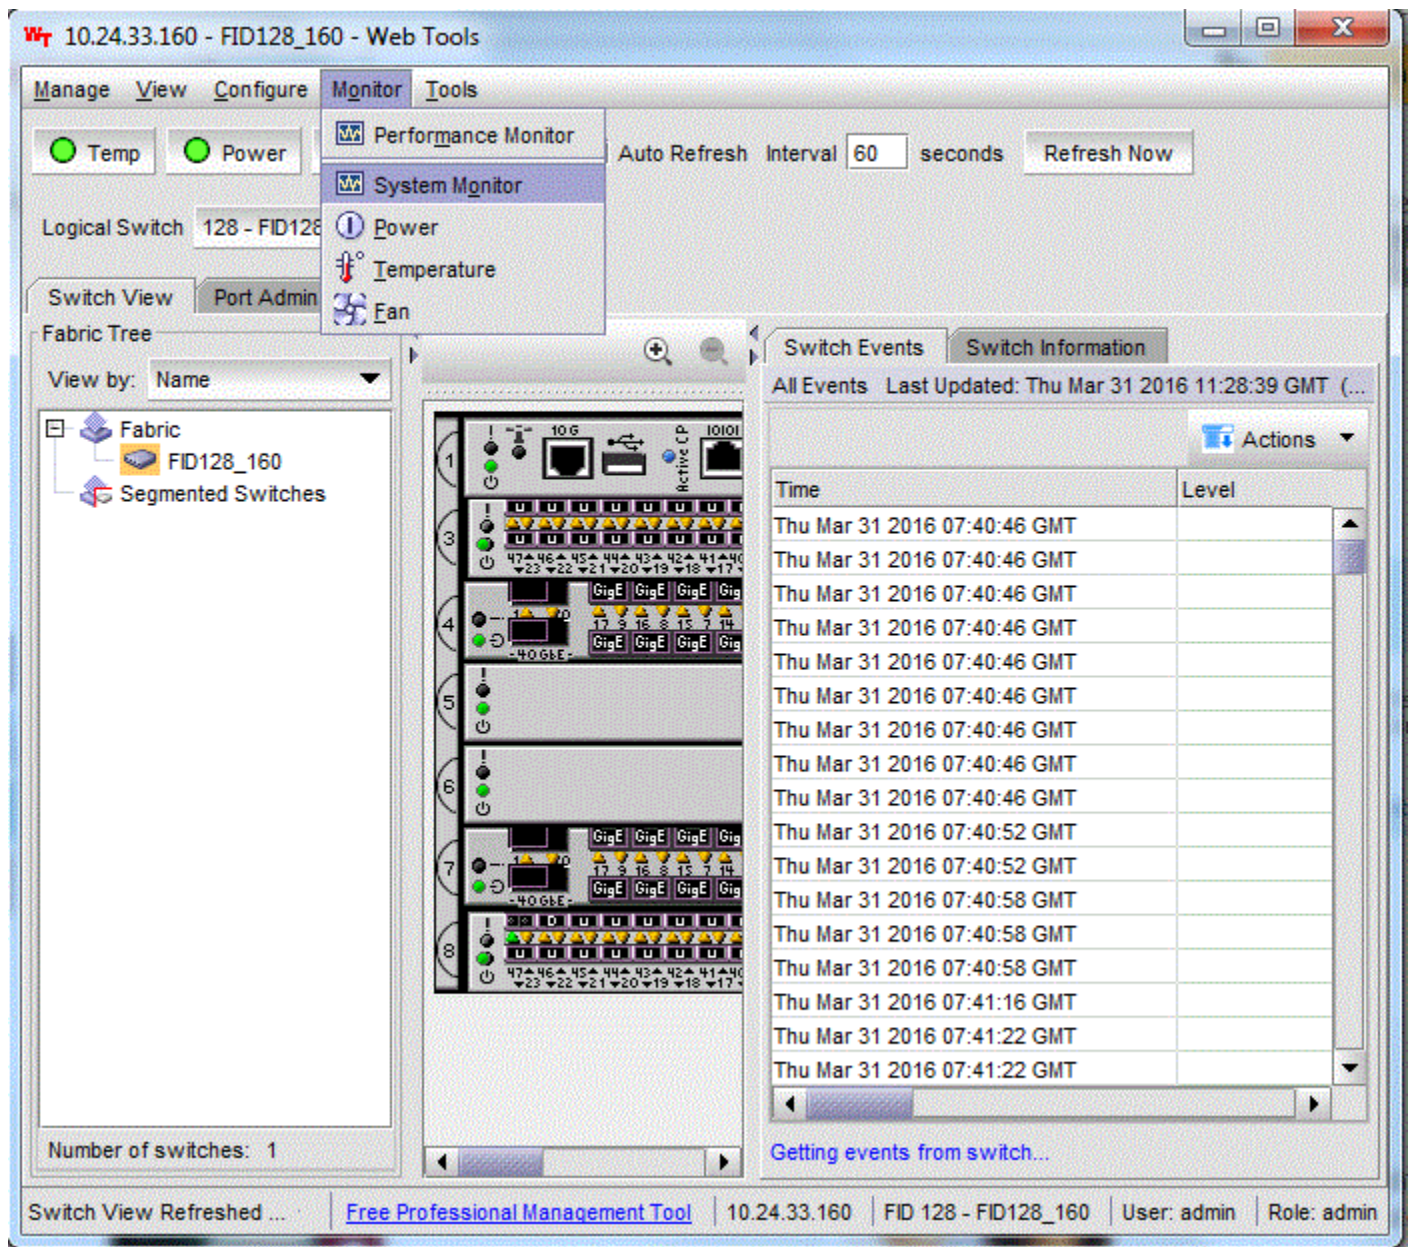
3. Select **Service**. The event service menu is enabled.
4. Select either **Switch** or **Chassis** from the menu to show only those messages.
5. Click **OK**.

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

System Monitor

The Web Tools System Monitor allows you to monitor memory usage and the CPU usage on the switch.

FIGURE 17 System Monitor



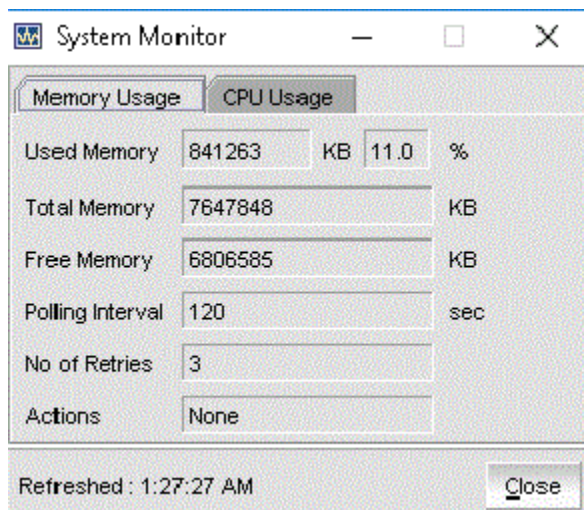
Monitoring memory usage

To monitor memory usage, perform the following steps.

1. Select **Monitor > System Monitor**.

The **System Monitor** dialog box displays as shown in the following figure.

FIGURE 18 Memory Usage tab



2. Select the **Memory Usage** tab. The following fields are displayed:
 - **Used Memory (KB %)** indicates the memory usage in KB and in % for the available resources.
 - **Total Memory (KB)** indicates the total memory of available resources.
 - **Free Memory (KB)** indicates the free memory of available resources.
 - **Polling Interval (sec)** indicates the polling interval in seconds. The default value is 120 seconds.
 - **No of retries** indicates the number of retries. The default value is 3.
 - **Actions** indicates the actions to be taken if system resources fall outside the boundaries defined by the high and low thresholds.


Monitoring CPU usage

To monitor CPU usage, perform the following steps.

1. Select **Monitor > System Monitor**.

The **System Monitor** dialog box displays as shown in the following figure.

FIGURE 19 CPU Usage tab



2. Select the **CPU Usage** tab. The following fields are displayed:
 - **CPU Usage (%)** indicates the CPU usage as a percentage of available resources.
 - **Polling Interval (sec)** indicates the polling interval in seconds. The default value is 120 seconds.
 - **No of retries** indicates the number of retries. The default value is 3.
 - **Actions** indicates the actions to be taken if system resources fall outside the boundaries defined by the high and low thresholds.

Displaying 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 you can drag the column divider to resize a column. You can also right-click a column heading to resize one or all columns, to sort the information in ascending or descending order, or to select which columns are displayed.

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

To display the Name Servers, perform the following steps.

1. In the **Switch Explorer** window, select **Name Server**.
The **Name Server** tab displays.
2. To set an auto-refresh rate for the **Name Server** entries, select the **Auto Refresh** check box and enter an auto-refresh interval (in seconds).

The minimum interval is 45 seconds, and the default interval is 60 seconds.

NOTE

The Name Server information is updated only when you refresh the **Name Server** tab.

Printing Name Server entries

To set up printing preferences, perform the following steps.

1. In the **Switch Explorer** window, select **Name Server**.
The **Name Server** tab 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

To display Name Server information for a particular device, perform the following steps.

1. In the **Switch Explorer** window, select **Name Server**.
The **Name Server** tab 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

To display zone members for a particular device, perform the following steps.

1. In the **Switch Explorer** window, select **Name Server**.
The **Name Server** tab 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 a role with admin permissions to initiate beaconing on a switch or chassis.

To use beaconing, perform the following steps.

1. Select a logical switch from the **Logical Switch** list in the top-right corner of the **Switch Explorer** window.
The selected switch displays in the **Switch View**.

2. Select **Tools > Beacon > Switch Beacon** for a switch, or Select **Tools > Beacon > Chassis Beacon** for a chassis.
 - When a logical switch is selected, the switch beacon flashes the respective ports owned by that logical switch (which can span blades on directors).
 - When **Chassis Beacon** is selected, all LEDs flash in a specific pattern for chassis mode.
3. Look at the physical switches in your installation location to identify the switch.

Locating logical switches using chassis beaconing

For Brocade DCX 8510-4 and Brocade DCX 8510-8 chassis, all LEDs on the chassis illuminate in sequence to identify the chassis.

For Brocade X6 chassis, select **Tools > Beacon > Chassis Beacon**. The two white beacon LEDs on each CP illuminate, and all port LEDs flash green in unison.

Virtual Fabrics overview

Virtual Fabrics is an architecture that virtualize hardware boundaries. Traditionally, SAN design and management are done at the granularity of a physical switch. Each switch and all 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 6510
- Brocade 6520
- Brocade 7840
- Brocade G620
- Brocade G630
- Brocade DCX 8510-4
- Brocade DCX 8510-8
- Brocade X6-4
- Brocade X6-8

Virtual Fabrics cannot be configured or managed from Web Tools. Configuration and management are done from either the Brocade Network Advisor or the Fabric OS command-line interface. For information about configuring and managing Virtual Fabrics, refer to the *Brocade Network Advisor User Manual* if you are using Brocade Network Advisor, or the *BrocadeFabric OS Administration 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 [Introducing Web Tools](#) on page 15, 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 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 **Switch View**.
- **Allow XISL Use** indicates whether 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 the CLI using the **configure** command.

To select a logical switch, perform the following steps.

1. Use the **Logical Switch** list to select the fabric ID.
2. Click **Yes** to confirm.

The selected logical switch displays.

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 Admin** tab.

1. Select the **Port Admin** tab. The **Port Admin** tab displays. Logical ports are displayed in the **FC Ports Explorer** tree structure.
2. To view logical port properties, expand the **Logical Ports** branch, and select a port. The **General** properties are displayed.

FIGURE 20 Logical ports

Port ID	Port Name	Port Type	Port Protocol	Speed (Gbps)	Port Status	Additional Port Info	Health
0(dx0)	sw03 port0	U-Port	FC	N/A	No_Light		OFFLINE
1(dx1)	sw03 port1	U-Port	FC	N/A	No_Light		OFFLINE
2(dx2)	sw03 port2	U-Port	FC	N/A	No_Module		OFFLINE
3(dx3)	sw03 port3	U-Port	FC	N/A	No_Module		OFFLINE
4(dx4)	sw03 port4	U-Port	FC	N/A	No_Module		OFFLINE
5(dx5)	sw03 port5	U-Port	FC	N/A	No_Module		OFFLINE
6(dx6)	sw03 port6	U-Port	FC	N/A	No_Module		OFFLINE
7(dx7)	sw03 port7	U-Port	FC	N/A	No_Module		OFFLINE
8(dx8)	sw03 port8	U-Port	FC	N/A	No_Module		OFFLINE
9(dx9)	sw03 port9	U-Port	FC	N/A	No_Module		OFFLINE
10(dxA)	sw03 port10	U-Port	FC	N/A	No_Module		OFFLINE
11(dxB)	sw03 port11	U-Port	FC	N/A	No_Module		OFFLINE
12(dxC)	sw03 port12	U-Port	FC	N/A	No_Module		OFFLINE
13(dxD)	sw03 port13	U-Port	FC	N/A	No_Module		OFFLINE
14(dxE)	sw03 port14	U-Port	FC	N/A	No_Module		OFFLINE
15(dxF)	sw03 port15	U-Port	FC	N/A	No_Module		OFFLINE
16(dx10)	sw03 port16	U-Port	FC	N/A	No_Module		OFFLINE
17(dx11)	sw03 port17	U-Port	FC	N/A	No_Module		OFFLINE
18(dx12)	sw03 port18	U-Port	FC	N/A	No_Module		OFFLINE
19(dx13)	sw03 port19	U-Port	FC	N/A	No_Light		OFFLINE
20(dx14)	sw03 port20	U-Port	FC	N/A	No_Module		OFFLINE
21(dx15)	sw03 port21	U-Port	FC	N/A	No_Module		OFFLINE
22(dx16)	sw03 port22	U-Port	FC	N/A	No_Module		OFFLINE
23(dx17)	sw03 port23	U-Port	FC	N/A	No_Module		OFFLINE
24(dx18)	sw03 port24	U-Port	FC	N/A	No_Module		OFFLINE
25(dx19)	sw03 port25	U-Port	FC	N/A	No_Module		OFFLINE
26(dx1A)	sw03 port26	U-Port	FC	N/A	No_Module	Speed Mismatch / Inc.	OFFLINE
27(dx1B)	sw03 port27	U-Port	FC	N/A	No_Module		OFFLINE
28(dx1C)	sw03 port28	U-Port	FC	N/A	No_Light		OFFLINE
29(dx1D)	sw03 port29	U-Port	FC	N/A	No_Light		OFFLINE
30(dx1E)	sw03 port30	U-Port	FC	N/A	No_Module		OFFLINE
31(dx1F)	sw03 port31	U-Port	FC	N/A	No_Module		OFFLINE
32(dx20)	sw03 port32	U-Port	FC	N/A	No_Module		OFFLINE

MAPS limited monitoring support

Web Tools provides MAPS limited support to monitor switch performance and switch status.

Beginning with Fabric OS 7.4.0, Fabric Watch is not supported for monitoring the performance and status of switches. You can use Monitoring and Alerting Policy Suite (MAPS) unlicensed support for monitoring the performance and status of switches. You can perform a firmware upgrade from Fabric OS 7.3.0 to Fabric OS 7.4.0 without the Fabric Watch license. MAPS is enabled implicitly for monitoring the unlicensed features. To monitor the port status, you must have the Fabric Vision license installed on the switch.

MAPS is enhanced to monitor the following switch components or features without a license:

- FRU
- Flash
- Fan
- Power
- Temperature
- CPU usage
- Memory usage

NOTE

Because Fabric Watch is no longer supported, switch status policy and switch status information are not displayed.

Maintaining Configurations and Firmware

• Creating a configuration backup file.....	75
• Restoring a configuration.....	76
• Uploading and downloading from USB storage.....	77
• Performing a firmware download.....	77


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)" or "Secure File Transfer Protocol (SFTP)", you cannot specify "anonymous" in the **User Name** field.

An **info** link is enabled when USB is chosen as the source of the configuration file. If you click on **info**, an information message displays as shown in the following figure.

FIGURE 21 Information dialog box



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

1. Select **Configure > Switch Admin**.

The **Switch Administration** window displays.

2. Select **Show Advanced Mode**.
3. Select the **Configure** tab.
4. Select the **Upload/Download** tab.

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.

The default path for Windows is **Folder Name\FileName.txt** or **FileName.txt**.


The default path for Linux is **Directory Name\FileName.txt** or **FileName.txt**.

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

- Use the **Fabric ID** selector to select the fabric ID of the logical switch from which the configuration file is to be 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 77 for more information.

- 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 that 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)" or "Secure File Transfer Protocol (SFTP)," you cannot specify "anonymous" in the **User Name** field.

To restore a configuration, perform the following task.

- Open the **Switch Administration** window.
- Select **Show Advanced Mode**.
- Select the **Configure** tab.
- Select the **Upload/Download** tab.

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

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

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

An **info** link is enabled when USB is chosen as the source of the configuration file. If you click **info**, an information message displays.

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

7. Use the **Fabric ID** selector to select the fabric ID of the logical switch to which the configuration file is to be 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.
8. 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 77 for more information.

9. Click **Apply**.

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

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.

To update your USB storage, perform the following steps.

1. Select **Mount USB Device**, and select **Yes** at the confirmation prompt.
2. Right-click a configuration file to access **Export**, **Copy**, and **Search** 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 Web Tools windows and log in again to avoid a firmware version mismatch. Note that for chassis-based switches, you might get pop-up messages that imply that the loss of connection is temporary and will soon be resolved. You must still close all windows and log in again.

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.

To download a new firmware version, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Firmware Download** tab.
3. Select **Firmware**.

The download source can be located on the **Network** or **USB** device.

NOTE

When you select the **USB** button, you can specify only a firmware path or directory name. The default path for Windows and Linux is /usb/usbstorage/Brocade/firmware/<version>. 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
```

In this syntax, 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, for example:

```
//directory_1/my_directory/v8.1.0/release.plist
```

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

If you select "Secure Copy Protocol (SCP)" or "Secure File Transfer Protocol (SFTP)", 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 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.

Managing Ports

• Port management overview.....	79
• Configuring FC ports.....	85
• Assigning a name to a port.....	89
• Port beaconing.....	90
• Port peer beaconing.....	90
• Enabling and disabling a port.....	92
• Persistently enabling and disabling ports.....	92
• Configuring NPIV ports.....	95
• Enabling Target Driven Zoning Mode.....	95
• Port activation.....	96
• Port swapping index.....	98
• Configuring port binding.....	100
• Configuring BB credits on an F_Port.....	102
• Configuring ALPA	102
• Configuring the port octet speed combination	103
• Configuring CSCTL.....	104
• Configuring compression and encryption.....	106
• Forward Error Correction.....	107
• In-band management.....	108
• GigE port modes.....	109

Port management overview

This chapter describes how to manage FC and Gigabit Ethernet (GbE) ports. Refer to [Viewing EX_Ports](#) on page 162 for information on how to view and configure EX_Ports.

The **Port Admin** tab is refreshed automatically every 60 seconds and is refreshed immediately after 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 fabricadmin. If you are logged in with a user, securityadmin, or zoneadmin role, you can only view the port information.

For information about creating unique user account roles, refer to [User-defined accounts](#) on page 187.

Opening the Port Admin tab

Select the **Port Admin** tab in the **Switch Explorer** window. The **Port Admin** tab displays information about the ports on the switch. Refer to [Switch View](#) on page 35 for information about accessible ports.

The **Port Admin** tab displays in **Basic** mode. To view more port management options, select **View > Advanced**.

NOTE

You can drag the column divider to resize a column, or you can drag columns to re-arrange them in a custom order. You can also right-click a column heading to resize one or all columns or to sort the information in ascending or descending order.

Port Admin tab components

The **Port Admin** tab has the following subtabs in the top left corner:

- **FC Ports** displays all FC ports on the switch (physical FC ports and logical ports).

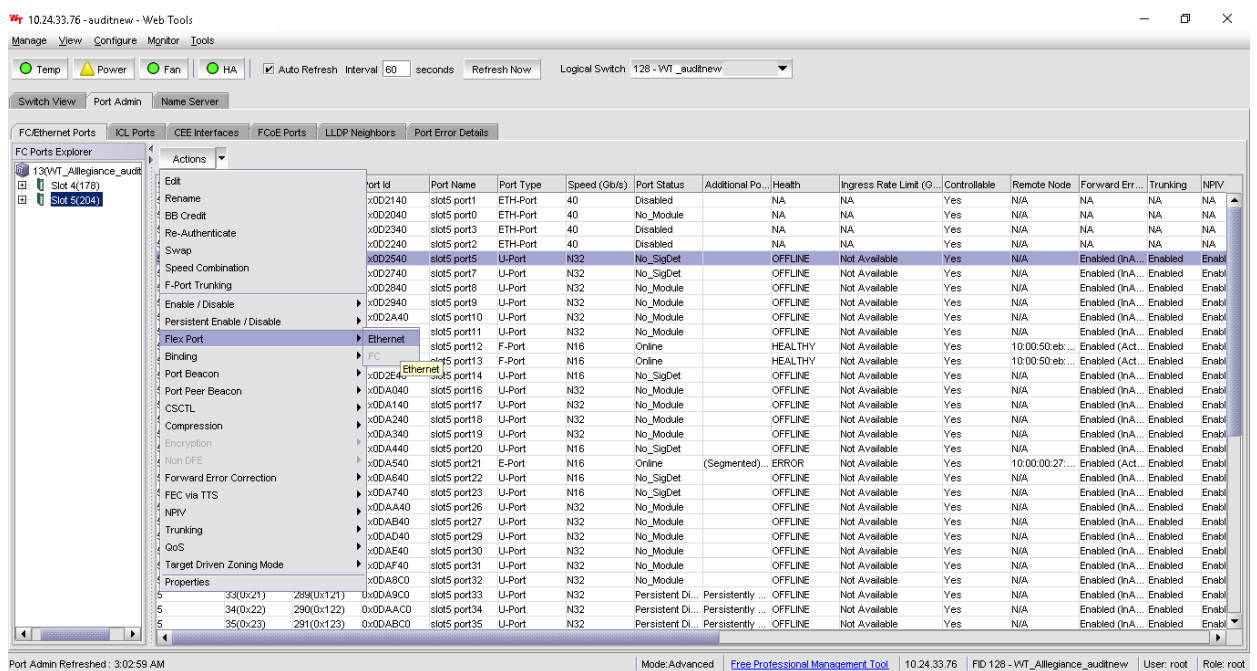
NOTE

In case of FCoE supported platforms, **FC ports** is renamed as **FC/Ethernet ports**. This also includes Flex-FC and Flex-Ethernet ports, as shown in the following figure. You are allowed to configure Flex ports to either FC or Ethernet.

NOTE

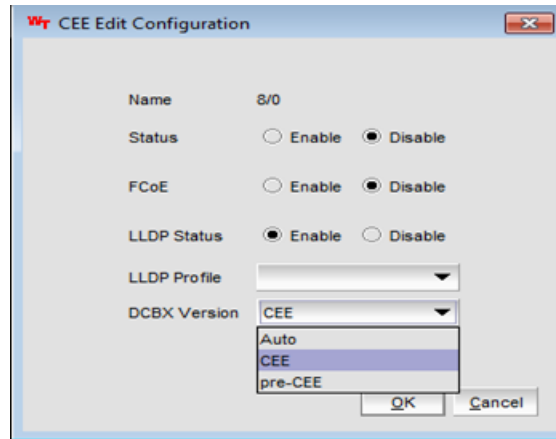
For Fabric OS 8.2.0, Web Tools does not support break-out configurations.

FIGURE 23 Flex port configuration



- The **CEE Interfaces** subtab is available only when Flex ports are available on the switch. This subtab allows the user to view and edit the CEE physical interface configuration, as shown in the following figure.

FIGURE 24 CEE Edit Configuration window



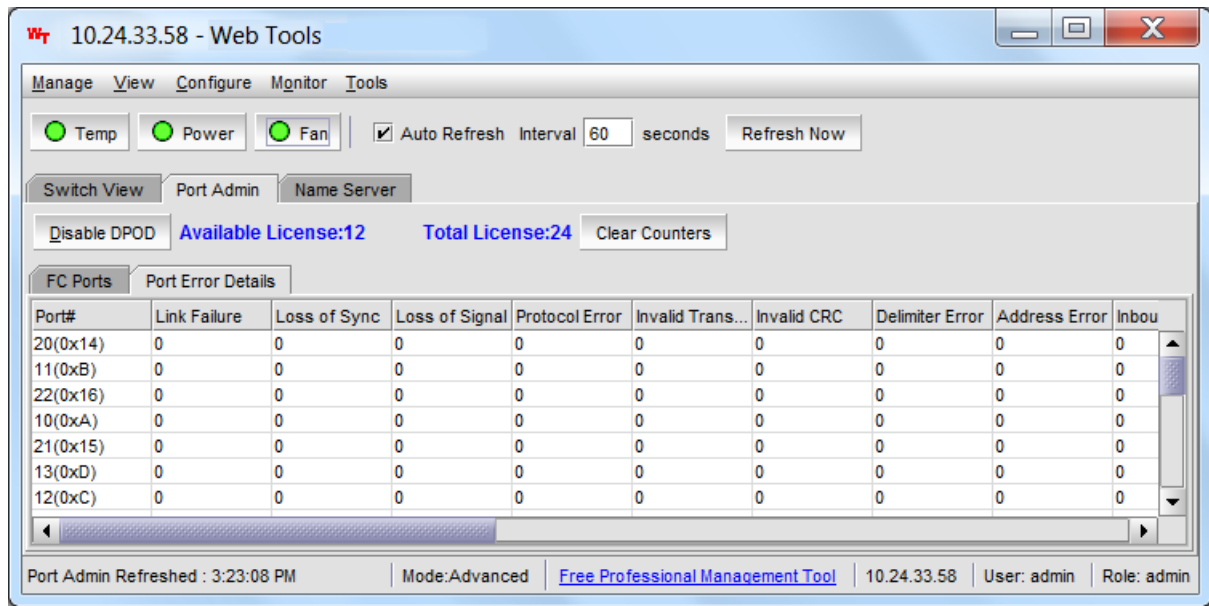
The following options can be enabled/disabled in the **CEE Edit Configuration** window:

- Status
- FCoE
- LLDP Status
- LLDP Profile
- DCBX Version

CEE interfaces has the following subtabs:

- General: General information about the CEE physical interfaces.
- FCoE Statistics
- LLDP Statistics
- Connected Devices
- **FCoE Ports** displays FCoE port information and allows the user to enable or disable FCoE ports.
- **LLDP neighbors** displays the remote port details for each local interface available on the switch.
- **Port Error Details** displays the detailed error statistics of all ports available on the switch.

FIGURE 25 Port Error Details subtab



- If the device has Gigabit Ethernet (GigE) ports, a "GigE Errors" subtab will also be displayed.
- The **VE/VEx Ports** subtab displays all VE_Ports and VEx_Ports on the switch. If the switch does not have VE_Ports and VEx_Ports, the **VE/VEx Ports** subtab does not display.
- The **ICL Ports** subtab displays all ICL ports on the switch. If the switch does not have ICL ports, the subtab does not display.
- The **GigE Ports** subtab displays all Gigabit Ethernet ports. If the switch does not have Gigabit Ethernet ports, the **GigE Ports** subtab does not display.

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

- General: General information about the Gigabit Ethernet ports.
- SFP: Displays information about SFP ports.
- Port Statistics: Displays statistics about the ports.
- IP Interfaces: Lets you view interfaces.
- IP Routes: Lets you view routes.
- Inband IP Interfaces: Lets you configure interfaces.
- Inband IP Routes: Lets you configure routes.
- FCIP Tunnels: Lets you view FCIP tunnels. This tab has two buttons: **Go to Extension 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
- GigE Port
- Source IP
- Destination IP
- Gateway

- VLAN ID
- MTU Size
- HA GigE Port
- HA Source IP
- HA Destination IP
- HA VLAN ID
- HA MTU Size
- Compression Mode
- Data L2COS Value
- DSCP Data
- IKE Policy Number
- IPsec Policy Enabled
- Keep Alive Timeout
- MaximumCommunicationRate (Mbps)
- MinimumCommunicationRate (Mbps)
- MaxRetransmitRate
- MinRetransmitRate
- Metric
- Pre-Shared Key
- QOS Mapping
- Selective Ack

Ports Explorer tree

The **Ports Explorer** tree displays on the left side of the window. Items in the tree are displayed as follows:

- Switches: Switch ID with the switch name in parentheses; for example, 3(MapsSW_202).
- Blades: Slot number of the blade with the blade ID in parentheses; for example, Slot 7(24).
- Ports: Port number; for example, Port 2.
- 10G SFP ports: A yellow triangle displays to visually distinguish the 10-Gbps SFP+ ports.
- QSFP Ports: Port number and QSFP number in parentheses; for example, (QSFP 0).

NOTE

The Brocade G610 Switch does not support QSFP ports.

Actions

The **Actions** list contains options for all tasks that you can perform on the selected ports. If you select more than one port, options are available for only those tasks that you can perform on all of the selected ports. Options are 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 ports for the slot or switch. If you select a port, the system displays detailed information about the port.

Subtabs

You can view either **Basic Mode** or **Advanced Mode**, and you can view the subtabs that contain additional information about the port. The available subtabs depend on the type of port selected.

To view basic mode, select **View** > **Basic**. When viewing detailed information about a port, **Basic Mode** provides the following subtabs:

- **General:** Under this tab, the **Actions** list provides the following options:
 - Edit
 - Rename
 - Enable/Disable
 - Persistent Enable/Persistent Disable
- **SFP:** Physical ports only (FC, CEE, and GbE)
 - Basic information about the port equipment
- **QSFP:** Quad Small Form-factor Pluggable ports
 - Basic information about the port
 - UnitNumber
 - ChannelIndex
 - DeviceTech

NOTE

From Fabric OS 8.2.0 and later, the SFP and QSFP tabs display current and voltage information along with other basic information.

- **Port Statistics:** All ports
 - Basic port information and statistics

Note that on the **Port Statistics** subtab, you can view 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

To view advanced mode, select **View** > **Advanced**. When viewing detailed information about a port, the **Advanced Mode** provides the following additional subtabs:

- **General:** Under this tab, the **Actions** list provides the following options:
 - Edit
 - Rename
 - BB Credit
 - Re-Authenticate
 - Swap
 - Reserve License
 - Release License
 - F-Port Trunking
 - Enable/Disable
 - Persistent Enable/Disable
 - Flex Ethernet/FC
 - Binding-Bind PID/Un-Bind PID
 - Port Peer Beacon--Port peer Beacon Enable/Port peer Beacon Disable

- CSCTL-Enable/Disable
- Beacon-Enable/Disable
- Compression-Enable/Disable
- Encryption-Enable/Disable
- Non DFE-Enable/Disable
- Forward Error Correction-Enable/Disable
- Transmitter Training Signal-Enable/Disable
- NPIV-Enable/Disable/Max Login
- Trunking-Enable/Disable
- QoS-Enable/Disable/Default (You can select more than one port to configure default QoS configuration)
- Target Driven Peer Zone-Enable/Disable
- Speed combination
- **SFP:** Physical ports only (FC, CEE, and GbE)
 - Basic information about the port
 - Advanced information about the port equipment
- **QSFP:** Quad Small Form-factor Pluggable ports
 - Basic information about the port
 - Advanced information about the port equipment
 - UnitNumber
 - ChannelIndex
 - DeviceTech
 - MaxCaseTemp
- **Port Statistics**
 - Advanced port statistics
 - Error details
 - FCIP Tunnels--GbE ports and logical extension ports only (not available for the FR4-16IP)

Controllable ports

All ports have a **Controllable** attribute visible from the **Advanced Mode**, which represents the RBAC permission.

The **Controllable** attribute is **No** when 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. Additionally, if you are logged in with read-only permission, the **Controllable** attribute displays **No** for all ports.

The **Controllable** attribute is **Yes** if your role gives you modify permission for ports. If a port is controllable, all configuration functionality is enabled.

Configuring FC ports


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

The following procedure describes how to open the **FC Port Configuration** wizard.

1. Click a port in the **Switch View** to open the **Port Admin** tab.
2. Select the **FC Ports** tab.
3. Select the port that you want to configure from the tree on the left.

4. Click the **General** subtab.
5. Select **Edit** from the **Actions** list.

FIGURE 26 Configuring FC ports



The **FC Port Configuration** wizard displays. The fields are populated with the current configuration values.

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 Admin** tab 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 with 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 E_Port, F_Port, or L_Port. For a logical FC port, the port can be 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 negotiates to its preferred type when the switch comes up, depending on what type of device or switch 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. An FC port cannot be configured as an E_Port or L_Port.

L_Ports are not supported on the Brocade FC16-32, Brocade FC16-48, Brocade FC16-64, Brocade 6505, Brocade 6510, and Brocade 6520.

NOTE

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

NOTE

VE_Port configuration is not supported in the Harpoon blade and Brocade 7840 Switch.

Speed

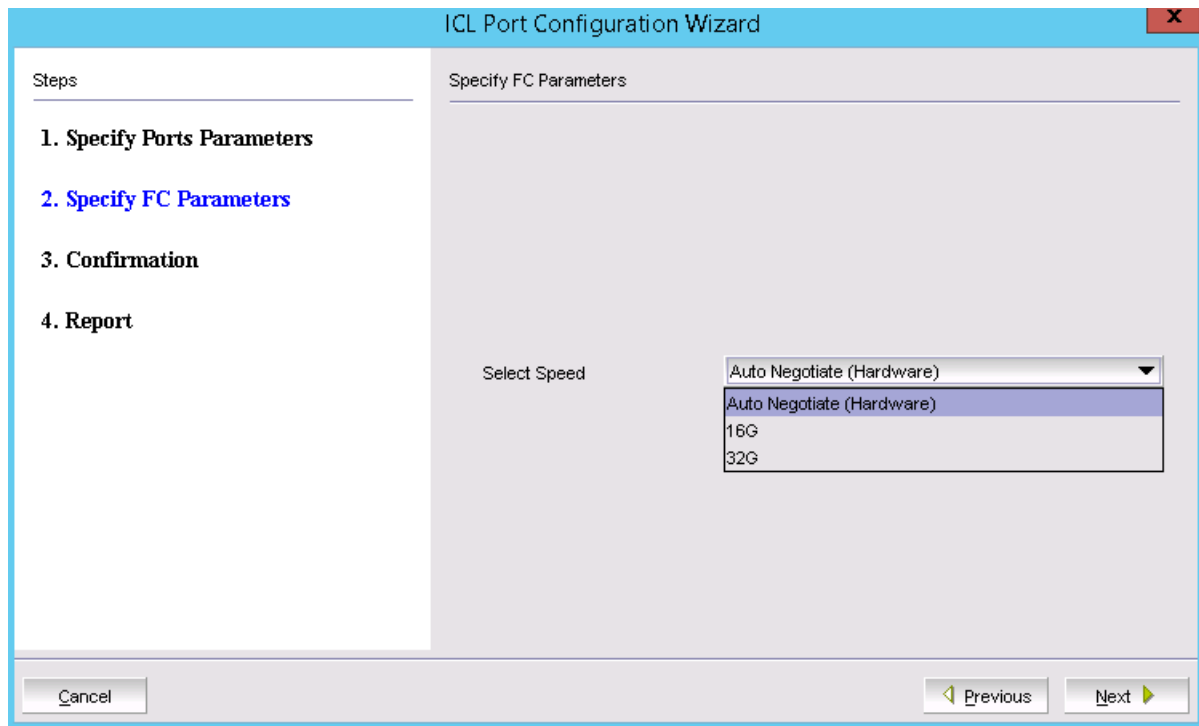
The **FC Port Configuration** tab provides the option to set the port speed. You can configure a 4G, 8G, 16G, or 32G port speed or set the port to auto-negotiate the highest possible port speed. The **Auto Max** options are displayed only when you set the port speed as auto-negotiation, and these options allow you to set the speed limit that the port can auto-negotiate. The following **Auto Max** speed levels are supported:

- Auto Max 4G
- Auto Max 8G
- Auto Max 16G
- Auto Max 32G

NOTE

From Fabric OS 8.2.0, and later the ICL port supports speed configuration of 32G, 16G, or Auto Negotiate, and you cannot configure 8G or any other port speed. In a ICL QSFP port, all four ports can have only a single speed.

FIGURE 27 ICL port speed configuration



When a Flex port is configured as Ethernet, the Flex-Ethernet port supports speed configuration. If the Flex-Ethernet port has the break-out option disabled, it supports only 40G speed configuration. When the Flex-Ethernet port has the break-out option enabled, it can support 10G and 25G speed configurations.

NOTE

Auto Max is not supported on QSFP and Flex-Ethernet ports.

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.

For information about long-distance mode settings, refer to [Administering Extended Fabrics](#) on page 175.

Available buffer credit calculation

The **FC Port Configuration** wizard provides non-editable **Recommended Buffer** and **Remaining Buffer** fields to check the available buffer credit for a port.

Recommended buffer

The number of recommended buffers. The recommended buffer is calculated based on the following values:

- Speed (not based on the auto-negotiate speed)
- Frame size
- Desired distance (km)

Remaining buffer

The number of remaining buffers. If the configured port buffer exceeds the remaining buffer, an error message displays.

Assigning a name to a port

Port names are optional. You can assign a name to an FC or Extension port to make port grouping easier. You can also rename FC and Extension ports. You cannot rename GbE ports. The **Port Name** column in the **General** tab displays the default port name.

Port names can be from 1 through 128 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.

NOTE

Duplicate port names are not allowed in FMS mode.

You can assign a Dynamic Port Name feature to display the switch name, port type, port index, and alias name as part of the port name for allowed port types. You can configure dynamic port name for E_Ports, F_Ports, EX_Ports, LE_Ports, and D_Ports. The default port name format for the dynamic port name is as follows:

- E_Ports: <switch name>.E_PORT. <Port index>
- F_Ports : <switch name>.F_PORT. <Port index><alias name>

If no alias is found for the F_Port, the "(null)" string displays in the <alias name> field. If any error occurs while finding the alias name, the "(none)" string displays in the <alias name> field.

To assign a name to a port, perform the following steps.

1. Click a port in the **Switch View** to open the **Port Admin** tab.
2. Select the **FC Ports** tab.
3. From the tree on the left, select the switch or slot that contains the port that you want to rename.
4. From the table, select a port or multiple ports that you want to rename.
5. Select **Rename** from the **Actions** list.
The **Rename** dialog box displays the selected port details.
6. Edit the names in the **Port Name** column and click **OK**.

Port beaconing

Individual FC ports can be set to beacon using the **Port Admin** tab. The port beaconing status displays in the **Port Beaconing** column. The **Switch View** reflects the port beaconing status by flashing the port amber and green in an alternating pattern for 2.5 seconds each.

To configure beaconing for an FC port, perform the following steps.

1. Select the **Port Admin** tab.
2. Click **View > Advanced**, if the **Port Admin** tab is in **Basic** mode.
3. Select the switch in the **FC Ports Explorer** list.
4. From the table, select a port or multiple ports that you want to set to beacon.
5. Select **Beacon > Enable** from the **Actions** list.

NOTE

You may select all ports on the switch, but if you select a port that is not valid for beaconing, the **Beacon** option is disabled.

While enabling port beacon, an error message displays for the following conditions:

- If switch beacon or chassis beacon is enabled on the switch.
- If the Port Peer Beacon feature is enabled on the port.

Port peer beaconing

The Port Peer Beacon feature in Web Tools allows you to physically identify the interconnections between FC ports.

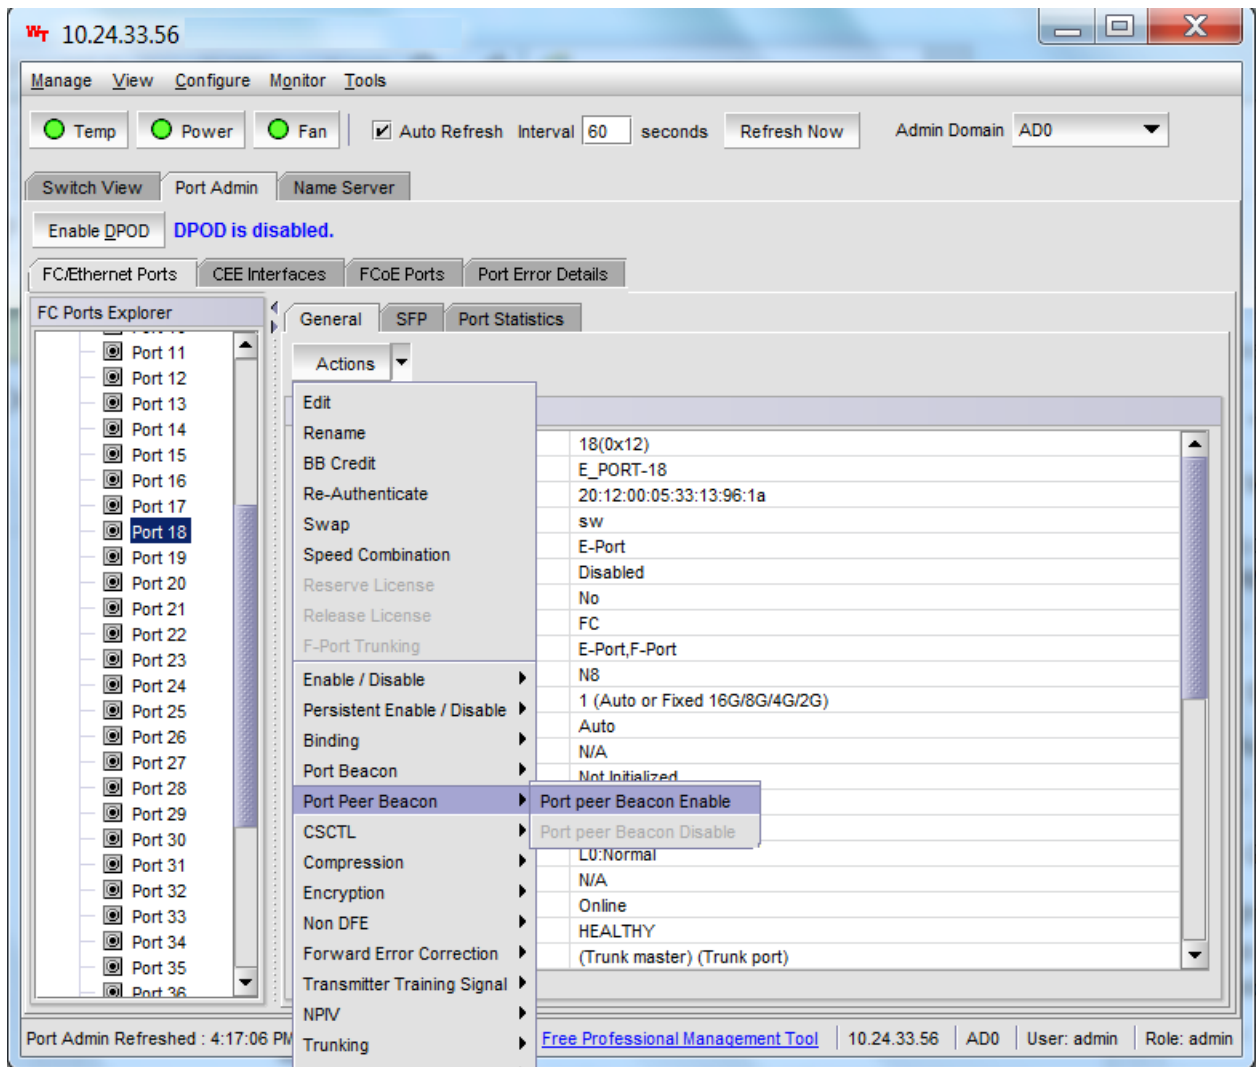
You can configure the Port Peer Beacon feature for a single port or for multiple ports. The Port Peer Beacon is supported on the E_Port, EX_Port, F_Port, N_Port, AE_Port, and Trunk ports. When you enable the Port Peer Beacon on any port that is part of a trunk group, it will enable port peer beaconing on all trunk ports in the same trunk. Port Peer Beacon configuration is supported in native switch mode and in Access Gateway mode.

To configure port peer beaconing for an FC port, perform the following steps.

1. Select the **Port Admin** tab.
2. Click **View > Advanced** if the **Port Admin** tab is in **Basic** mode.
3. Select the switch in the **FC Ports Explorer** list.
4. From the table, select a port or multiple ports that you want to set to beacon.

5. Select **Port Peer Beacon > Port peer Beacon Enable** from the **Actions** list.

FIGURE 28 Port peer beaconing

**NOTE**

You may select all ports on the switch, but if you select a port that is not valid for beaconing, an error message is displayed.

NOTE

You can configure the Port Peer Beacon only for a single AE_Port.

While enabling Port Peer Beacon, an error message displays if the switch beacon or chassis beacon is enabled on the switch.

Enabling and disabling a port

To enable or disable a port, perform the following steps.

1. Click a port in the **Switch View** to open the **Port Admin** tab.
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 that 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. Select either **Enable/Disable** > **Enable** or **Disable** from the **Actions** list.

NOTE

If the **Enable** or **Disable** option is unavailable, the port is already in the enabled or disabled state. For example, if the **Enable** option is unavailable, the port is already enabled. If you select multiple ports in both enabled and disabled states, both options are active. When you click either option, 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 Gigabit Ethernet ports to either copper or optical.
 - a) Select the **GigE Ports** tab.
 - b) Select the **GE0** or **GE1** port.
 - c) Select **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:

- If a port is not licensed, you cannot enable it until you install the appropriate license, such as a Ports on Demand, N_Port ID Virtualization license, or Q-Flex license (refer to [Port activation](#) on page 96 for more information). The **Licensed** field located in the **General** tab in the **Port Admin** tab 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 the *Brocade Fabric OS Administration Guide*.

Persistently enabling and disabling ports

To enable or disable a port so that it remains enabled or disabled across switch restarts, perform the following steps.

NOTE

Ports cannot be persistently enabled or disabled when FMS is enabled.

1. Select a port in the **Switch View** to open the **Port Admin** tab.
2. Select the **FC Ports**, **VE/VEx Ports**, **ICL 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.

- Select either **Persistent Enable/Disable** > **Enable** or **Disable** from the **Actions** list.

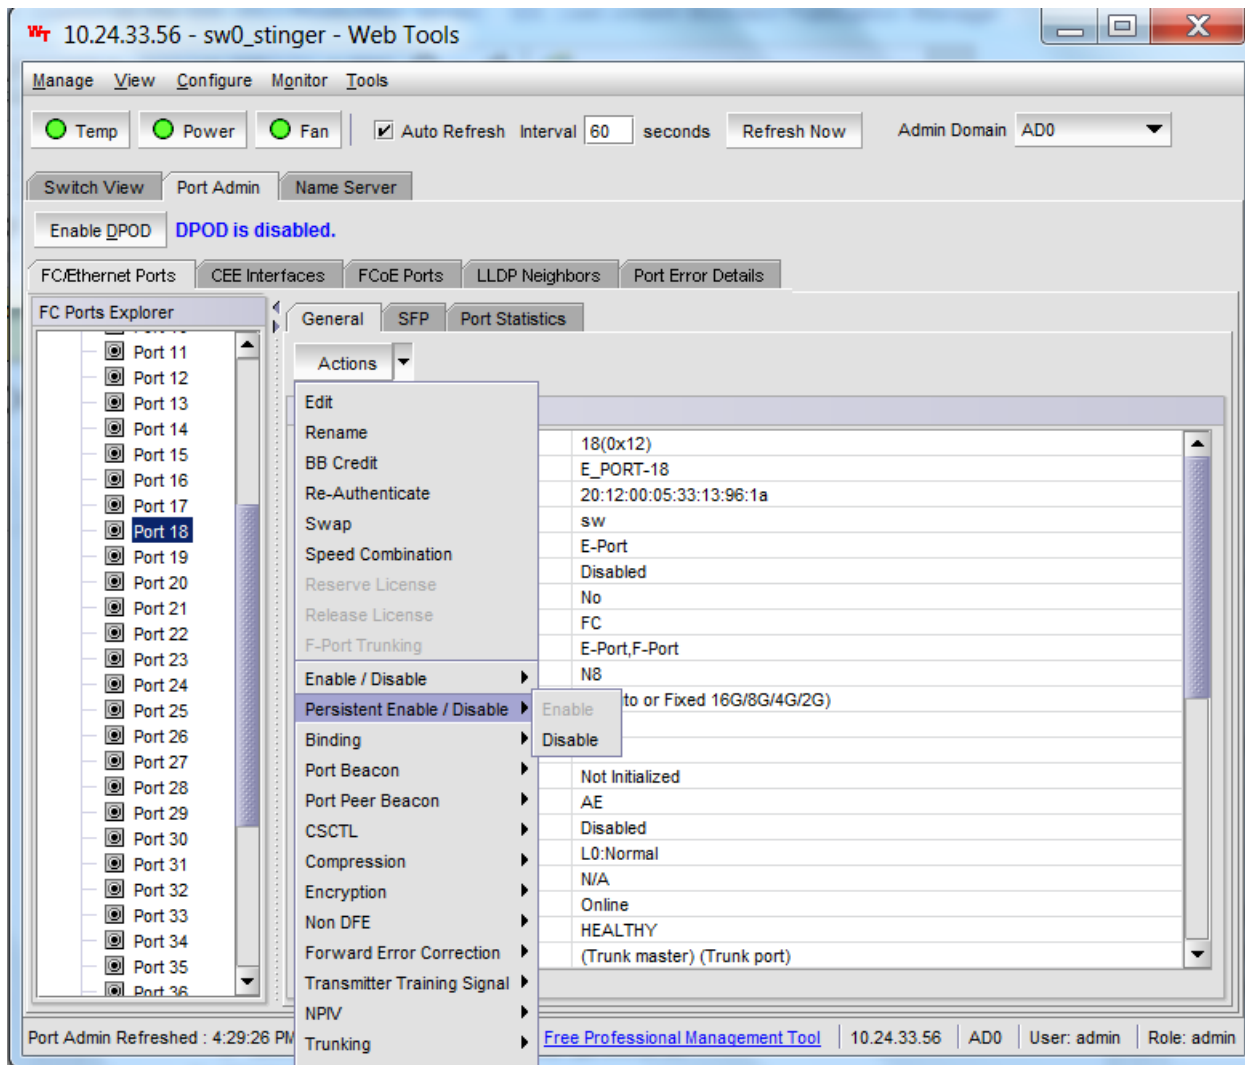
NOTE

Persistent enable or disable is not supported in FMS mode.

NOTE

If the **Enable** or **Disable** option is unavailable, the port is already in that state or FMS mode is enabled on the switch. For example, if the **Enable** option is unavailable, the port is already enabled. If you select multiple ports in both enabled and disabled states, both options are active. When you click either option, the action is applied to all selected ports.

FIGURE 29 Persistently enabling and disabling ports



6. Optional: If you are accessing a Brocade 7800 switch, you can set the media type for the GE0 and GE1 Gigabit Ethernet ports to copper or optical.
 - a) Select the **GigE Ports** tab.
 - b) Select the **GE0** or **GE1** port.
 - c) Select **Copper** or **Optical** from the **Media Type** selection list.
7. Click **Yes** in the confirmation window.

Configuring NPIV ports

For detailed information about understanding and configuring NPIV ports, refer to the *BrocadeFabric OS Administrator's Guide*.

NOTE

The NPIV feature cannot be disabled when Access Gateway mode is enabled.

The **NPIV Max Login Limit** option configures the maximum number of permitted logins per NPIV port. Each NPIV port can support up to 255 logins. The range of valid values is from 1 through 255 logins per port. The default value is 126 logins.

The NPIV feature supports virtual switches, but not on physical switches. Each port can have a different NPIV login limit value in each logical switch. The NPIV Max Login column displays the value assigned to each port.

To configure an NPIV port, perform the following steps.

1. Select the **Port Admin** tab.
2. Select **View > Advanced**.
3. Select the **FC Ports** tab.
4. From the tree on the left, select the switch or slot.
5. From the table, select one or more logical ports.
6. Select **NPIV > Max Login** from the **Actions** list.
The **NPIV Max Login** dialog box displays.
7. Enter the number of logins to be allowed in the **NPIV Max Login** field.
8. Select a port or ports under **Selected Ports**.
9. Click the right arrow to set the maximum login limit for the selected ports, and click **OK**.

Enabling Target Driven Zoning Mode

You can enable Target Driven Zoning Mode for target-driven peer-zone-enabled ports.

To enable Target Driven Zoning Mode on a port, perform the following steps.

1. Select the **Port Admin** tab.
2. Select **View > Advanced** if the **Port Admin** tab is in **Basic** mode.
3. Select a port or ports to configure.

4. Select **Target Driven Zoning Mode > Enable** from the **Actions** list.
The Target Driven Zoning Mode feature is enabled on the port.

NOTE

By default, Target Driven Zoning Mode is disabled on a port.

Port activation

Brocade switches come with a preset number of enabled ports. 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 paperpack 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.

The following table 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, as well as the ports that can be enabled with the Dynamic Ports on Demand feature.

TABLE 11 Ports enabled with POD licenses and the DPOD feature

Switch name	Enabled by default	Enabled with POD licenses	Enabled with the DPOD feature
Brocade 6505	0-11	12-23	
Brocade 6510	0-23	24-35, 36-47	
Brocade 6520	0-47	48-71, 72-95	

Brocade Gen 6 platforms support only DPOD, and therefore you can release and reserve a license using **Port Admin > Actions > Reserve License** or **Release License** in a single port or multiple ports.

In the **Port Admin** tab, the **Licensed** attribute for a port indicates whether a port is licensed (yes), whether it can be licensed (possible) because there are free licenses available (only applicable with the Dynamic Ports on Demand 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 92. Alternatively, you can disable and re-enable the switch to activate all ports as described in [Enabling and disabling a switch](#) on page 51.

To unlock a Ports on Demand license, you can use the supplied license key or you can generate a license key. For information on generating a license key, refer to the *Brocade Fabric OS Software Licensing Guide*.

Enabling Ports on Demand

To enable Ports on Demand, perform the following steps.

1. Install the Brocade Ports on Demand licensed product. For instructions, refer to [Enabling Ports on Demand](#).
2. Enable the ports as described in [Enabling and disabling a port](#) on page 92.

If you remove a Ports on Demand license, the licensed ports are disabled after the next platform restart or the next port deactivation.

Diagnostic ports

Diagnostic ports (D_Ports) are used for running diagnostics to isolate link-level faults and inter-switch link testing in fabric, optical, and remote loopback modes. A D_Port is not part of any fabric, and it does not carry any data or protocol traffic with it. It is used only for running diagnostic traffic for isolating link-level faults. The D_Port can be used to get estimated link distance measurements as done for long-distance mode links. For information on configuring a D_Port, refer to the *Brocade Fabric OS Administration Guide*. Web Tools cannot configure a D_Port.

The following list of features is not supported when a port is configured as a D_Port:

- Port swap
- Port bind
- Port trunk
- QoS enable/disable
- BB credit
- NPIV enable/disable/max login
- Allow/prohibit matrix

D_Ports do not take part in zoning. If a D_Port is added to a zone, it does not take part in the fabric.

Reserving and releasing licenses on a port basis

NOTE

The Dynamic Ports on Demand feature is supported on the Brocade X6-4 and Brocade X6-8 Directors and the Brocade 6505, Brocade 6510, Brocade 6520, Brocade G610, Brocade G620, and Brocade G630 Switches.

To reserve and release licenses on a port basis, perform the following steps.

1. Click a port in the **Switch View** to open the **Port Admin** tab.
2. Click the **FC Ports** tab.
3. From the tree on the left, click the switch or the slot that contains the port.

The **Licensed** column identifies the port license status:

- If the port has a license allocated, the **Licensed** 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 **Licensed** 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 **Licensed** field contains the value **Possible**.

You can reserve or release a license on any port with a license allocated. You must be logged in as Admin to reserve and release licenses.

NOTE

You must disable the port or switch before reserving or releasing a license.

To reserve a license, click **Reserve License** in the **Port Admin** tab.

To release a license, click **Release License** in the **Port Admin** tab.

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.

NOTE

Port swapping is not applicable to GE or ICL ports because there are no areas assigned to these ports.

The following restrictions apply to all ports:

- Ports can be swapped only once.
- A swapped port can only be un-swapped.
- Port binding is not supported on swapped ports.
- The swapped ports cannot be moved to other logical switches.
- The logical switch addressing mode cannot be changed if the logical switch contains the swapped 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 as shown in the following figure.

FIGURE 30 Port swapped label

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)

To swap ports, perform the following steps.

1. Select a port in the **Switch View** to open the **Port Admin** tab.
2. Select the **FC Ports** tab.
3. Select **View > Advanced**.
4. From the tree on the left, select the port that you want to swap.
5. Select **Enable/Disable > Disable** from the **Actions** list.

You must disable the ports used for port swapping. If the port is not in the disabled state, the port swap operation internally disables and re-enables the port.

6. Select **Swap** from the **Actions** list.

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 swapping on an FC16-48 is supported only on ports 0 through 15.

8. Click **OK**.

Determining if a port index was swapped with another switch port

To determine whether a port was swapped, perform the following steps.

1. Select a port in the **Switch View** to open the **Port Admin** tab.
2. Select the **FC Ports** tab.
3. Select **View > Advanced**.
4. From the tree on the left, select the port that you want to swap.

- 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 *Value* (Swapped), as shown in the following figure. The value indicates the port index with which the port was swapped.

FIGURE 31 Port swapping index

General SFP Port Statistics	
Actions ▼	
General	
Port Number	4(0x4)
Port Name	port4
Port WWN	20:04:00:05:33:e7:15:80
Port Media	--
Port Type	U-Port
Bound Status	No
Port protocol	FC
Allowed Port Type	E-Port,F-Port,L-Port
Speed (Gb/s)	N32
Speed Combination	1 (Auto or Fixed 32G/16G/8G/4G)
Speed Configured	Auto
Physical Port	N/A
Ingress Rate Limit (Gb/s)	Not Available
QoS Status	AE
CSCTL Mode	Disabled
Long Distance Mode	L0:Normal
Desired Distance (km)	N/A
Port Status	No_Module
Health	OFFLINE
Additional Port Info	
Controllable	Yes
Licensed	Yes
Port Index	5(0x5) (Swapped)
Trunking	Enabled

Configuring port binding

To bind a port or ports, perform the following steps.

- Select a port in the **Switch View** to open the **Port Admin** tab.
- Select the **FC Ports** tab.
- Select **View > Advanced**.
- From the table, select the port or ports that you want to bind.

5. Select **Binding** > **Bind PID** from the **Actions** list.

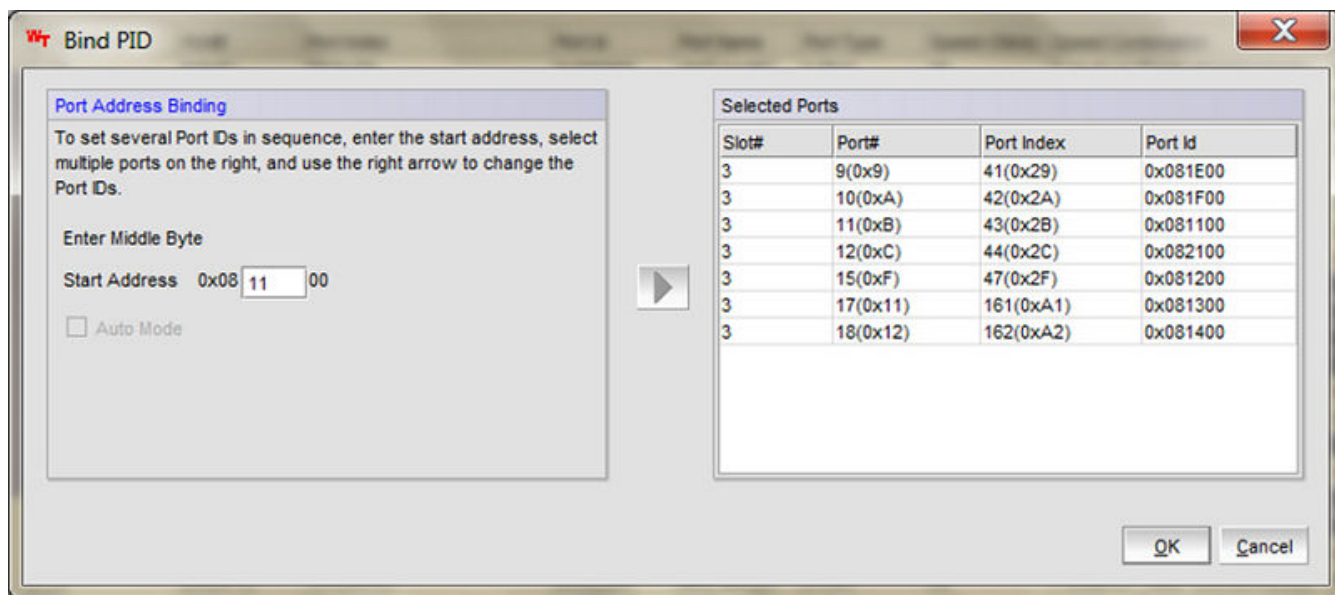
NOTE

If a port is already bound, a warning message is displayed that the port ID is already bound. The bound port is not listed in the **Bind PID** dialog box.

6. Enter the middle byte to be set in the **Start Address** field.
7. Select the ports under **Selected Ports** for which you want to set the middle bytes and click the right arrow as shown in the following figure.

For the selected ports, the middle bytes are assigned sequentially.

FIGURE 32 Port address binding



8. Click **OK**.

In the port list table for the bound ports, the Port ID column displays as Port ID *value* (Bound). You can also check the **Bound Status** attribute on the **General** tab to know if a port is bound or not.

Unbinding a port

To unbind a port or ports, perform the following steps.

1. Select a port in the **Switch View** to open the **Port Admin** tab.
2. Select the **FC Ports** tab.
3. Select **View** > **Advanced**.
4. From the **FC Ports** table, select the port or ports that you want to unbind.
5. Select **Binding** > **Un-Bind PID** from the **Actions** list.

For ports that are already bound, a warning message is displayed that the ports are already bound.

Configuring BB credits on an F_Port

To configure the BB credits value on an F_Port, perform the following steps.

1. Select a port in the **Switch View** to open the **Port Admin** tab.
2. Select the **FC Ports** tab.
3. Select **View > Advanced**.
4. From the tree on the left, select the switch or slot.
5. From the table, select one or more ports for which you want to set the BB credit.
6. Select **BB Credit** from the **Actions** list.

The **BB Credit** dialog box displays.

7. Enter the BB credit value in the **Enter BB Credit** field (the default value is 8).

NOTE

BB credit is not applicable for VE and ICL ports.

8. Select a port or ports under **Selected Ports**.
9. Click the right arrow to set the BB credit value for the selected ports, and click **OK**.

The value displays in the table of the **Port Admin** tab. If no value is configured, the **F-Port BB Credit** column displays the default value.

Configuring ALPA

The PID is the address assigned to the host when it performs a login with a fabric. The 24 bits of the PID are built from three 1-byte fields. The most significant byte is the Domain ID, the second byte is the area that device belongs to, and the least significant byte is the ALPA.

Persistent ALPA provides the hosts with the same ALPA that they received the first time they logged in. If they log in using the same port, the domain and the area for that device are still the same. This ensures that whenever a host logs in using the same port, it receives the same PID. The hosts can select their ALPA, and the switch provides the same value if it is available.

By default, persistent ALPA is disabled on Access Gateway switches. The Access Gateway always tries to request the same ALPA that the host has requested to the edge switch, but there is a possibility that the ALPA value has already been taken by another host. Therefore, the device either can use a different ALPA value (FLEXIBLE ALPA) that is available or can stick to the same requested ALPA value (STRINGENT ALPA). As the Access Gateway controls the assignment of ALPA values to the devices, it knows which ALPA value has been taken and which is free. With the FLEXIBLE ALPA option, the host login is accepted with either the requested ALPA value or a different ALPA value. With STRINGENT ALPA, if the requested ALPA value is not available, the login is rejected.

The Enable/Disable of Persistent ALPA feature is available on the **Switch** tab of the **Switch Admin** dialog box. The Persistent ALPA tables start populating as soon as the Access Gateway boots and the devices start logging in.

NOTE

Persistent ALPA is supported on all Access Gateway platforms. Persistent ALPA is not supported in non-Brocade fabrics.


To configure Persistent ALPA, perform the following steps.

1. Select the **Configure > Switch Admin > Switch** tab.

2. Select **Enable** for **Persistent ALPA Mode**.
This enables the **Stringent ALPA** and **Flexible ALPA** options.
3. Select **Stringent ALPA** or **Flexible ALPA**.
4. Click **Apply**.
5. Close the **Switch** page.
6. Select **Port Admin** tab.
7. Select an F_Port or U_Port from the device tree or Port List table.
8. Select **ALPA Map** from the **Actions** list.

The **Port WWN Map** dialog box launches listing the Port WWN to ALPA Map with the host. The Port WWN map automatically populates.

FIGURE 33 Port WWN Map dialog box



9. Optional: Click **Remove All** to clear all Port WWN maps.

Configuring the port octet speed combination

The **Port Admin** tab provides an option to set the port octet speed combination. This option is available only on the following platforms:

- Brocade DCX 8510-8 and Brocade DCX 8510-4, and Brocade X6-8 and Brocade X6-4 with the FC16-32, FC16-48, FC32-32, FC32-48, and FC32-64 port blades.
- Brocade 6510
- Brocade 6520
- Brocade G620
- Brocade G630

The ports on these hardware models are segregated into 8-port octets. The port octet speed combination is applied to the eight ports to which the selected port belongs. Based on this port octet speed combination, the speed options will be available in the **Port Configuration** wizard.

TABLE 12 Port octet speed combinations

Port octet in combination	Available port speeds within the octet
1	Auto or Fixed 32G 16G 8G 4G

TABLE 12 Port octet speed combinations (continued)

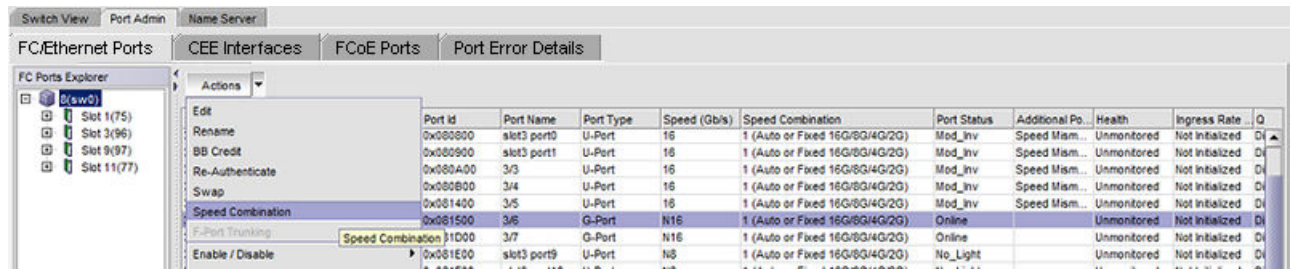
Port octet in combination	Available port speeds within the octet
2	Auto or Fixed 32G 10G 8G 4G

You can change the octet combination of a blade or switch. The octet speed must be set consistently across all members of the port octet.

To configure the port octet speed combination, perform the following steps.

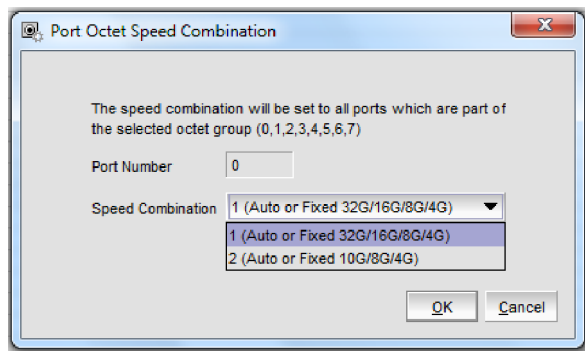
- 1. Select the **Port Admin** tab.
- 2. Select **View > Advanced mode**.
- 3. Select the **FC Ports** tab.
- 4. In the **FC Ports Explorer** dialog box, select a port to configure.

FIGURE 34 FC Ports Explorer dialog box



- 5. Select **Speed Combination** from the **Actions** list.
- The **Port Octet Speed Combination** dialog box displays.

FIGURE 35 Port Octet Speed Combination dialog box



- 6. Select a speed combination, and click **OK**.

Configuring CSCTL

Unlike the QoS zone-based FC flow prioritization method, CSCTL enables the same SID/DID pair exchange frames with different priorities.

To be able to prioritize a frame flow between two end nodes, Fabric OS 7.0.0 and later provide support for up to 32 virtual channels (VCs) per port. This categorizes the frames that enter a fabric on the basis of preset behavior defined with these VCs, and it conserves the frame's behavior until it is transmitted out of the fabric. However, of the 32 VCs for each external port, only 16 are used.

With the CSCTL method of prioritization, there is no need to have explicit traffic segregation, such as QOS_H, QOS_M, and QOS_L. The classification is entirely based upon the CSCTL database that is programmed into the ASIC. As the name suggests, CSCTL bits in each frame are used to define the VC number on the transmit port.

In order to achieve this kind of classification, Fabric OS 7.0.0 and later provide a CSCTL database table on each chip, which is capable of storing 256 entries. Each entry in the database table is populated with a VC number that, if this feature is enabled, is retrieved by indexing the CSCTL value into the table for each frame entering the fabric.

Irrespective of the type of frame classification method used, the flow priority of a frame is primarily determined by the VC number used to transmit the frames across the ISL ports. In both methods of classification, the VC number for a frame is determined at the ingress Fabric port (F_Port) or Fabric Loop port (FL_Port) when the frame enters the fabric for the very first time. To maintain the same flow priority for a frame across all ISL hops in a fabric, the same VC number is used while transmitting the frame at the egress E_Port until it leaves the fabric through an F_Port or FL_Port. The main difference between the QoS zone method of classification and the CSCTL VC-based method of classification is how the VC number is computed when the frame enters the fabric through an F_Port or FL_Port port and, of course, the manner of setting up these two frame classification methods.

Once CSCTL mode is enabled on an F_Port or FL_Port in a switch, the CSCTL value in the frame header of all incoming frames on that F_Port or FL_Port is used to index into the ASIC's CSCTL database table to compute the VC number. The VC number will define the frame's flow priority throughout its life in the fabric until it exits of the fabric through another F_Port or FL_Port. The QoS links (ISLs) preserve this classification during the frame's traversal across all hops in the fabric.

NOTE

When CSCTL mode and QoS zones are enabled, QoS zones lose priority to CSCTL mode.

NOTE

When the QoS zone is configured as default, CSCTL mode is disabled.

Enabling CSCTL mode

To enable CSCTL mode, perform the following steps.

1. Select the **Port Admin** tab.
2. Select the **View> Advanced** mode.
3. Select a port or ports to configure.
4. Select **CSCTL > Enable** from the **Actions** list.

Disabling CSCTL mode

To disable CSCTL mode, perform the following steps.

1. Select the **Port Admin** tab.
2. Select the **View> Advanced** mode.
3. Select a port or ports to configure.
4. Select **CSCTL > Disable** from the **Actions** list.

Configuring compression and encryption

Encryption provides security for the frames while they are in-flight, and compression allows better bandwidth utilization over long distance. For directors, the total bandwidth supported per blade for encryption is 32 Gbps and for compression is 64 Gbps. For a fixed-port switch, total bandwidth for encryption or compression is 32 Gbps. The Brocade G610 switch does not support in-flight compression and encryption. The Brocade G620 switch supports compression on an unlimited number of ports, but it does not support encryption.

NOTE

The compression and encryption feature is not supported in Access Gateway mode.

Enabling or disabling encryption

To configure encryption for an FC port, perform the following steps.

1. Click **Configure** > **Switch Admin**.
2. Select the **Security Policies** tab and then **Authentication** from the left panel.
3. Select **Active** or **On** from the **Switch Authentication Policy Mode** list.

NOTE

To enable or disable encryption on a port, the following criteria must be satisfied.

- The switch authentication policy should be active or on.
- **DH Group** must be set to **0, 1, 2, 3, 4** or **4 (2048 bit key)**.

4. Click **Apply** and **Close**.
5. Select the **Port Admin** tab from the **Switch Explorer** window.
6. Select **View** > **Advanced**.
7. Select a port from the **FC Ports Explorer**.
8. Select **Encryption** > **Enable** or **Disable** from the **Actions** list.
9. Click **Yes** on the confirmation dialog box to enable or disable encryption.

One of the following encryption statuses is displayed in the **General** tab of the port:

- **Enabled (Active)** - Encryption is enabled on a port, and the configuration is Active.
- **Enabled (Inactive)** - Encryption is enabled on a port, and the configuration is Inactive.
- **Disabled** - Encryption is not enabled on a port.

Enabling or disabling compression

NOTE

The compression feature is not supported on the Brocade G610 switch.

To configure compression for an FC port, perform the following steps.

1. Select the **Port Admin** tab from the **Switch Explorer** window.
2. Select **View** > **Advanced**.
3. Select a port from the **FC Ports Explorer**.

4. Select **Compression** > **Enable** or **Disable** from the **Actions** list.
5. Click **Yes** on the confirmation dialog box to enable or disable compression.

One of the following compression statuses is displayed in the **General** tab of the port:

- **Enabled (Active)** - Compression is enabled on a port, and the configuration is Active.
- **Enabled (Inactive)** - Compression is enabled on a port, and the configuration is Inactive.
- **Disabled** - Compression is not enabled on a port.

Displaying the compression ratio

32-Gbps-capable FC platforms that support compression and 16-Gbps-capable FC platforms that support compression and encryption also provide the compression ratio. Under the **Port Admin** tab, the **FC Ports** tab, and the **General** tab of a port, displays one of the following values for the **Compression Ratio**.

TABLE 13 Compression ratio

Value	Description
<numeric value>	Compression is enabled for a port.
- -	Compression is supported but not configured.
N/A	Compression is not supported.

Forward Error Correction

Forward Error Correction (FEC) allows the recovery of error bits in a 10-Gbps, 16-Gbps, or 32-Gbps data stream. This feature is enabled by default on all ISLs and ICLs of 32-Gbps FC platforms. FEC is supported in Access Gateway mode.

To configure FEC for an FC or ICL port, perform the following steps.

1. Select the **Port Admin** tab from the **Switch Explorer** window.
2. Select **View** > **Advanced**.
3. Select a port from the **FC Ports Explorer** or **ICL Ports Explorer**.
4. Select **Forward Error Correction** > **Enable** or **Disable** from the **Actions** list.

One of the following FEC statuses is displayed in the **General** tab of the FC port or ICL port:

- **Enabled (Active)** - FEC is enabled on a port, and the configuration is Active.
 - **Enabled (Inactive)** - FEC is enabled on a port, and the configuration is Inactive.
 - **Disabled** - FEC is not enabled on a port.
 - **NA** - FEC is not supported.
5. Click **OK** on the warning message that prompts for enabling or disabling of FEC results in traffic disruption.

In-band management

In-band management is designed to allow the management of the switch through Gigabit Ethernet ports. This allows a management station located on the WAN side of an Extension platform to communicate with the control processor for management tasks, such as launching Web Tools, SNMP polling, SNMP traps, troubleshooting, and configuration. To provide this communication, new interfaces have been added to the control processor that have an external IP address, allowing IP connectivity through the port processor to the control processor.

The in-band management interface is protocol-independent, so any traffic destined for these in-band management interfaces is passed through the distribution point to the control processor. It is then handled on the control processor according to the rules set forth for the normal management interface and following any security rules that may be in place on the control processor.

To provide redundancy, there is one in-band management interface per Gigabit Ethernet port. This allows the management station on the WAN side of the network to have multiple addresses with which to reach that switch, and it allows redundancy if one of the Gigabit Ethernet ports becomes unreachable for any reason.

Communication is handled through external addresses that are configured independently for each in-band management interface. The in-band management interfaces share the routing table on the control processor. This is separate from the routing table for each Gigabit Ethernet port that exists. So, there are certain limits to the addresses that are allowed and to the routes that are allowed for the in-band management interfaces and route entries.

In-band management is supported on the Brocade FX8-24. In-band management is not supported in Fabric OS 7.3.0 or earlier. Only one IP interface entry can be configured per Gigabit Ethernet port.

To configure in-band management, perform the following steps.

1. Select **Port Admin > GigE Ports > In-band IP Interface**.
2. Click **Add** to configure a new in-band management entry.
3. Set the **IP Address Type** to **IPv4**.
4. Set the address options:
 - IP address
 - Subnet mask
 - MTU size
5. Click **OK**.
6. Select the **Inband IP Routes** tab.
7. Click **Add** to configure a new route entry.
8. Set the **IP Address Type** to either **IPv4** or **IPv6**.
9. Set the address options of the management station on the WAN side of the Extension platform:
 - Destination IP address
 - Subnet mask
 - Gateway IP address
10. Click **OK**.
11. Select the **General** subtab.
12. Select the **Enable** option from the **Inband** selection list to activate in-band management.

GigE port modes

Web Tools allows you to set the GigE port mode for the FX8-24 extension blades to 1-Gbps, 10-Gbps, or dual modes.

To configure the GigE port mode, perform the following steps.

NOTE

You must install the FX8-24 extension blade in a slot that contains a 10-GE license to configure the mode.

1. Select **Port Admin** > **GigE Ports**.
2. Select **View** > **Advanced**.
3. Select a slot from the **GigE Ports Explorer** panel.
4. Select one of the following modes from **Mode** under the **Actions** list.
 - **1G** - To enable the ge0 through ge9 ports.
 - **10G** - To enable the xge0 and xge1 ports.
 - **Dual** - To enable the ge0 through ge09 and xge0 ports.
5. Click **Yes** on the confirmation dialog box.

Enabling ISL Trunking

• ISL trunking overview.....	111
• Disabling or enabling ISL Trunking	111
• Viewing trunk group information.....	111
• F_Port trunk groups.....	112

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 60. For additional information about ISL trunking, refer to the *Brocade Fabric OS Administration Guide*.

For detailed information about ISL trunking configurations and criteria, refer to the *Brocade Fabric OS Administration Guide*.

Disabling or enabling ISL Trunking

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.

To disable trunking on a port, or to re-enable trunking if it has been disabled, perform the following steps.

1. Select a port in the **Switch View** to open the **Port Admin** tab.
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. Select **View** > **Advanced**.
6. Select **Trunking** > **Enable** or **Disable** from the **Actions** list.

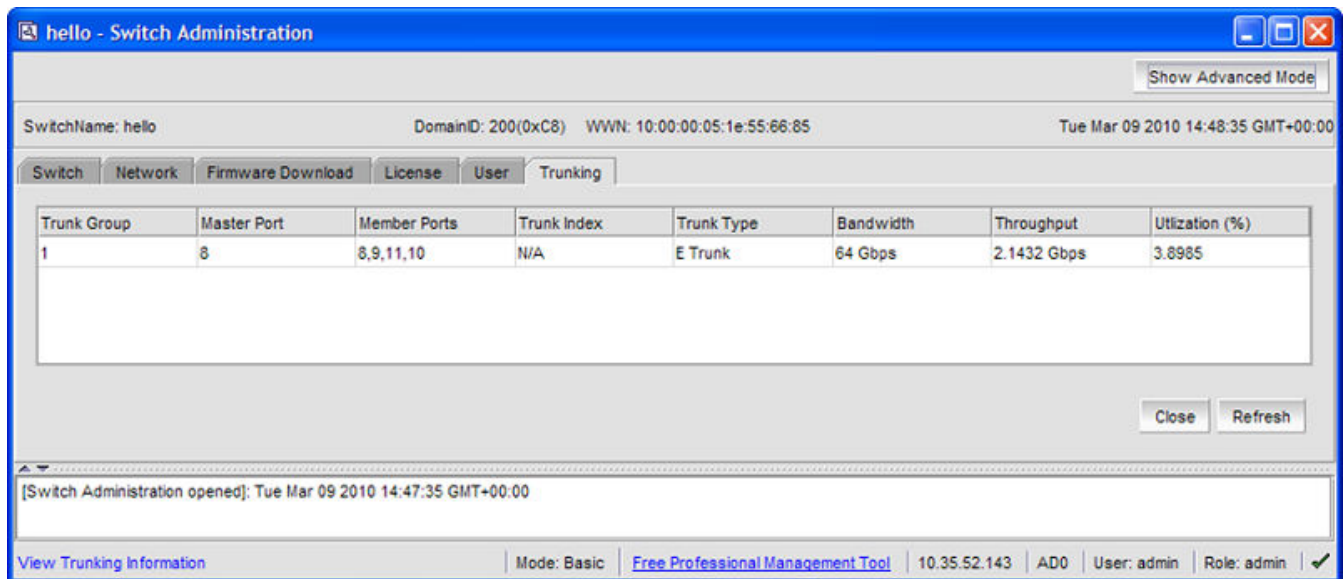
If the option is unavailable, the selected port is already in that state.

7. Click **Yes** in the confirmation dialog box.

Viewing trunk group information

Use the **Trunking** tab on the **Switch Administration** window to view trunk group information.

FIGURE 36 Trunking tab



The following trunking attributes can be displayed from the **Switch Administration** 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 for E_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** tab in **Advanced** mode:

- 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 four general criteria for establishing F_Port trunking:

- 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.

Creating and maintaining F_Port trunk groups

On the Brocade G620 Switch, the trunk groups are in the port ranges 0-7, 8-15, 16-23, 24-31, 32-39, 40-47, 48-55, and 56-63, which are applicable on front-end ports.

Use this procedure to create an F_Port trunk group and to add or remove member ports.

1. Select the **Port Admin** tab.
2. Select **View > Advanced**.
3. Select any port from the port group in which you want to create the trunk group.
4. Select **F-Port Trunking** from the **Actions** list.

The **F-Port Trunking** dialog box displays.

5. Select one or more ports in the **Ports for trunking** pane.
6. Click **Create Trunk Group**.

The **Select Trunk Index** dialog box displays asking you to select a trunk index.

7. Select the trunk index from the list populated with the index for all ports.

A trunk group is created, identified by the trunk index, and containing the ports that you selected.

8. Select the trunk group that you just created.
9. Additional ports can be added by selecting a port from the **Ports for trunking** table and then clicking **Add Members**.

NOTE

To remove a port from the trunk group, select the port from the **Trunk Groups** table and then click **Remove Members**.

10. Click **OK** to save your changes.


CEE switch level configuration

- CEE tab configuration overview..... 115
- Link Aggregation.....115
- FCoE Management..... 117
- LLDP..... 118

CEE tab configuration overview

From 8.2.0 and later the CEE related switch level configuration will be positioned under the **Switch Admin** panel. This tab will have the provision to configure Link Aggregation (LAGs), FCoE Management and LLDP. The CEE tab is available in both basic mode and advanced mode on FCOE capable platforms.

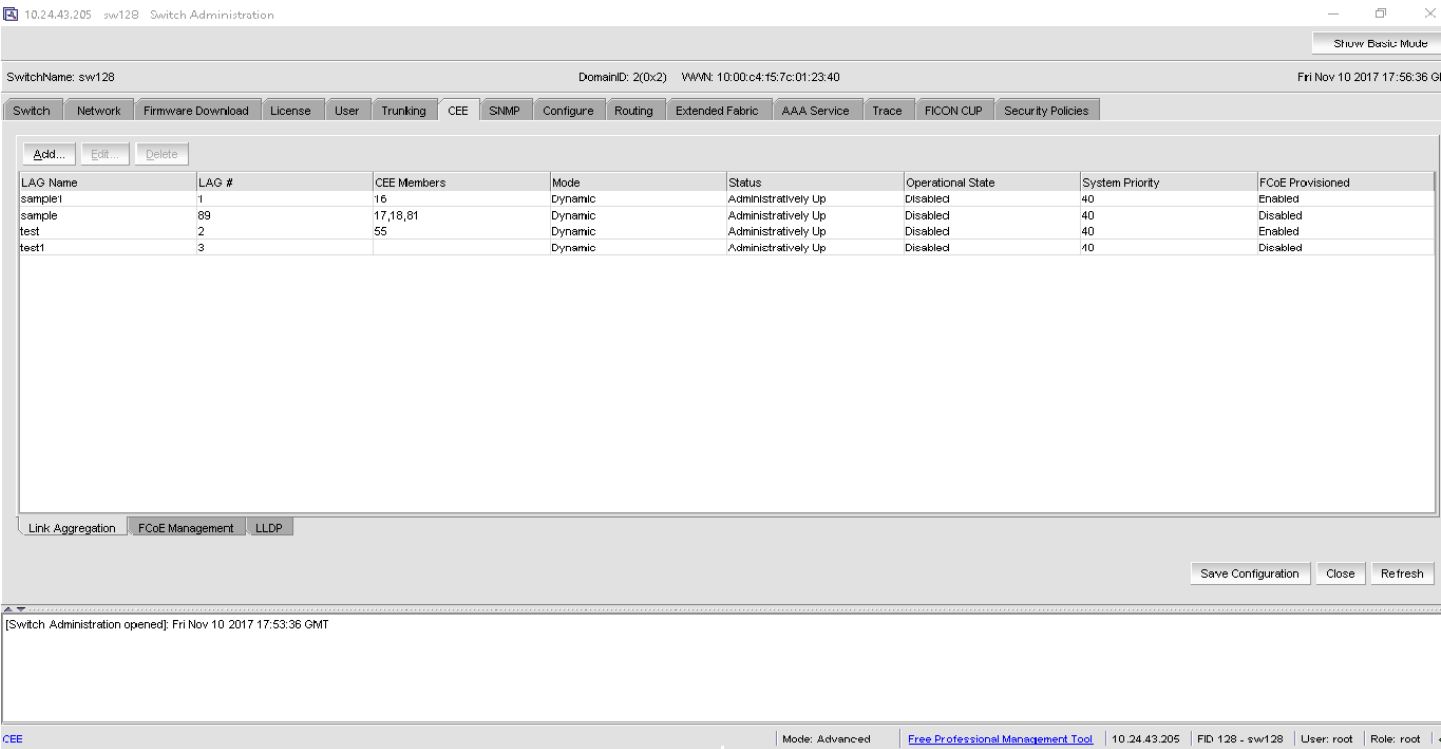
FIGURE 37 CEE tab



Link Aggregation

Web Tools will display ' Link Aggregation' sub-tab under CEE tab, where you can view and configure port channels and their properties. You can add, edit, and delete Link Aggregation settings.


FIGURE 38 Link Aggregation



Adding Port Channel

You are allowed to add port channels. The interface speeds will also be shown along with the interfaces in the list. If you are configuring a Brocade X6-8 and Brocade X6-4 Director, without a Brocade FC32-64 port blade, you can only create empty LAGs, because no Ethernet ports are available.

FIGURE 39 Add LAG Configuration



Editing Port Channel

You can also edit the LAG configuration. After editing you can save the configuration.

NOTE

You can add or remove only one interface at a time.

FIGURE 40 Edit LAG Configuration

Edit LAG Configuration

LAG Name:

LAG #:

Interfaces

Selection List: ☐ Interface

Selected List: 8/11, 8/14

Only one Interface can be Added or Deleted at a time

System Priority:

LAG Status: ☐ Enable ☒ Disable

FCoE Provisioned: ☒ Enable ☐ Disable

FCoE Management

Web Tools will display 'FCoE Management' sub-tab under CEE tab, which includes two sections FCoE Fabric Map and FCoE Enodes.

You can view and edit the FCoE Fabric Map settings in FCoE Management.

The following options can be edited in the FCoE Fabric Map settings:

- Name
- VLAN
- Priority Bits
- FC Map
- FKA
- Timeout

You are also allowed to view the Domain ID and view/edit the number of FCoE Enodes as shown in following figure.

FIGURE 41 FCoE Management

The screenshot displays the 'FCoE Management' configuration page. At the top, there is a navigation bar with tabs: Switch, Network, Firmware Download, License, User, Trunking, CEE, SNMP, Configure, Routing, Extended Fabric, AAA Service, Trace, FICON CUP, and Security Policies. The 'Configure' tab is active.

The main configuration area is divided into two sections:

- FCoE Fabric Map:**
 - VLAN (2-4093): 1002
 - VFID: 128
 - Priority Bits (0-6): 3
 - FC MAP (0x0efc00-0x0efcff): 0x0efc00
 - FKA (250-90000): 8000
 - Timeout: ☒ Enable ☐ Disable
- FCoE Enodes:**
 - Domain ID: 2
 - MAC Address: c4:f5:7c:00:d1:d2
 - FCoE Enodes (0-1600): 20

At the bottom, there are three sub-tabs: Link Aggregation, FCoE Management (selected), and LLDP. On the right side, there are three buttons: Save Configuration, Close, and Refresh.

LLDP

Web Tools provisions the LLDP sub-tab under CEE tab in **Switch Admin** panel.

The two levels of LLDP configuration are:

- Global configuration
- LLDP Profile

Global configuration: The configuration can be performed at the global level and applied to every CEE interface in the system.

FIGURE 42 LLDP- Global Configuration

The screenshot shows the 'LLDP' configuration page under the 'Global' tab. The left sidebar has 'Global' and 'LLDP Profile' tabs. The main area contains the following configuration fields:

- LLDP**: ☒
- System Name**: test
- System Description**: a_b
- Mode**: Both (dropdown)
- Transmit Interval (4-180)**: 30
- Multiplier (2-10)**: 4
- Hold Time for Advertise**: 120
- Advertise Optional-tlv**:
 - ☒ Chassis ID TLV
 - ☒ Port ID TLV
 - ☒ TTL TLV
 - ☒ port-description
 - ☐ system-name
 - ☒ system-description
 - ☒ management-address
 - ☒ Advertise dot1-tlv
 - ☒ Advertise dcbx-fcoe-logical-link
 - ☒ Advertise dot3-tlv
 - ☒ Advertise dcbx-tlv

At the bottom, there are tabs for 'Link Aggregation', 'FCoE Management', and 'LLDP'.

LLDP Profile: You can create customize the LLDP profiles locally and apply them over specific DCB interfaces of your choice to override the Global configurations. You can also view all the LLDP profiles created locally.

FIGURE 43 LLDP Profile


The screenshot shows the 'LLDP Profile' configuration page. At the top, there are 'Add...', 'Edit...', and 'Delete' buttons. Below them is a table with the following columns: Name, Mode, Multiplier, Transmit Interval, Hold Time for Advertise, port-descript..., system-name, system-cap..., system-des..., management..., and Dot1-tlv. The table contains three rows of data:

Name	Mode	Multiplier	Transmit Interval	Hold Time for Advertise	port-descript...	system-name	system-cap...	system-des...	management...	Dot1-tlv
sample-	Both				Not Advertis...	Not Advertis...	Not Advertis...	Not Advertis...	Not Advertis...	Not Advertis...
bew-	Both				Not Advertis...	Not Advertis...	Not Advertis...	Not Advertis...	Not Advertis...	Not Advertis...
1111_	Both				Not Advertis...	Not Advertis...	Not Advertis...	Not Advertis...	Not Advertis...	Not Advertis...

At the bottom, there are tabs for 'Link Aggregation', 'FCoE Management', and 'LLDP'.

You can add the LLDP configuration as shown in the following figure.

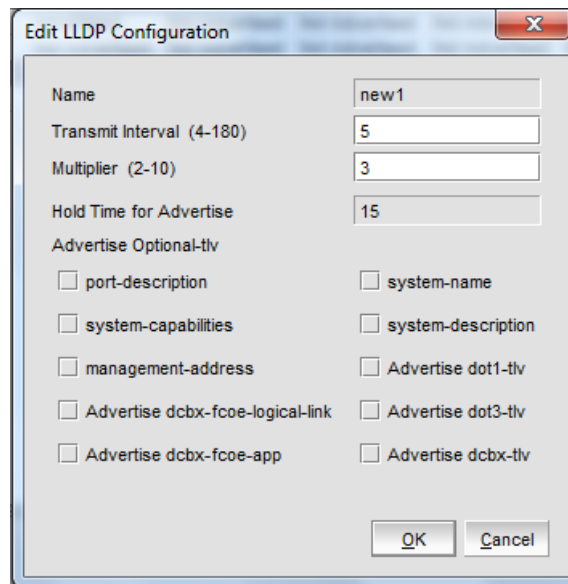
FIGURE 44 Add LLDP Configuration



The following options can be edited in the Edit LLDP Configuration window as shown in the following figure:

- Transmit Interval
- Multiplier

FIGURE 45 Edit LLDP Configuration



Monitoring Performance

• Performance Monitor overview.....	121
• Opening the Performance Monitor window.....	123
• Creating basic performance monitor graphs.....	123
• Customizing basic monitoring graphs.....	123
• Tunnel and TCP performance monitoring graphs.....	126
• Printing graphs.....	127

Performance Monitor overview

The Web Tools Performance Monitoring tool graphically displays throughput (in megabytes per second) for each port and for the entire switch.

Basic monitoring

The **Basic Monitoring** menu is standard in the Web Tools software. Any user logged in to Web Tools with a role of zoneadmin or securityadmin cannot open **Performance Monitor**. The roles user, operator, basicswitchadmin, and properly configured user-defined roles are allowed to perform basic monitoring tasks. 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).

NOTE

Beginning with Fabric OS 7.4.0, canvas display and save operations are not supported.

Performance graphs

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

Graphs within the **Performance Monitor** 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 Monitor** window, you can perform the following tasks:

- Select **Window** > **Tile** to view all graphs at once, tiled in the **Performance Monitor** 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 Monitor** 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 under the **Performance Graphs** window displays real-time performance monitoring charts for the Brocade 7840 Extension Switch, Brocade SX6 Extension Blade, and Brocade FX8-24 Extension Blade. This option is not available on other platforms.

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.

You can access the basic monitoring graphs on all switches. The following table lists the basic monitoring graphs available.

TABLE 14 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 for the Brocade DCX 8510 and Brocade X6 enterprise-class platforms.
Switch Throughput Utilization	The port throughput, in Gbps at the time the sample is taken. For the Brocade DCX 8510 and X6 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 DCX 8510 and X6 enterprise-class platforms, 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 ports on a switch. For the Brocade DCX 8510 and Brocade X6 enterprise-class platforms, this graph displays the CRC error rate for each slot. You can customize this graph to display information for particular ports.

The following table lists each graph and indicates the supported port types for each graph. The port selection columns for each graph display the supported ports.

TABLE 15 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	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

The labeling of the axes in the graphs depends on the switch type:

- For the Brocade G620, the X-axis scales up to 32.0 Gbps in increments of 0.32 Gbps.
- For the Brocade DCX 8510-4, Brocade DCX 8510-8, and the Brocade X6-4 and Brocade X6-8 enterprise-class platforms, slot numbers are displayed with expansion arrows next to them. Click the arrows to expand and contract the list of ports per slot.
- Switches such as the Brocade 7840 Extension Switch and the Brocade G620 do not have slot numbers because they have no blade FRUs, and therefore there is no need for slot numbering.
- For the Brocade DCX 8510-4, Brocade DCX 8510-8 and the Brocade X6-4 and Brocade X6-8 enterprise-class platforms, the X-axis scales up to 409.6 Gbps in multiples of 2.
- For the Brocade 6505, Brocade 6510, and Brocade 6520, the X-axis scales up to 16.0 Gbps in increments of 0.16 Gbps.
- For the Brocade 7840 Extension Switch, the X-axis scales up to 40.0 Gbps in increments of 0.4 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.

Opening the Performance Monitor window

To open the **Performance Monitor** window, perform the following steps.

1. Select a switch from the **Fabric Tree** and log in when prompted.
2. Click **Monitor** > **Performance Monitor**.

The **Performance Monitor** window displays.

Creating basic performance monitor graphs

To create the basic performance monitor graphs listed in [Predefined performance graphs](#) on page 122, perform the following steps.

1. Open the **Performance Monitor** window.
2. Select **Performance Graphs** > **Basic Monitoring** > **Graph Type**.

Depending on the type of graph that 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 the Brocade 7840 Extension Switch, enter only a port number.

4. Click **OK**.

The graph is displayed in a window in the **Performance Monitor** window.

Customizing basic monitoring graphs

You can customize some of the basic performance monitoring graphs to display information for particular ports. For the Brocade DCX 8510-4, Brocade DCX 8510-8, and Brocade X6-4 and Brocade X6-8 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 that you want to customize.

Refer to [Creating basic performance monitor graphs](#) on page 123 for instructions on creating a graph.

2. For the Brocade 6505, Brocade 6510, Brocade 6520, Brocade G610, Brocade G620, Brocade G630, and the Brocade 7840 Extension Switch, 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 the Brocade 7840 Extension Switch, proceed to Step 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 the following figure.

The title of the dialog box varies, depending on the type of graph that you are customizing, but the layout of the dialog box is the same. The following figure displays an example of the setup dialog box for the **Edit Switch Throughput Utilization** graph.

FIGURE 46 Select Ports for customizing the Switch Throughput Utilization graph



You can perform the following in the dialog box:

- Double-click the domain to expand the slot or port list.

NOTE

For the Brocade DCX 8510-4, Brocade DCX 8510-8, Brocade X6-4, and Brocade X6-8 enterprise-class platforms, click the plus signs (+) to expand the ports under each slot, as shown in the previous figure.

- b) Click the port that 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 that 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.

NOTE

The performance monitoring graphs does not display Ethernet port statistics in the graph.

Tunnel and TCP performance monitoring graphs

This section describes how to generate the Tunnel and TCP performance monitor graphs. You can launch a 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 available TCP graphs are:

- Sender RoundTrip
- Sender RoundTripVariance
- TCP DupAck
- TCP OOS
- TCP SlowStart
- TCP FastRetransmit
- TCP Tx(MB/sec)
- TCP Rx(MB/sec)

The available Tunnel graphs are:

- Throughput(MB/sec)
- Effective Throughput(MB/sec)
- CompressionRatio

For TCP connection graphs, tool tips are displayed only for all selected connections.

To create a Tunnel and TCP graph, perform the following steps.

1. Select **Monitor > Performance Monitor**.
The **Performance Monitor** 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** list for which you want to generate the graphs.

For the Brocade 7840 Extension Switch, you can have a maximum of four circuit connections in a tunnel, and for the BrocadeFX8-24 Extension Blade, you can have a maximum of ten circuit connections in a tunnel.

4. In the **Tunnel and TCP** area at the bottom of the window, select the required check boxes for the statistic that 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 from 3 through 30 seconds. The X-axis is limited to 30-minute. The graph scale starts with 0 minutes and auto-scales to draw the statistics. Once the 30-minute graph is drawn, the first minute data is removed to accommodate the 31st-minute values.
 - **Global auto scaling:** By default, this option is disabled. You can either enable or disable this option. If enabled, the graph's X-axis scales up to 30-minute, and if it is disabled, the X-axis will scale up to 10-minute.
 - **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.

NOTE

Brocade Network Advisor has an option for launching the **Tunnel and TCP Graph** dialog box from the **FCIP Tunnels** dialog box.

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 properties are not persistent. When you close the graph, the values reset to the default settings.

In addition, you can print the graph and save the graph to a file.

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.

To print a graph, perform the following steps.

1. Open the **Performance Monitor** window.
2. Create a basic or advanced Performance Monitor graph as described in [Creating basic performance monitor graphs](#) on page 123.
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**.

Administering Zoning

• Zoning overview.....	129
• Zoning configurations	131
• Zoning management.....	132
• Zone configuration and zoning database management.....	141
• Best practices for zoning.....	149

Zoning overview

This chapter describes zoning and provides the procedures for managing zoning. The **Zone Admin** window provides two zoning options on the left pane:

- Basic zones
- Traffic Isolation zones

You can perform basic zoning and Traffic Isolation zoning using Web Tools.

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 the 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.

Peer zones

Peer zoning introduces the concept of principal zone members and non-principal peer members defined within a single zone. Peer zoning allows the principal zone members to communicate with non-principal peer members.

Within a zone, principal and peer members can communicate with one another. However, peer members cannot communicate with other peer members, and principal members cannot communicate with other principal members. If multiple principal members are present within the same peer zone, they will not be visible to one another, nor will they be able to communicate with one another. Peer zoning supports LSAN and QoS peer zoning.

In Fabric OS 8.1.0 and later, you can add aliases to a peer zone. However, you cannot move a mixed-type alias to a peer zone. The alias members must be of the same type as the existing peer zone members; for example, if a peer zone is of type D, I, the alias must also be of the D, I type. If the switches in the fabric are not compatible with the configuration, an error message is displayed.

To configure the peer zone, perform the following steps.

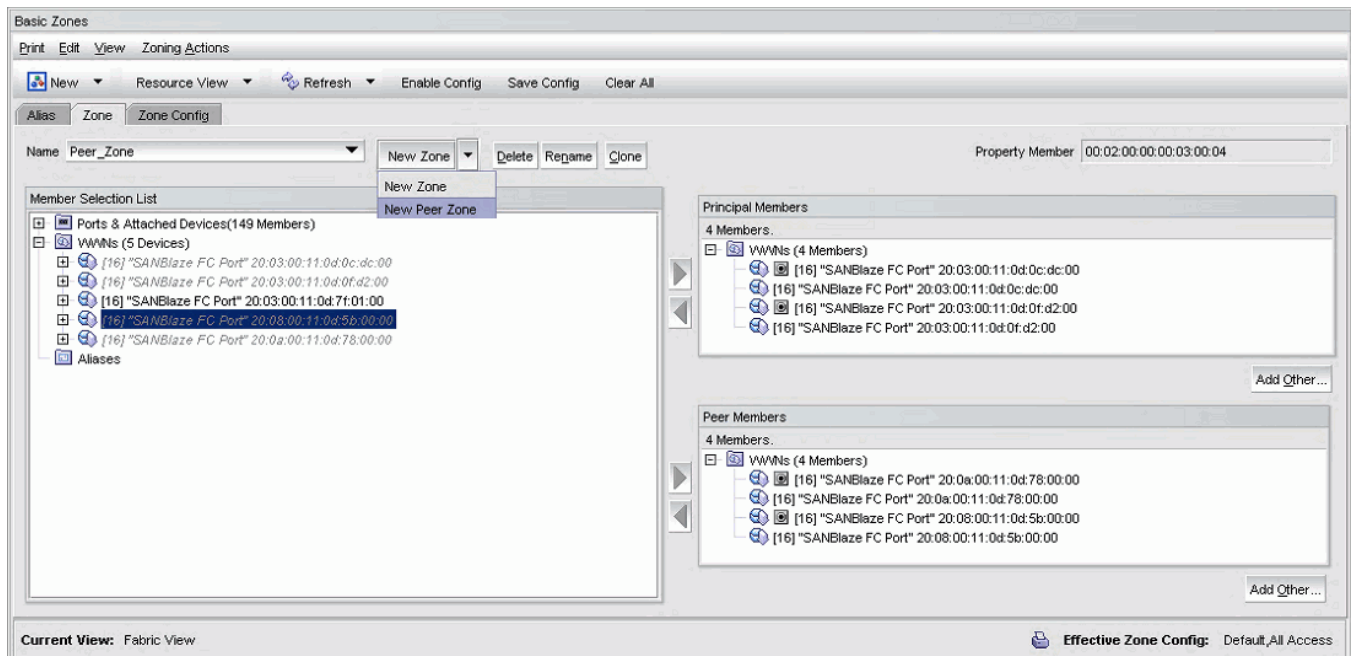
1. Select a switch from the Fabric Tree.
2. Click **Configure > Zone Admin**.
The **Zone Administration** window displays.
3. Select **Basic Zones > Zone > New Peer Zone**.
The **Basic Zones** dialog box is shown in the following figure.

NOTE

The zone object name can begin with a letter or numeral followed by alphanumeric characters, and it can contain one or more special characters such as -, \$, and ^. However, the name cannot start with a special character. This feature is available only when all switches in the fabric are upgraded to Fabric OS 8.1.0 or later.

4. Select the members from the **Potential Members** list, and click the right arrow button to move the members to the **Principal Members** list or the **Non Principal Members** list.

FIGURE 47 Peer Zones



The Property Member is read-only and is generated as per the is Fabric OS algorithm. All members have to be either of type D, I, or of type WWN. Mixed-type members within the same zone are not allowed. The default logical display order of the peer zone member is as follows:

1. Property member
2. Principal member
3. Non-principal member

Target Driven Zoning

A Target Driven peer zone is a peer zone that is configured in a fabric through a target. Target Driven Zoning is a variant of peer zoning wherein a device, usually a target, can manage peer zones by itself. The target device must be the principal device of the peer zone.

Target Driven Zoning mode can be used only in read-only mode. You can only read, delete, activate, or deactivate the members present in the Target Driven peer zone. The principal and non-principal members should be only the WWN type. You can view the Target Driven peer zones in the **Zone Administration** window.

You can configure Target Driven Zoning mode on multiple devices connected to the Target Driven peer zone-enabled ports. Refer to [Enabling Target Driven Zoning Mode](#) on page 95 to enable a Target Driven peer zone on a port. A Target Driven peer zone is supported on all port types.

NOTE

From Fabric OS 8.2.0 and later, Target Driven Zones are allowed to have aliases as members. Aliases that are members of Target Driven Zones are not allowed to be edited from Web Tools.

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 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. However, the LSAN zone name can contain one or more special characters such as -, \$, and ^.
- 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 the following table (the names are not case-dependent). The QoS zone name can contain one or more special characters such as -, \$, and ^.

TABLE 16 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 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 latest Fabric OS level. Refer to [Best practices for zoning](#) on page 149 for more recommendations about zoning.

Opening the Zone Administration window

Launching the **Zone Administration** window and performing any kind of zone configuration take more time if there are a large number of entries in the zone database. If the zone count is above 10,000, the time taken for completing the operation increases proportionately.

You cannot open the **Zone Administration** window from AD255 (physical fabric).

To open a **Zone Administration** window, perform the following steps.

1. Select a switch from the **Fabric Tree**.
2. Click **Configure > Zone Admin**.

The **Zone Administration** window displays, as shown in [Figure 48](#) on page 133.

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.

NOTE

You should not change the default zone mode from No Access to All Access if there is no effective zone configuration and more than 120 devices are connected to the fabric.

To set the default zoning mode, perform the following steps.

1. Open the **Zone Administration** window (refer to [Opening the Zone Administration window](#) on page 132).
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** window. 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 is not activated.

You must be logged in to the switch using a user name with one of the following roles to make changes to the zoning: zoneAdmin, admin, fabricAdmin, or any user-defined role with modify rights. All other roles allow only view or read-only access. Most of the zoning operations are disabled in read-only mode.

A snapshot is taken of all 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 135.

When you log in to a virtual switch or select a virtual switch using the drop-down list under the **Fabric Tree** section in the **Switch Explorer** window, only the ports that are associated with the virtual fabric ID that you selected are displayed in the member selection list, as shown in the following figure. You can use the **Add Other** button to add ports of other switches in the fabric.

FIGURE 48 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 that you make in the **Zone Administration** window to the zoning database on the switch, refer to [Saving local zoning changes](#) on page 135.

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.

The **Member Selection List** lists only the ports of the current switch and the devices of all switches in the fabric. Slot and port information of other switches is 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 Brocade Network Advisor to analyze zone configurations.

In the **Member Selection List**, you can differentiate between node WWNs and port WWNs by their icons, as shown in [Figure 49](#) and [Figure 50](#).

FIGURE 49 Port WWN icon for host

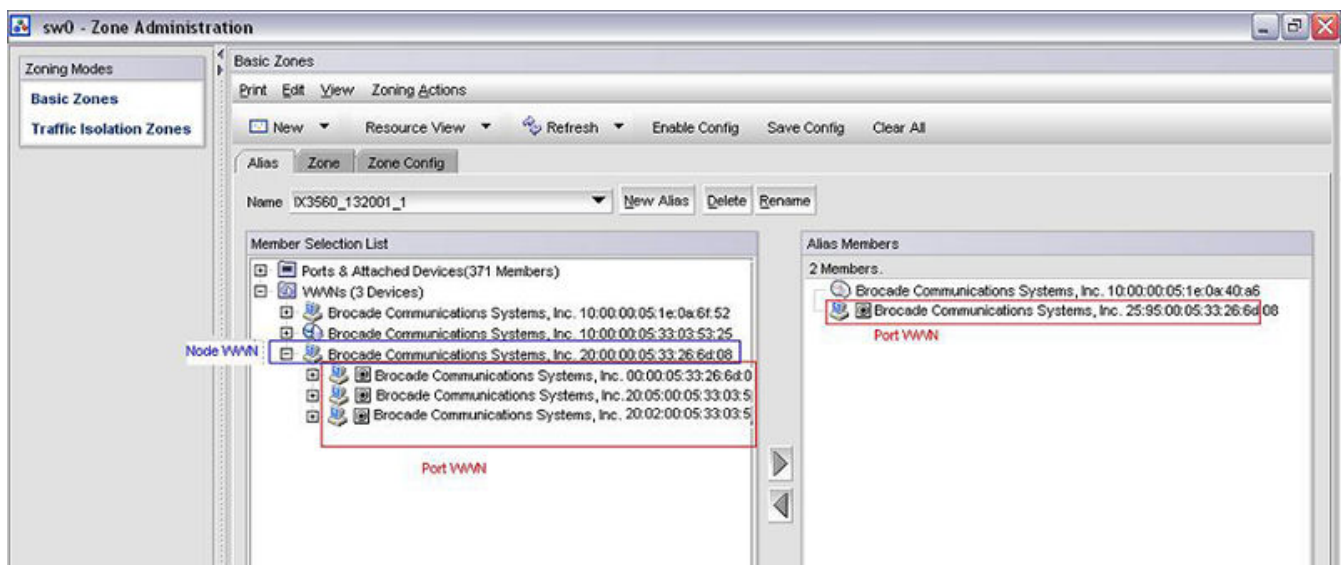



FIGURE 50 Port WWN icon for device



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 that you are logged in to, you may not be able to modify zones or ports in other 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 that 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 to 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 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 that you made and have not yet saved are erased from the buffer and replaced with the currently enabled zone configuration information that is saved on the switch.

To refresh the **Zone Administration** window, perform the following steps.

1. Launch the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
2. Select **View** > **Refresh Zoning** or click **Refresh**.

This redisplay 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 134. 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. Any other user looking at the zone information for the switch does 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.

To save the local zone changes, perform the following steps.

1. Make the zoning changes in the **Zone Administration** window.
2. Select **Zoning Actions > Save Config**.

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 144.

Selecting a zoning view

You can define how zoning elements are displayed in the **Zone Administration** window. The zoning view that you select determines how members are displayed in the **Member Selection List** panel ([Figure 48](#) on page 133). 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 that 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 and a list of the attached and imported physical devices (by WWN). 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.

To define the view of the fabric resource, perform the following steps.

1. Launch the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
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 providing 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 144.

To create a zone alias, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
2. Select a format to display zoning members in the **Member Selection List** as described in [Selecting a zoning view](#) on page 136.
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**.

NOTE

The zone object name can begin with a letter or numeral followed by alphanumeric characters, and it can contain one or more special characters such as -, \$, and ^. However, the name cannot start with a special character. This feature is available only when all switches in the fabric are upgraded to Fabric OS 8.1.0 or later.

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 right arrow becomes active.

You can move only port WWNs (WWPNs) to include in the alias. However, WWNNs (node WWNs) in the member selection list are grayed out.

7. Click the right arrow 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.

NOTE

In Fabric OS 8.1.0 or later, you can define a zone alias as a peer zone member. If the fabric consists of switches with firmware before Fabric OS 8.1.0, Web Tools blocks the creation of the peer zone alias.

Adding and removing members of a zone alias

For more information on enabling the configuration, refer to [Enabling zone configurations](#) on page 144.

NOTE

When you assign a node WWN to an alias or zone, only the WWPNs associated to that node are moved. However, the Node WWN (WWNN) cannot be moved. If the zone has a WWNN in the member selection list, the WWNN is grayed out. This functionality is supported only for IMO mode. This behavior is duplicated in Brocade Network Advisor zoning. This functionality is supported only by selecting the node WWN and assigning it to the alias or zone.

To add or remove zone alias members, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
2. Select the **Alias** tab.
3. Select the alias that 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 the right arrow to add the selected alias member, or click the left arrow 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 144.

To change the name of a zone alias, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
2. Select the **Alias** tab and select the alias that 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**.

NOTE

The zone object name can begin with a letter or numeral followed by alphanumeric characters, and it can contain one or more special characters such as -, \$, and ^. However, the name cannot start with a special character. This feature is available only when all switches in the fabric are upgraded to Fabric OS 8.1.0 or later.

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.

To delete a zone alias, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
2. Select the **Alias** tab.
3. Select the alias that 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 144.

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.

To create a zone, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
2. Select a format to display zoning members in the **Member Selection List** as described in [Selecting a zoning view](#) on page 136.
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**.

NOTE

The zone object name can begin with a numeral followed by alphanumeric characters, and it can contain one or more special characters such as -, \$, and ^. However, the name cannot start with a special character. This feature is available only when all switches in the fabric are upgraded to Fabric OS 8.1.0 or later.

LSAN zones and QoS zones have specific naming requirements:

- For LSAN zones, refer to [LSAN zone requirements](#) on page 131.
- For QoS zones, refer to [QoS zone requirements](#) on page 131.

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 right arrow becomes active.
8. Click the right arrow 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 144.

Adding and removing members of a zone

For information on enabling the configuration, refer to [Enabling zone configurations](#) on page 144.

NOTE

When you assign a node WWN to an alias or zone, all WWPNS associated to that node are also moved. This functionality is supported only for IMO mode. This behavior is duplicated in Brocade Network Advisor zoning. This functionality is supported only by selecting the node WWN and assigning it to the alias or zone.

To add or remove zone members, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
2. Select the **Zone** tab.
3. Select the zone that 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 the right arrow to add a zone member, or click the left arrow 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 144.

To change the name of a zone, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
2. Click the **Zone** tab.
3. Select the zone that 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.

NOTE

The zone object name can begin with a letter or numeral followed by alphanumeric characters, and it can contain one or more special characters such as -, \$, and ^. However, the name cannot start with a special character. This feature is available only when all switches in the fabric are upgraded to Fabric OS 8.1.0 or later.

6. Select **Zoning Actions** > **Save Config** to save the configuration changes.

Cloning zones

To clone a zone configuration, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
2. Click the **Zone** tab.
3. Select the zone that 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 144.

To delete a zone, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
2. Click the **Zone** tab.
3. Select the zone that 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 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 that the following conditions are satisfied:

- All switches in the fabric should have Fabric OS 6.4 or later.
- A port can be assigned to multiple TI zones that have the same failover state.

To create and populate an enhanced TI zone, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
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 right arrow becomes active.

6. Click the right arrow 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 6.4.0 or later, and all 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 Isolation 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.

The following figure 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 51 Sample zoning database



Creating zone configurations

To create a zone configuration, perform the following steps. 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 144.

NOTE

Any changes made to the currently enabled configuration do not display until you re-enable the configuration.

To create zone configurations, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
2. Select a format to display zoning members in the **Member Selection List** as described in [Selecting a zoning view](#) on page 136.
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**.

NOTE

The zone object name can begin with a letter or numeral followed by alphanumeric characters, and it can contain one or more special characters such as -, \$, and ^. However, the name cannot start with a special character. This feature is available only when all switches in the fabric are upgraded to Fabric OS 8.1.0 or later.

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 right arrow becomes active.

7. Click the right arrow 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 144.

To add or remove members of a zone configuration, perform the following steps.

NOTE

You can make changes to a configuration that is currently enabled; however, changes do not display until you re-enable the configuration.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
2. Select the **Zone Config** tab.
3. Select the configuration that 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 the right arrow to add a configuration member or the left arrow 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.

NOTE

You cannot rename the currently enabled configuration.

To rename the zone configuration, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
2. Select the **Zone Config** tab.

3. Select the configuration that 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**.

NOTE

The zone object name can begin with a letter or numeral followed by alphanumeric characters, and it can contain one or more special characters such as -, \$, and ^. However, the name cannot start with a special character. This feature is available only when all switches in the fabric are upgraded to Fabric OS 8.1.0 or later.

The configuration is renamed in the configuration database.

5. Select **Zoning Actions** > **Save Config** to save the configuration changes.

Cloning zone configurations

To clone a zone configuration, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
2. Select the **Zone Config** tab.
3. Select the zone configuration that 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 are made to the effective configuration. You do not need to enable the configuration.

Deleting zone configurations

To delete a zone configuration, perform the following steps.

NOTE

You cannot delete an enabled configuration.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
2. Select the **Zone Config** tab.
3. Select the configuration that 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 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.

To enable the zone configuration, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
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.

To disable a zone configuration, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
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 window 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 Admin** 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

To view the enabled zone configuration name, perform the following steps.

Select a logical switch from the **Logical Switch** list in the top-right corner of the **Switch Explorer** window. The selected switch displays in the **Switch View**.

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

To view detailed information about the enabled zone configuration, perform the following steps.

1. Open the **Zone Administration** window, as described on [Opening the Zone Administration window](#) on page 132.

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 if 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 your Zone Admin session, those changes are applied.

2. From Fabric OS 8.2.0, the Web Tools zone configuration loads the latest checksum in the **Zone Administration** window.

When you perform a **save config** operation, the configuration will be saved only if the checksum value matches the checksum currently calculated on the switch; otherwise, Web Tools will show an error message of checksum mismatch as shown in the following figure, and it will refresh the page with the latest information.



CAUTION



3. To identify the most recently effective zone configuration without saving or applying any changes that 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.

4. Optional: Click **Print** in the **Print Effective Zone Configuration** dialog box to print the enabled zone configuration details.

NOTE

You must use Brocade Network Advisor to print the zone database summary configurations, to display zone configuration summaries, and to 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 before 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.

To add a WWN, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
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 zoning elements that include the new WWNs. All elements are selected by default.

4. Click items in the list to select or clear, and click **Add** to add the new WWNs to all selected zoning elements.

The WWN is added to the Zone Admin buffer and can be used as a member.

Different icons are used to differentiate between node WWN and port WWN.

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 132.
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 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 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.

To replace a WWN in multiple aliases and zones, perform the following steps.

1. Launch the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
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 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, 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, zones are also included, along with the elements that 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**.

To search for zone members, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
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.

To clear the zone database, perform the following steps.

1. Open the **Zone Administration** window as described in [Opening the Zone Administration window](#) on page 132.
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 BrocadeNetwork Advisor 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 latest Fabric OS-level switch.
Switches with earlier Fabric OS versions do not have the capability to view all the functionality that a newer Fabric OS provides since functionality is backward-compatible but not forward-compatible.
- Zone using the core switch versus an edge switch.
- Zone using a director instead of a switch.
A director has more resources to handle zoning changes and implementations.
- Zone on the switch that you connect to when bringing up Web Tools (the proxy switch).

Refer to the *Brocade Fabric OS Administration Guide* for further best practices.

Working with Diagnostic Features

• Trace dumps.....	151
• Displaying switch information.....	153
• Port LED interpretation.....	156

Trace dumps

A trace dump is a snapshot of the running behavior within a 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 runs on the switch, 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.
- A kernel panic occurs.
- The hardware watchdog timer expires.

For information about the **tracedump** command, refer to the *Brocade Fabric OS Command Reference*.

For Gen 5 platforms: 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.

For Gen 6 platforms: The trace dumps are created in types of dumps and are listed on the switch. If AutoFTP is enabled, the data is deleted once it is transferred to the FTP host.

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, Fabric OS runs the **supportsave** command on the affected switch. This command packages all error logs, the **supportshow** output, and the trace dump into a single file, and if automatic trace dump transfers are enabled, it moves the file to your FTP server. Refer to [Setting up automatic trace dump transfers](#) on page 152 for instructions on configuring your switch to automatically copy trace dumps to your FTP server.

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 **supportsave** command, refer to the *Brocade 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 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

To specify a remote server, perform the following steps.

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 path for Windows is *Folder Name\FileName.txt* or *FileName.txt*.

The path for Linux is *Directory Name/FileName.txt* or *FileName.txt*.

The password is optional if you log in as an anonymous user.

5. Click **Apply**.

Enabling automatic transfer of trace dumps

To enable the automatic transfer of trace dumps, perform the following steps.

NOTE

You cannot enable the automatic transfer feature until you have specified a remote server. Refer to [Specifying a remote server](#) on page 152.

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.

To disable automatic trace uploads, perform the following steps.

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 the **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 Web Tools login to enable the export and copy functionality.


Enter the text string in the box that displays on the table, as shown in [Figure 53](#) on page 155, and press **Enter**. This is an incremental search and allows a maximum of 24 characters including the 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.

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 the following figure.

FIGURE 52 Fan States window



Fan No.	State	Speed (RPM)	Airflow Direction
1	Ok	7039	Non-portside Intake
2	Ok	7187	N/A

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. The **Fan No** column indicates the fan FRU number when it is available; otherwise it displays the fan number.

The **Airflow Direction** column displays the direction state as either **Non-portside Exhaust (E)** or **Non-portside Intake (I)** for the Brocade 6510 and Brocade 6520. For all other hardware, the displayed value will be **N/A**.

NOTE

If the Fan States window has no "Fan Speed" column, *the speed is not monitored*.

To view the detailed fan status of a switch, perform the following steps.

1. Select a logical switch from the **Logical Switch** list in the top-right corner of 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 the **Fan** button.

The detailed fan status for the switch displays, as shown in the previous figure.

Viewing the temperature status

The icon on the **Temp** button indicates the overall temperature status. For more information regarding switch temperature, refer to the appropriate hardware documentation.

To view the temperature status, perform the following steps.

1. Select a logical switch from the **Logical Switch** list in the top-right corner of the **Switch Explorer** window.
The selected switch displays in the **Switch View**. The icon on the **Temp** button indicates the overall temperature status.


- Click **Temp** on the **Switch View**.

NOTE

Note that the **Absent** state indicates an empty blade slot.

The detailed temperature sensor states for the switch are displayed, as shown in the following figure.

FIGURE 53 Temperature Sensor States for Chassis window



Thermal Senso...	State	Centigrade	Fahrenheit	Sensor Index	Slot
1	Ok	28	82	0	3
2	Ok	33	91	1	3
3	Ok	55	131	2	3
4	Ok	53	127	3	3
5	Absent	0			4
6	Ok	45	113	0	1
7	Ok	37	98	1	1
8	Ok	44	111	2	1
9	Ok	43	109	3	1
10	Absent	0			2
11	Ok	27	80	0	5
12	Ok	51	123	1	5
13	Ok	46	114	2	5
14	Ok	28	82	0	6
15	Ok	42	107	1	6
16	Ok	48	118	2	6
17	Ok	26	78	0	7
18	Ok	56	132	1	7
19	Ok	54	129	2	7
20	Ok	38	100	3	7
21	Ok	37	98	4	7
22	Ok	51	123	5	7
23	Ok	52	125	6	7
24	Ok	53	127	7	7
25	Ok	27	80	0	8
26	Ok	28	82	1	8
27	Ok	51	123	2	8
28	Ok	54	129	3	8

Viewing the power supply status

The icon on the **Power** button indicates the overall status of the power supply. For more information regarding switch power modules, refer to the appropriate hardware documentation.

FIGURE 54 Power Supply States window

Power Supply No.	State	Type
1	Ok	AC
2	Ok	N/A

To view the power supply status, perform the following steps.

1. Select a logical switch from the **Logical Switch** list in the top-right corner of the **Switch Explorer** window. 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, as shown in the previous figure. The **Type** column displays either **AC**, **DC**, or **N/A** value.

Port LED interpretation

The Switch View that 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

32-Gbps Brocade switches and port blades do not have port speed LEDs, 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)
- Violet (port disabled)
- White (SFP connected but not online)

The following figures show port icons and associated LEDs of a switch. The layout is similar for all switches, but varies according to switch type. The blinking LEDs indicate port status.

FIGURE 55 Example for green (healthy) and red (critical) LEDs



The following figure shows the port icon toggling effect. When you mouse-over a port icon, the icon is highlighted with a yellow border.

FIGURE 56 Example for yellow (marginal) LED




FIGURE 57 Example for gray (unmonitored) and violet (port disabled) LEDs




FIGURE 58 Example for dimmed (unlicensed) LED




FIGURE 59 Example for white (SFP connected but not online) LED



Using the FC-FC Routing Service

• Fibre Channel Routing overview.....	159
• Supported switches for Fibre Channel Routing.....	160
• Setting up FC-FC routing.....	160
• FC-FC routing management.....	160
• Viewing EX_Ports.....	162
• Configuring an EX_Port.....	162
• Configuring FCR router port cost.....	163
• Viewing LSAN zones.....	163
• Configuring the backbone fabric ID.....	164

Fibre Channel Routing overview

Fibre Channel Routing (FCR) 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 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 the FC-FC Routing Service.
Inter-fabric link (IFL)	The link between an E_Port and an EX_Port or between 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 for the Fabric OS 8.2.0 release:

- Brocade 6510
- Brocade 6520
- Brocade G620
- Brocade G630
- Brocade 7840 Extension Switch
- Brocade DCX 8510-4 and Brocade DCX 8510-8, Brocade X6-4 and Brocade X6-8 directors when configured with FX8-24, FC16-32, FC16-48, FC16-64, FC32-48, FC32-64 blades or CR16-8 and CR32-8 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 164.
2. On the FC router, ensure that the ports to be configured as EX_Ports are either not connected or 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 162.

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 163.

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 and an Integrated routing license. You can manage FC-FC routing through the FC routing module. 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 last refresh time is displayed in the lower-left corner of the subtabs.

The switch must be FC router-capable, as described in [Fibre Channel Routing overview](#) on page 159.

You need to configure only EX_Ports and the backbone fabric ID on the FC router. You can configure LSAN zones on the fabrics 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.

To modify the data, you must log in as switchadmin, fabricadmin, basicswitchadmin, operator, or any user-defined role configured with modify rights. 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 device tabs continue to display the existing entries, but display the entries related to the backbone fabric only. All EX_Ports are disabled, and you cannot enable them until FC-FC Routing is enabled.

Opening the FC Routing module

The **FCR** submenu under **Configure** launches the FC Routing module. This module is displayed only for the following switches:

- Brocade 6510
- Brocade 6520
- Brocade 7840 Extension Switch
- Brocade G620
- Brocade G630
- Brocade DCX 8510 enterprise-class platform, when configured with FX8-24, FC16-32, FC16-48, FC32-64, CR16-8, or FC16-64 blades
- Brocade X6 enterprise-class platform, when configured with FC32-48, SX6, and Core blade (CR32-8 and CR32-4)

NOTE

When the Virtual Fabrics capability is enabled on the switch, Fabric ID cannot be set using the **Set Fabric ID** button.

To open the FC Routing module, perform the following steps.

1. Select a logical switch from the **Logical Switch** list in the top-right corner of the **Switch Explorer** window.
The selected switch displays in the **Switch View**.
2. Click **Configure** > **FCR**.

The FC Routing module displays. If FC-FC Routing is disabled, a message to that effect displays on all tabs in the module.

Viewing and managing LSAN fabrics

The **LSAN Fabric** tab displays all 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.

Viewing EX_Ports

The **EX_Ports** tab displays all 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

To disable FC Routing, you must disable all EX_Ports and VEX_Ports. You cannot enable these ports until FC Routing is enabled.

For more detailed information about a specific port, click a port name in the table, and select **Properties** from the **Actions** list. 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_Port configuration
- Rename an EX_Port
- Swap the port index of an EX_Port (described in [Port swapping index](#) on page 98)
- 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.

Configuring an EX_Port

NOTE

From Fabric OS 7.2.0 and later, EX_Ports can also be configured on ICL ports (only in the base switch).

To configure an EX_Port, perform the following steps.

1. Select **Configure** > **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.

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

To edit the configuration of an EX_Port, perform the following steps.

1. Select **Configure** > **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 a 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 less than the others, then the lowest cost path is used.

Every link has a default cost. For an EX_Port, 4-Gbps, 8-Gbps, 10-Gbps, 16-Gbps, and 32-Gbps links, the default cost is 1,000. For a VEX_Port, the default cost is 10,000. If the cost is set to 0, the default cost is used for that link.

To configure the FCR router port cost, perform the following steps.

1. Click **Configure** > **FCR**.
2. Click the **EX_Ports** tab.
3. Disable the EX_Port.
4. Click the **Router Port Cost** button.

Viewing LSAN zones

The **LSAN Zones** tab displays all LSAN zones, in both 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 a mapping of LSAN zones with the edge fabric with which they are going to communicate. 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 138 to create LSAN zones.

Viewing LSAN devices

The **LSAN Devices** tab displays information about the physical and proxy devices, and it 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 physical and proxy LSAN devices. Note that this count is for all 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. You should manually enable FCR after setting the backbone FID. However, you must first disable all EX_Ports before you begin this operation. After the fabric ID is changed, you must re-enable these ports.

NOTE

When Virtual Fabrics capability is enabled on the switch, the fabric ID cannot be set using the **Set Fabric ID** button.

To configure the backbone fabric ID, perform the following steps.

1. Select **Configure** > **FCR**.
2. Select the **EX_Ports** tab.
3. Select all EX_Ports in the table, and click **Disable**.
4. Select the **General** tab.
5. Click **Set Fabric ID** in the task bar.

The **Configure Backbone Fabric ID** window displays.

6. 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.

7. Click **OK**.
8. Click **Enable FCR** in the task bar.
9. Select all EX_Ports in the table, and click **Enable**.

Using the Access Gateway

- Access Gateway overview.....165
- Viewing Switch Explorer for Access Gateway mode..... 165
- Access Gateway mode167
- Enabling Access Gateway mode.....167
- Disabling Access Gateway mode..... 168
- Viewing the Access Gateway settings.....168
- Port configuration.....168
- Access Gateway policy modification.....172
- Enabling the Automatic Port Configuration policy.....173

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 to Access Gateway mode to transform 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 it simplifies configuration and management in a large fabric by reducing the number of domain IDs and ports.

For detailed descriptions of the Access Gateway, refer to the *Brocade Access Gateway Administration Guide*.

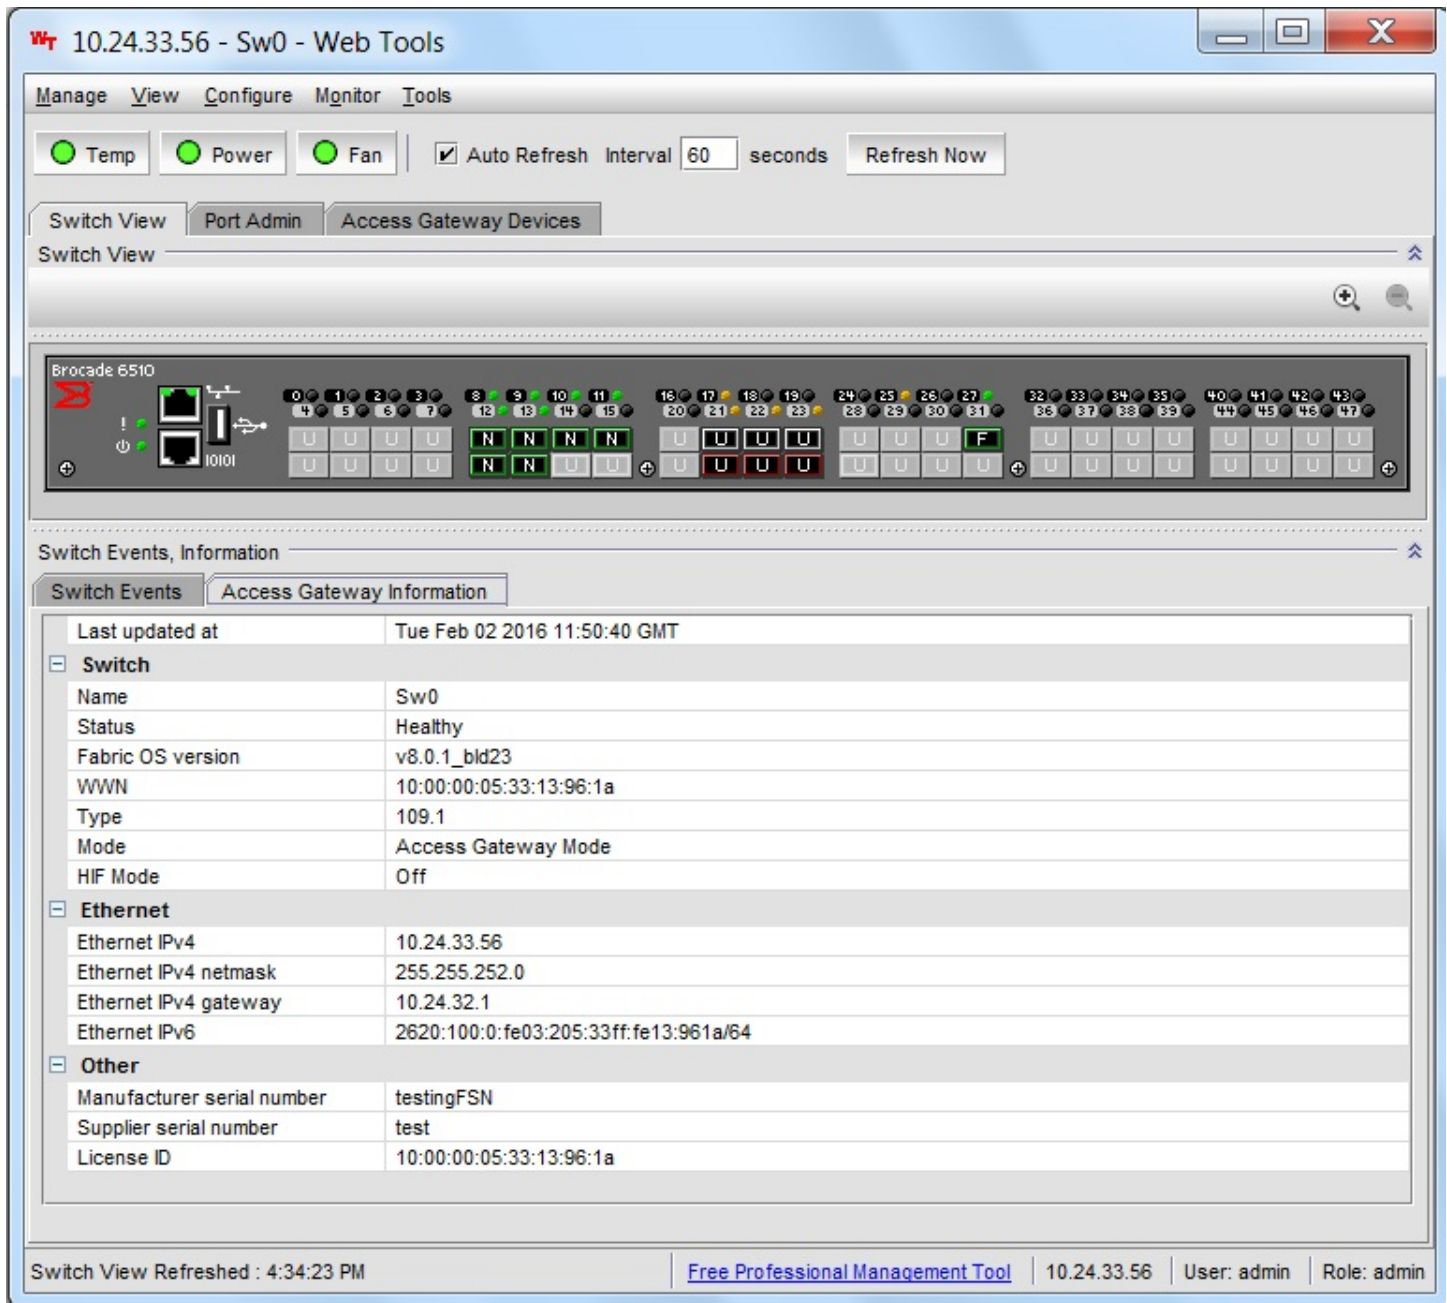
NOTE

When Access Gateway mode is enabled on switches that are 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 the following figure.

FIGURE 60 Switch Explorer view for Access Gateway mode



The Access Gateway mode Switch Explorer is divided into the following areas:

- Menu bar
- **Switch View** buttons
- **Switch View**, **Port Admin**, and **Access Gateway Devices** tabs
- **Switch Events** and **Access Gateway Information**
- Indicator bar

- Professional Management Tool offering

Access Gateway mode

The Access Gateway feature on the Brocade G620 switch enables interoperability with Cisco fabrics. The Access Gateway mode of the switch presents standard F_Ports to the hosts, but it connects to the enterprise fabric as an N_Port (rather than as an E_Port in the case of a regular switch).

Restricted access in the Port Admin tab

When Access Gateway mode is enabled, the following options can be configured in Access gateway mode:

- **Port Configuration Policy** – You can select Auto or Advanced mode (default). When Auto mode is selected, options like **Configure N-Port Groups**, **Configure F-N Port Mappings**, and **N Port configuration** are disabled.
- **Trunking** – You can enable and disable N_Port trunking.
- **Configure N-Port Groups** – You can configure the port group details from the Port Group Configuration window.
- **Configure F-N Port Mappings** – Add (right arrow) and Remove (left arrow) buttons are disabled for primary mappings and secondary failover mappings.
- **N Port Configuration** – By default, all ports are set to N_Ports, and failover and fallback are disabled. You can edit the speed. The following options are enabled 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 (AG) 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, in the **Switch Explorer** window, click **Configure** > **Switch Admin**, and then select the **Configure** > **Upload/Download** subtab and upload the configuration file.

You cannot enable Access Gateway mode if Management Server is enabled. To disable Management Server, enter the **MsplmgmtDeactivate** command.

Access Gateway mode is unavailable when VF is enabled.

NOTE

If any error is encountered while enabling AG mode, the switch is disabled and remains in the disabled state until you manually enable it.

To enable Access Gateway mode, perform the following steps.

1. Select a switch.
2. Click **Configure** > **Switch Admin**.
The **Switch Administration** dialog box displays.
3. Click **Enable** in the **Access Gateway Mode** section.

4. Click **Apply**.
Web Tools closes all open dialog boxes and re-launches again with AG mode enabled.

Disabling Access Gateway mode

To disable Access Gateway mode, perform the following steps.

1. Select a switch.
2. Click **Configure** > **Switch Admin**.
The **Switch Administration** dialog box displays.
3. Click **Disable** in the **Access Gateway Mode** section.
4. Click **Apply**.
Web Tools closes all open dialog boxes and re-launches again with AG mode disabled.

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 the **Access Gateway Devices** tab.

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 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 **Port Configuration** wizard 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.

Editing a port

Editing a port allows you to configure port types and port speed.

With the **Edit** dialog box, you can configure allowed port types and port speed for physical ports. To edit a port, perform the following steps.

1. Click a port in the **Switch View** to open the **Port Admin** tab.
2. Select the **FC Ports** tab.
3. Select the port that you want to configure from the tree on the left.

4. Select **Edit** from **Actions** list.
The **Edit** dialog box displays as shown in the following figure.

FIGURE 61 Edit dialog box

5. Select the port to configure the allowed port types and port speed.
6. Click **OK** to save the changes.

Port editing notes

Note the following when editing a port:

- Long distance is not displayed from the **Edit** window.
- The Auto Max speed levels are displayed only when you set the port speed as Auto Negotiate; these options allow you to set the speed limit that the port can auto-negotiate.

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 the N_Port grouping.

To create port groups, perform the following steps.

1. Click a port in the **Switch View** to open the **Port Admin** tab.
2. Select **Advanced** from **Configure** > **Port Configuration Policy**.
3. Select a port or ports to configure.

4. Select **Configure N-Port Groups** from the **Actions** list.

NOTE

Configure N-Port Groups is unavailable if you select **Automatic** from the **Port Configuration Policy**.

5. In the **Port Group Configuration** dialog box, click **Add**.
The **Add Port Group** window displays.
6. Enter the ID for the new port group in the **Port Group ID*** field.
7. Enter the name for the new port group in the **Port Group Name** field.
8. Select the **Login Balancing** check box to enable login balance for the port group.
9. Select the **Fabric Name Monitoring** check box to manually configure managed fabric name monitoring.
10. Under the **Select Members (N-Port)*** section, select the required ports to group.
11. Click **Save**.

Editing or viewing port groups

To edit port groups, perform the following steps.

1. Click a port in the **Switch View** to open the **Port Admin** tab.
2. Select **Configure N-Port Groups** from the **Actions** list.
3. On the **Port Group Configuration** dialog box, select the group that you want to edit, and 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 the **Manual Balancing** button become enabled.
6. Click **Failover Enable**.
A confirmation dialog box displays.
7. Click **Yes** to enable failover to all 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 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 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 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 to group, and clear the check boxes for the ports 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).

To delete port groups, perform the following steps.

1. Click a port in the **Switch View** to open the **Port Admin** tab.
2. Select **Configure N-Port Groups** from the **Actions** list.
3. In the **Port Group Configuration** dialog box, select the group to delete, and then click **Delete**. A confirmation dialog box displays.
4. Click **Yes** to confirm the action.

Defining custom primary F-N port mappings

To manually change primary F-N port mappings, perform the following steps.

1. Click a port in the **Switch View** to open the **Port Admin** tab.
2. Click the **FC Ports** tab.
3. Select **Configure F-N Port Mappings** from the **Actions** list.
4. Select the **Primary Mappings** subtab on the right side of the dialog box.
5. In the **Primary Mappings** area, select ports and use the **Add** (right arrow) button to map F_Ports or U_Ports to N_Ports.
6. Optional: Use the **Remove** (left arrow) button to delete an F_Port mapping from an N_Port.
7. 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.

8. After you have made the appropriate changes, click **Save**.

Defining custom static F-N port mappings

NOTE

Static mappings and custom WWN-N port mappings are mutually exclusive.

To manually change static F-N port mappings, perform the following steps.

1. Click a port in the **Switch View** to open the **Port Admin** tab.
2. Click the **FC Ports** tab.
3. Select **Configure F-N Port Mappings** from the **Actions** list.
4. Select the **Static Mappings** subtab on the right side of the dialog box.
5. In the **Primary Mappings** area, select ports and use the **Add** (right arrow) button to map F_Ports or U_Ports to N_Ports.
6. Optional: Use the **Remove** (left arrow) button to delete an F_Port mapping from an N_Port.
7. After you have made the appropriate changes, click **Save**.

Defining custom WWN-N port mappings

NOTE

Static mappings and custom WWN-N port mappings are mutually exclusive.

To manually change WWN-N port mappings, perform the following steps.

1. Select the **Port Admin** tab.
2. Click the **FC Ports** tab.
3. Select **Configure WWN-N Port Mappings** from the **Actions** list.
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 pane and select the WWN, and then use the **Remove** (left arrow) button to remove the mapping.
7. Optional: Define a failover in the **Secondary Failover Mappings** area by selecting the ports using the **Add** (right arrow) and **Remove** (left arrow) 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 provides only the ability to enable and disable the policies. For more information on these policies, refer to the *Brocade Access Gateway Administration 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 to which these policies are mapped 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

To modify Path Failover and Failback policies, perform the following steps.

1. Click a port in the **Switch View** to open the **Port Admin** tab.
2. Select the N_Port for which you want to modify the policy.
3. Select **Edit** from the **Actions** list.
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 Mappings**, and **Configure WWN-N Port Mappings** options are unavailable.

To enable auto-rebalancing from the **Switch Administration** window, perform the following steps.

1. Click a port in the **Switch View** to open the **Port Admin** tab.
2. Select **Automatic** from **Configure > Port Configuration Policy**.

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 Admin** tab.

3. Click **Yes** in the confirmation window.

4. In the **Switch Explorer** window, click **Configure** > **Switch Admin**. The **Switch Administration** window displays.

FIGURE 62 Access Gateway Auto Rebalancing

SW300_155 - Switch Administration

SwitchName: SW300_155 VMM: 10:00:00:05:1e:07:b5:bf Thu Jul 09 2009 20:03:38 GMT+00:00

Switch Network Firmware Download License User Trunking

Switch Name

Name: SW300_155 Manufacturer Serial #: ALJ0613D04K Supplier Serial #: none

Switch Status

☒ Enable ☐ Disable

Report

View Report

DNS Configuration

DNS Server 1: 10.32.2.10 DNS Server 2: 10.32.2.11 Domain Name: englab.brocade.com Remove All

Reboot/Fastboot

Reboot Fastboot

Access Gateway Mode

☒ Enable ☐ Disable ☒ N Port Auto Rebalancing ☒ F Port Auto Rebalancing Manual Balancing

Apply Close Refresh

[Switch Administration opened]: Thu Jul 09 2009 19:49:21 GMT+00:00

Change current switch settings Mode: Basic Free Professional Management Tool 10.35.52.155 User: admin Role: admin

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 the **F Port Auto Rebalancing** check box to enable F_Port rebalancing.
 - Click **Manual Balancing**. 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.

Administering Extended Fabrics

- [Extended link buffer allocation overview.....](#)175
- [Configuring a port for long distance.....](#)178

Extended link buffer allocation overview

If a 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 **Extended Fabric** tab displays information about the port speed, long-distance settings, and buffer credits, as shown in the following figure. For detailed information on managing extended fabrics, refer to the *Brocade Fabric OS Administration Guide*.

The **Extended Fabric** tab displays the following columns:

- **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:
 - 4G – 4 Gbps
 - 8G – 8 Gbps
 - 10G – 10 Gbps
 - 16G – 16 Gbps
 - 32G – 32 Gbps
 - N4 – Negotiated 4 Gbps
 - N8 – Negotiated 8 Gbps
 - N16 – Negotiated 16 Gbps
 - N32 – Negotiated 32 Gbps
 - Auto-Negotiation

NOTE

Auto-Negotiation is not supported on QSFP ports.

- **Buffer Needed** – The number of buffers needed. You can edit the buffer needed for LD and LS modes. When you change the **Buffer Needed** value, the **Frame Size** and **Desired Distance(km)** values cannot be changed.
- **Buffer Allocated** – The number of buffers actually allocated.
- **Recommended Buffer**– The number of recommended buffers. The recommended buffer value is non-editable. The default port configuration value displays when the Extended Fabric is launched for the first time. When you change **Frame Size** and **Desired Distance(km)**, the recommended buffer value changes according to the current port configuration. When the number of buffers needed is configured for a port, the recommended buffer value is set to N/A for the same port.
- **Remaining Buffer**– The number of remaining buffers. The remaining buffer value is non-editable. If the **Buffer Needed** value exceeds the remaining buffer value, a warning message displays.

- **Frame Size** – The size of the frame. When you edit the frame size value, the desired distance value can also be changed for LD and LS modes and vice versa. But the buffer value cannot be changed.
- **VC Link Init** – The fill words used on long-distance links. When set to IDLE (0) mode, the link uses IDLE fill words. When set to ARB (1) mode, the link uses the default ARB fill words.
- **Fill Word** --Fill word comprises the following modes:
 - **Mode 0** : Use IDLE in link in it and IDLE as the fill word.
 - **Mode 1**: Use ARB in link in it and ARB as the fill word.
 - **Mode 2** : Use IDLE in link in it and ARB as the fill word.
 - **Mode 3** : Try Mode 1 first; if it fails, then try Mode 2.
- **Link Distance(km)** --The actual distance of the link.
- **Desired Distance(km)** -- Required for a port configured in LD or LS mode (as shown in the following table), 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** -- The following table describes the long-distance settings and identifies which settings require a Brocade Extended Fabrics license.

FIGURE 63 Extended Fabric tab

SwitchName: DCX_1 DomainID: 14(0x72) VMM: 10.00.00:05:33:40:28:00 Thu Apr 10 2014 02:29:14 GMT

Switch Network Firmware Download License User Blade Trunking SNMP Configure Routing Extended Fabric AAA Service Trace FICON CLIP Security Policies

Extended Fabric Administration

Port Number	Buffer Limited	Port Speed	Buffer Needed	Buffer Allocated	Recommended Buffer	Remaining Buffer	Frame Size	VC Link Init	Filterword	Link Distance(km)	Desired Distance(km)	Long Distance
0	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
1	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
2	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
3	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
4	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
5	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
6	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
7	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
8	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
9	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
10	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
11	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
12	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
13	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
14	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
15	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
16	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
17	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal
18	No	16G	0	0	25	4950	0	IDLE(0)	N/A	-	N/A	LO: Normal

Slot_1 Slot_2 Slot_3 Slot_4 Slot_5 Slot_6 Slot_7 Slot_8 Slot_9 Slot_10 Slot_11 Slot_12

Apply Close Refresh

[Switch Administration opened] Thu Apr 10 2014 02:08:10 GMT
Retrieve Distribution Parameters from the switch.
Retrieve IPFilter Policies from the switch.

Long Distance Port Configuration. Mode: Advanced Free Professional Management Tool 10.24.33.1 FID: 128 - DCX_1 User: admin Role: admin

For the Brocade DCX 8510-4 and Brocade DCX 8510-8, the slots for CPs are not available.

The Brocade G610, Brocade G620, and Brocade G630 Switches support auto-negotiated link speeds of 4, 8, 16, and 32 Gbps.

TABLE 17 Long-distance settings and license requirements

Value	Description	Extended Fabrics license required?
LO	No long-distance setting is enabled. The maximum supported link distance is: <ul style="list-style-type: none"> 10 kilometers at 1 Gbps 5 kilometers at 2 Gbps 2.5 kilometers at 4 Gbps 1 kilometer at 10 Gbps 500 meters at 16 Gbps 	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	Yes

TABLE 17 Long-distance settings and license requirements (continued)

Value	Description	Extended Fabrics license required?
	Gbps, depending on the switch platform and the availability of frame buffers within the port group.	
LS	<p>Static setting is enabled. Buffer credits for the given E_Port are statically configured based on the desired link distance.</p> <p>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 Gbps, depending on the switch platform and the availability of frame buffers within the port group.</p> <p>For the Brocade G620 Switch, the buffer credits are 10 through X; where X is proportional to the available buffers.</p> <p>For the Brocade G610 Switch, you can statically specify a distance that is greater than 10 km.</p>	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.

To configure a port for long distance, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Click **Show Advanced Mode**.
3. Select the **Extended Fabric** tab.
4. This step is switch-specific:

For the Brocade DCX 8510-8, Brocade DCX 8510-4, Brocade X6-8, and Brocade X6-4 enterprise-class platforms, click the slot subtab that corresponds to the correct slot for the logical switch.

For the Brocade 6505, Brocade 6510, Brocade 6520, and Brocade 7840 Extension Switch, proceed 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 17](#) on page 177.

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. For LD or LS options, the **Buffer Needed** column is made editable to specify the buffer needed value. On changing the buffer needed value, the **Frame Size** and **Desired Distance** values cannot be changed.

- a) Double-click the **Desired Distance** field for the port, as shown in [Figure 63](#) on page 177.
- 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.
 - For the Brocade G620, the buffer credits are 10 through X; where X is proportional to the available buffers.

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**.
7. Click **Yes** to apply the changes, or click **No** to close the confirmation message window.

Routing Traffic

• Routing overview.....	181
• Viewing Fabric Shortest Path First routing.....	182
• Configuring dynamic load sharing.....	182
• Specifying frame order delivery.....	183
• Configuring the link cost for a port.....	184
• E_Port balance priority.....	184

Routing overview

For Fabric OS 7.0.0 and later, the supported routing policies are:

- Port-based routing - Port-based routing assigns a "static route" in which the path chosen for traffic never changes.
- Exchange-based routing - The exchange-based routing policy is the default. The exchange-based routing policy always employs "dynamic path selection" in which the software defines a path based on current traffic conditions.
- Device-based routing - Device-Based Routing (DBR) is a read-only option. If DBR is set on the switch, the DBR radio button appears auto-selected and is unavailable in Web Tools.

Refer to the *Brocade Fabric OS Administration Guide* for more information.

To optimize port-based routing, the dynamic load sharing (DLS) feature 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. The following figure displays the **Routing** tab.

FIGURE 64 Routing tab

SwitchName: sw0 DomainID: 5(0x5) WWN: 10:00:00:05:33:65:ab:83 Thu Sep 11 2014 02:05:43 GMT+00:00

Switch Network Firmware Download License User Trunking SNMP Configure **Routing** Extended Fabric AAA Service Trace FICON CUP Security Policies

Advanced Performance Tuning (APT) Policy
☐ Port-Based-Routing [Info](#)
☒ Exchange-Based-Routing [Info](#)
☐ Device-Based-Routing [Info](#)

Dynamic Load Sharing
☒ On
☐ Off

In Order Delivery (IOD)
☐ On
☒ Off

Loss Less
☐ On
☒ Off

E-port Balance Priority
☐ On
☒ Off

Rebalance/Rebalance All
 Rebalance
 Rebalance All

Routing Table

In Port	Destination Domain	Out Port	Metric	Hops	Flags	Next Domain	Next Port
0	2	4	1000	2	D	23	40
0	2	5	1000	2	D	23	41
0	2	6	1000	2	D	23	42
0	2	7	1000	2	D	23	43
0	4	4	1500	3	D	23	40
0	4	5	1500	3	D	23	41
0	4	6	1500	3	D	23	42
0	4	7	1500	3	D	23	43
0	23	4	500	1	D	23	40
0	23	5	500	1	D	23	41
0	23	6	500	1	D	23	42

Apply Close Refresh

[Switch Administration opened: Thu Sep 11 2014 02:04:37 GMT+00:00]

Configure Routing Information Mode: Advanced [Free Professional Management Tool](#) 10.24.33.71 FID 128 - sw0 User: root Role: root

Viewing Fabric Shortest Path First routing

The **Routing** tab of the **Switch Administration** window displays information about routing paths.

To view the Fabric Shortest Path First (FSPF) routing, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Routing** tab.

Displaying routing path information is switch-specific. For the Brocade 6505, Brocade 6510, Brocade 6520, Brocade G610, Brocade G620, and Brocade G630 click a slot number under 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 (DLS) feature for dynamic routing path selection. When this policy is in force, DLS is always enabled and cannot be disabled.

When the port-based routing 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 or E_Port balancing, refer to the **dlisset** command in the *Brocade Fabric OS Command Reference*.

To configure dynamic load sharing, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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 on the following platforms:

- Brocade 6505
- Brocade 6510
- Brocade 6520
- Brocade 7840 on FC ports
- Brocade G610
- Brocade G620
- Brocade G630
- Brocade DCX 8510-4 Director
- Brocade DCX 8510-8 Director
- Brocade X6-4 Director
- Brocade X6-8 Director

NOTE

For the Brocade FX8-24 Extension Blade, the lossless feature is supported only on FC ports.

You can enable this lossless feature from Web Tools. If you try to enable it when DLS is off, an error message displays.

To enable or disable lossless DLS, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Routing** tab.
3. Select **On** in the **Loss Less** area to enable the mode, or select **Off** to disable dynamic load sharing.

When the exchange-based routing policy is in effect, the lossless DLS buttons display on the **Routing** tab.

4. Click **Apply**, and then 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 *Brocade Fabric OS Administration Guide*.

NOTE

Enabling in-order delivery can cause a delay in the establishment of a new path when a topology change occurs, and therefore it should be used with care.

To specify frame order delivery, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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 inter-switch 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 ISLs along each path. FSPF defines the path with the 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.

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 45.
2. Select the **Routing** tab.
Displaying routing path information is switch-specific. For the Brocade 6505, Brocade 6510, Brocade 6520, Brocade 7800, Brocade 7840, Brocade G610, Brocade G620, and Brocade G620 click the slot number of the logical switch under **Link Cost** in the navigation tree.
3. Double-click in the row of the **Cost** column that corresponds to the appropriate port.
4. Enter the link cost. Valid values for link cost are from 1 through 65534. Setting the value to 0 sets the link cost to the default value for that port.
5. Click **Apply**.

E_Port balance priority

E_port balance priority allows you to balance the E_port load.

You can enable the E_port balance priority feature from Web Tools. When you enable the E_port balance priority feature, the E_Port load will be even across all E_Ports with the same domain during the topology change. You can select **Rebalance** or **Rebalance ALL** to re-balance the E_Port load on a particular logical switch or on all logical switches, without waiting for a topology change to occur.

The E_port balance priority is supported on the following platforms.

- Brocade DCX 8510-4 Director

- Brocade DCX 8510-8 Director
- Brocade X6-4 Director
- Brocade X6-8 Director
- Brocade 6520
- Brocade 7800
- Brocade 7840
- Brocade G610
- Brocade G620
- Brocade G630

NOTE

- The **E_port Balance Priority** will be available on all platforms, but will be grayed out for unsupported platforms.
- When the **Dynamic Load Sharing (DLS)** is disabled, **Lossless Dynamic Load Sharing (DLS)** is not supported and the **E_port Balance Priority** feature is also disabled; but the **E_port Balance Priority** can be enabled even if the DLS is in the **Off** state.

To enable or disable the E_port balance priority, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Routing** tab.
3. Select **On** in the **E-port Balance Priority** area to enable E_Port load balancing, or select **Off** to disable E_Port load balancing.
Clicking the **Rebalance** button will perform E_Port balancing on the current logical switch only, and clicking the **Rebalance All** button will perform E_Port balancing on all logical switches available.
4. Click **Apply**, and then click **OK**.

Configuring Standard Security Features

• User-defined accounts.....	187
• User-defined roles.....	193
• Access control list policy configuration.....	197
• Fabric-Wide Consistency Policy configuration.....	200
• Authentication policy configuration.....	200
• SNMP configuration.....	203
• RADIUS management.....	205
• Active Directory service management.....	207
• TACACS+ management.....	209
• IPsec concepts.....	210
• IPsec over management ports.....	215
• Establishing authentication policies for HBAs.....	220

User-defined accounts

In addition to the default accounts (admin and user), Fabric OS 7.0.0 and later support 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 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.

Access rights for any user session are determined by the user's role-based access rights. Refer to [Introducing Web Tools](#) on page 15 for additional information about Role-Based Access Control (RBAC).

The **User** tab of the **Switch Administration** window [Opening the Switch Administration window](#) on page 45 displays account information. You can create and manage accounts depending on your role. The roles and permissions are listed in the following table.

TABLE 18 User role and permissions

Role	Permissions
admin	Create and manage all predefined accounts (except the root account) and user-defined accounts.
operator	Change your own password but cannot create, modify, or view predefined or user-defined accounts.
securityadmin	Create and manage all security roles.
switchadmin	Change your own password but cannot create, modify, or view predefined or user-defined accounts.
zoneadmin	Change your own password but cannot create, modify, or view predefined or user-defined accounts.
fabricadmin	Change your own password but cannot create, modify, or view predefined or user-defined accounts.
basicswitchadmin	Change your own password but cannot create, modify, or view predefined or user-defined accounts.
user	Change your own password but cannot create, modify, or view predefined or user-defined accounts.

NOTE

Web Tools displays an error message when a user tries to access the root account.

Virtual Fabrics considerations

If no home logical fabric ID is specified for a user, the system provides a default home ID of 128.

Viewing user account information

To view user account information, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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

To create user-defined accounts, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **User** tab.

3. Click **Add**.

The **Add User Account** dialog box displays. For switches that support Virtual Fabrics, refer to the following figure.

FIGURE 65 Add User Account dialog box (VF)

Logical Fabric ID	User Role
1	No Access
2	No Access
3	No Access
4	No Access
5	No Access
6	No Access
7	No Access

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 period (.), 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 27 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 period (.), 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 192 for more information.)

9. Retype the password in the **Confirm Password** field for confirmation.
10. Check the available Virtual Fabrics that you can access.

For Virtual Fabrics, all logical fabrics IDs (1-128) are displayed, even if they have not all been 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.

The default home logical fabric ID is 128.

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

To delete user-defined accounts, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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

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.

To change user account parameters, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **User** tab.

3. Select the account to modify.

NOTE

You cannot modify the default root account, even if you are logged in as root.

4. Click **Modify**.

The **Modify User Account** dialog box displays.

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. Only the root account can disable itself. If you disable an account, all active CLI sessions for that account are logged out.

8. Select a home domain for the user from the **Home** menu.

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. Click **OK** and click **Apply** to apply your changes.

Maintaining passwords

When a password expires, the next time that a 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 procedure in [Unlocking a password](#) on page 193.

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. But you cannot change the root account password.

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 period, 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 192 for more information.)

NOTE

You must change your account password when Web Tools asks you to change it. You cannot proceed further unless you change your existing password.

To change the password of an account, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **User** tab.
3. Select the account to modify.

If you are logged in as a switchadmin, you can change only 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 account, 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

To set rules for passwords, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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 that 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 for which the password is locked (0-99999 minutes)
- Number of days to warn a user before password expiration (0-999)

5. Select whether to enable or disable the lockout administration features.

If you select to disable 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

To set a password as expired, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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

To unlock a password, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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 45.
2. Select the **User** tab.
3. Select an account.
4. Select **Show Role and VF**. The role mapping for that user displays.

User-defined roles

User-defined roles provide the ability to dynamically create roles on the switch. The default roles, such as Root, Admin, User, SwitchAdmin, ZoneAdmin, FabricAdmin, BasicSwitchAdmin, SecurityAdmin, and Operator, are defined by giving different permissions for different features or by restricting access to various features. The default roles cannot be edited for assigning different privileges. However, user-defined roles provide the ability to create new roles and define permissions for the RBAC classes.

Guidelines and restrictions

Follow these guidelines and restrictions when creating and configuring user-defined roles:

- In order for the user-defined role to be able to edit the Port Admin and FCR configuration, you must assign the RBAC_SwitchPortManagement and RBAC_SwitchPortConfiguration RBAC classes to the role.
- In order for the user-defined role to be able to set the fabric ID, you must assign the RBAC_FabricRouting and RBAC_SwitchConfiguration RBAC classes to the role.
- In order for the user-defined role to be able to view reports, you must assign the RBAC_SwitchManagement, RBAC_SwitchConfiguration, and RBAC_FRUManagement RBAC classes to the role.

For some functionality and operations, that need chassis-level access, the user-defined role privileges must be assigned at both the chassis level and the Logical Fabric level to have the corresponding tab enabled:

- In order for the user-defined role to have access to the **System Monitor**, which displays **CPU** and **Memory Usage** under the **Monitor** tab, you must assign the read/write RBAC_FabricWatch permission and the CHASSIS_CONTEXT context type to the Chassis Access Role.
- In order for the user-defined role to have access to the **Configure** tab, you must assign the RBAC_ConfigManagement, RBAC_SwitchConfiguration, or RBAC_Configure classes to the user-defined role, which is applied at the Logical Fabric level. Any of these three classes is sufficient.
- In order for the user-defined role to have access to the **Security Policy** tab, you must assign the RBAC_Authentication, RBAC_FabricDistribution, RBAC_Security, RBAC_IPSec, RBAC_AG, or RBAC_IPfilter classes to the user-defined role, which is applied at the Logical Fabric level. Any of these six classes is sufficient.
- In order for the user-defined role to have access to the **Switch** tab, you must assign the RBAC_SwitchConfiguration, RBAC_SwitchManagement, RBAC_FRUManagement, RBAC_AG, or RBAC_Configure classes to the user-defined role, which is applied at the Logical Fabric level. Any of these five classes is sufficient.

Creating a user-defined role

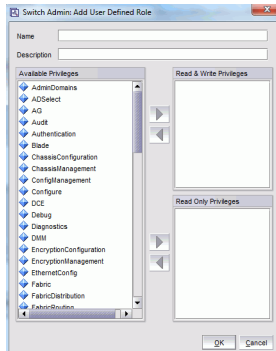
To add a user-defined role, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **User** tab.
3. Select the **Role** subtab.

- Click the **Add** button.

The **Switch Admin: Add User Defined Role** dialog box displays.

FIGURE 66 Switch Admin: Add User Defined Role dialog box



- Enter a role name in the **Name** field.
- Enter a description of the role in the **Description** field.
- To grant the role a read/write privilege, select the privilege and click the right arrow next to the **Read & Write Privileges** section.
You can select multiple privileges.
- To grant the role a read privilege, select the privilege and click the right arrow next to the **Read only Privileges** section.
You can select multiple privileges.
- To delete a privilege, select it and click the left arrow.
- Click **OK** to save your changes.

Editing a user-defined role

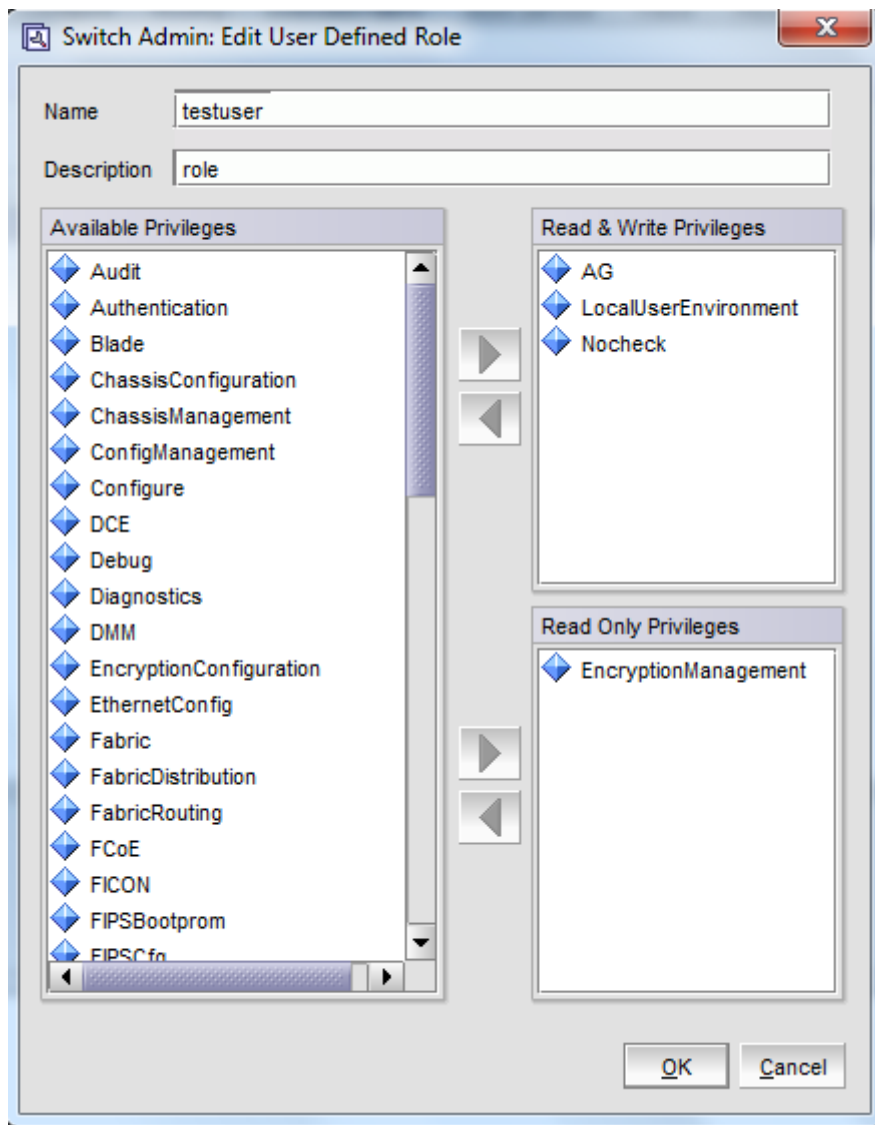
To edit a user-defined role, perform the following steps.

- Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
- Select the **User** tab.
- Select the **Role** subtab.
- Select an existing user-defined role.

- Click the **Edit** button.

The **Switch Admin: Edit User Defined Role** dialog box displays.

FIGURE 67 Switch Admin: Edit User Defined Role dialog box



- To grant the role a read/write privilege, select the privilege and click the right arrow next to the **Read & Write Privileges** section. You can select multiple privileges.
- To grant the role a read only privilege, select the privilege and click the right arrow next to the **Read only Privileges** section. You can select multiple privileges.
- To delete a privilege, select it and click the left arrow.
- Click **OK** to save your changes.

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. The SCC and DCC policy configuration in base Fabric OS is performed on a switch-local basis.

The 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.

Creating an SCC, DCC, or FCS policy

You can create the FCS policy only once.

To create an SCC, DCC, or FCS policy, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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 that 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 devices in the fabric are also listed in the tree.

11. In the **ADD Domain, Port Index** field, enter the value in 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 198 for instructions.

Editing an SCC, DCC, or FCS policy

To edit an SCC, DCC, or FCS policy, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Make sure that 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**.

The **ACL Policy Configuration** wizard is launched.

6. Select the policy type that 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 active policy.

To delete all SCC, DCC, or FCS policies, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Security Policies** tab.
3. Click **Delete All**.
A warning message displays.
4. Click **OK** to delete all 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.

To activate all SCC, DCC, or FCS policies, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Security Policies** tab.
3. Click **Activate All** to activate all 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

An SCC or DCC policy can be distributed only for a primary switch.

To distribute an SCC, DCC, or FCS policy, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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. Click **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.

To move an FCS policy switch position, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Click **Show Advanced Mode**.
3. Select the **Security Policies** tab.
4. Select the **FCS** tab.
5. Click **Move FCS Switch**.
6. Select the appropriate from and to positions.
7. Click **Apply**.
8. After you move all member switches, click **Apply** and **Close**.

Configuring the Advanced Device Security policy

The Advanced Device Security (ADS) policy allows you to restrict devices that are logged in to the fabric using a particular F_Port. When this policy is enabled, only authorized devices are allowed to log in to the fabric. This can be achieved by allowing all devices, blocking all devices, or giving access to selected devices. ADS is supported only in Access Gateway mode.

The restrictions to device login are:

- **All Access** – Allows all devices to log in to the fabric through that F_Port.
- **No Access** – Blocks all devices trying to log in to the fabric through that F_Port.
- **WWNs** – Allows only selected WWNs to log in to the fabric through that F_Port. NPIV-capable device port WWNs can also be added to the allowed list of device port WWNs for the particular F_Port.

When the ADS policy is enabled the first time, all F_Ports are set to **All Access** and all the devices are allowed to log in to the fabric. This configuration persists for subsequent logins from all devices. Existing devices that are already logged in to the fabric are not affected.

When the ADS policy is disabled, all allowed lists are cleared and all the devices are allowed to log in to the fabric.

To configure the ADS policy, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Click **Show Advanced Mode**.
3. Select the **Security Policies** tab to configure the ADS policy in Access Gateway mode.
4. Select the **ADS** option.
5. Select the **Enable ADS Policy** option.

The **Configure Advanced Device Security Port WWN** table displays.

6. Optional: Select an F_Port from the table and click the **Edit** button.

The **ADS Port WWN Configuration** dialog box displays. You can configure device port WWNs that can be allowed to log in to a particular F_Port by adding them to the **Selected WWN** list.

7. Select either **All Access** or a list of selected WWNs.
8. Optional: You can add the detached port WWN to the selected WWNs list by adding the WWN in the **detached WWN** text field and clicking **Add**.
9. Optional: For a selected F_Port, if you select the **Show device WWN connected to this port** check box of the **ADS Port WWN Configuration** dialog box, only connected devices are listed in the **Available WWNs** list. When you clear the check box, all connected device port WWNs and detached WWNs added to the AG are listed in the **Available WWNs** list.

Fabric-Wide Consistency Policy configuration

Fabric-Wide Consistency Policy (FWCP) configures the fabric wide consistency behavior of distributable ACL policies. The FWCP 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, the policies are managed on a per-switch basis.

To set the fabric-wide consistency policy for an SCC, DCC or FCS policy, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Security Policies** tab.
3. Click **FWCP**.
4. Select one of the following consistency behaviors for the required policy type (SCC, DCC, FCS).
 - Absent
 - Tolerant
 - Strict

NOTE

You can change the consistency behavior of an 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 you can then distribute the authentication policy to other switches in the fabric. You can also set shared secret keys.

Configuring authentication policies for E_Ports

To configure authentication policies for E_Ports, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Security Policies** tab.
3. Select **Authentication**.
4. In the **Authentication Type** field, select **FCAP** or **DHCHAP**.
5. Select a switch authentication policy mode listed in the following table.

TABLE 19 Switch Authentication Policy modes

Policy mode	Description
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, SHA 256, 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

To configure authentication policies for F_Ports, perform the following steps.

1. Open the **Switch Administration** window and click **Show Advanced Mode**, if not selected.
2. Select the **Security Policies** tab.
3. Select **Authentication**.
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 selected switches accept the distribution. Only the policy mode is distributed to the selected switches. The switch that initiates the distribution must accept distribution.

To distribute authentication policies, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Security Policies** tab.
3. Select **Authentication**.
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.

To re-authenticate policies, perform the following steps.

1. Click a port in the **Switch View** to open the **Port Admin** tab.
The **Port Admin** tab displays with the port selected.
2. Select **Re-Authenticate** (active only for F_Ports and E_Ports) from the **Actions** list.
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.

To set a shared secret key pair, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Security Policies** tab.
3. Select **Authentication**.
4. Select the **Shared Secret Keys** subtab.
5. Click **Add**.
The **Add Shared Secret Keys** dialog box displays.
6. Enter the switch or HBA 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**.

NOTE

In the **Security Policies** subtabs (**Authentication Policies** and **Shared secret keys**), user configurations will be retained in the inner tab navigation. **Apply** will wrap the configuration in both tabs.

Modifying a shared secret key pair

You can edit and modify the secret key pairs by switch.

To modify a shared secret pair, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Security Policies** tab.
3. Select **Authentication**.
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 authentication is required when a switch logs in to a fabric.

To set the Switch Policy Authentication mode, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **Security Policies** tab.
3. Select **Authentication**.
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, access control, and system group 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 *Brocade Fabric OS Command Reference*.

Setting SNMP trap levels

To set SNMP trap levels, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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 that 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

To change the systemGroup configuration parameters, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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

To set SNMPv1 configuration parameters, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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.

To set SNMPv3 configuration parameters, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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 inform requests are 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.

To change the access control configuration, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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 (AAA) service. 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:

- **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** - The switch user login database is checked only if the physical connection to the RADIUS server fails.
- **None** - The switch user login database is never checked. Only a RADIUS server can be used for authentication.

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 RADIUS authorization fails**
- **Switch Database when RADIUS times out**

When RADIUS login fails, even though the RADIUS server is available, the additional service allows you the option to use the switch database as a 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.

To enable or disable RADIUS, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **AAA Service** tab.
3. To enable RADIUS, select **RADIUS** from the **Primary AAA Service** menu.

4. Select **None**, **Switch Database when RADIUS Login Fails**, or **Switch Database when RADIUS Login Times out** 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. The configuration is saved in a configuration upload, and it 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.

To configure RADIUS, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **AAA Service** tab.
3. Click **Add**.

The **RADIUS/ADLDAP/TACACS+ 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. Select **RADIUS** from **Server Type**.
5. Enter the RADIUS server name as a valid IP address (in either IPv4 or IPv6 format) or a Dynamic Name Server string.
Each RADIUS server must have a unique IP address or DNS name for the RADIUS server.
6. Enter the port number.
7. Enter the secret string.
8. Enter the timeout time in minutes.
9. 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.

10. Click **OK** to return to the **AAA Service** tab.
11. Click **Apply**.

Modifying the RADIUS server

To change the parameters of a RADIUS server that is already configured, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **AAA Service** tab.
3. Select a RADIUS server from the **RADIUS Configuration** list.
4. Click **Modify**.

The **RADIUS/ADLDAP/TACACS+ 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 in which they are listed, starting from the top of the list and moving to the bottom.

To modify the RADIUS server order, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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

To remove a RADIUS server, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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 user profiles. Active Directory provides user authentication and authorization using LDAP as the authentication protocol. Active Directory provides better security while using a remote authentication mechanism.

You can add, remove, and modify the settings of the Active Directory server.

Enabling an Active Directory service

To add a new Active Directory server, you must provide the server IP address, port number, secret string, timeout value, and you must choose 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.

To enable an Active Directory service, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **AAA Service** tab.
3. To enable an Active Directory service, select **Active Directory** from the **Primary AAA Service** menu.
4. Select **None**, **Switch Database when Active Directory authentication failes**, or **Switch Database when Active Directory times out** from the **Secondary AAA Service** menu.

NOTE

To disable the **Active Directory** service, select **Switch Database** from the **Primary AAA Service** menu, and select **None** from the **Secondary AAA Service** menu.

5. Click **Apply**.

Modifying an Active Directory service

To change the parameters of an Active Directory service that is already configured, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **AAA Service** tab.
3. Select a server from the **ADLDAP Configuration** list.
4. Click **Modify**.

The **RADIUS/ADLDAP/TACACS+ 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 an Active Directory service

To remove an Active Directory server, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
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.

TACACS+ management

TACACS+ provides user authentication and authorization using TACACS+ as the authentication protocol. You can add, remove, and modify the settings of the TACACS+ server.

Enabling and disabling TACACS+

At least one TACACS+ server must be configured before you can enable TACACS+.

To enable or disable TACACS+, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **AAA Service** tab.
3. To enable TACACS+, select **TACACS+** from the **Primary AAA Service** menu.
4. Select **None**, **Switch Database when TACACS+ Login Failed**, or **Switch Database when TACACS+ Login Timeout** from the **Secondary AAA Service** menu.

NOTE

To disable TACACS+, select **Switch Database** from the **Primary AAA Service** menu and select **None** from the **Secondary AAA Service** menu.

5. Click **Apply**.

Configuring TACACS+

To enable TACACS+, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **AAA Service** tab.
3. Click **Add** to configure the TACACS+ service.
4. Select **TACACS+** from **Server Type**.
5. Enter the **Server**, **Port**, **Timeout(s)**, and **Secret String** details in the appropriate fields.
6. Select **CHAP** or **PAP** from the **Authentication** list.
7. Click **OK**.

The server details display in the **TACACS+ Configuration** list.

8. Click **Apply**.

Modifying TACACS+

To change the parameters of a TACACS+ service that is already configured, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **AAA Service** tab.
3. Select a server from the **TACACS+ Configuration** list.
4. Click **Modify**.

The **RADIUS/ADLDAP/TACACS+ Configuration** dialog box displays.

5. Enter new values for the fields that you want to modify.
6. Click **OK** to return to the **AAA Service** tab.
7. Click **Apply**.

Removing TACACS+

To remove a TACACS+ server, perform the following steps.

1. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
2. Select the **AAA Service** tab.
3. Select a server from the **TACACS+ 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 Service** 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 the following table.

TABLE 20 Relevant RFCs

RFC number	Title
RFC 3041	Privacy Extensions for Stateless Address Autoconfiguration in IPv6
RFC 3971	Secure Neighbor Discovery
RFC 3972	Cryptographically Generated Addresses
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 4306	Internet Key Exchange Version 2 (IKEv2) Protocol
RFC 4307	Cryptographic Algorithms for Internet Key Exchange Version 2 (IKEv2)

TABLE 20 Relevant RFCs (continued)

RFC number	Title
RFC 4309	Using Advanced Encryption Standard (AES) CCM Mode with IPsec Encapsulating Security Payload (ESP)
RFC 4869	Suite B Cryptographic Suites for IPsec


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 that specifies the addresses of the tunnel endpoints. IPsec is implemented between tunnel endpoints. IPsec is transparent to the actual endpoints within the IP header in the original packet.

The following figure provides a basic visual comparison of how transport mode and tunnel mode modify an IP datagram.

FIGURE 68 Transport mode and tunnel mode comparison



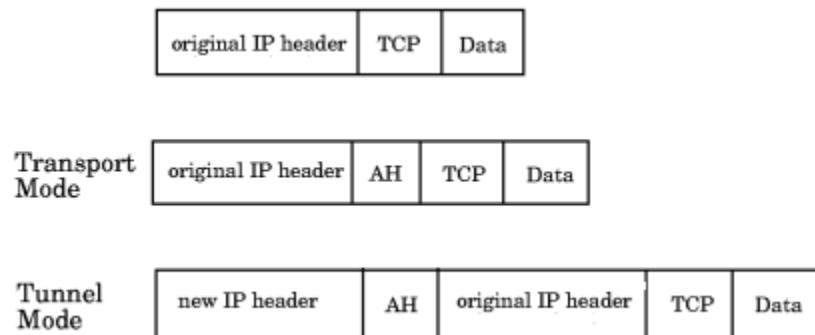
IPsec header options

IPsec adds headers to an IP datagram to enable authentication and privacy. There are two options:

- Authentication Header (AH)
- Encapsulating Security Payload (ESP)

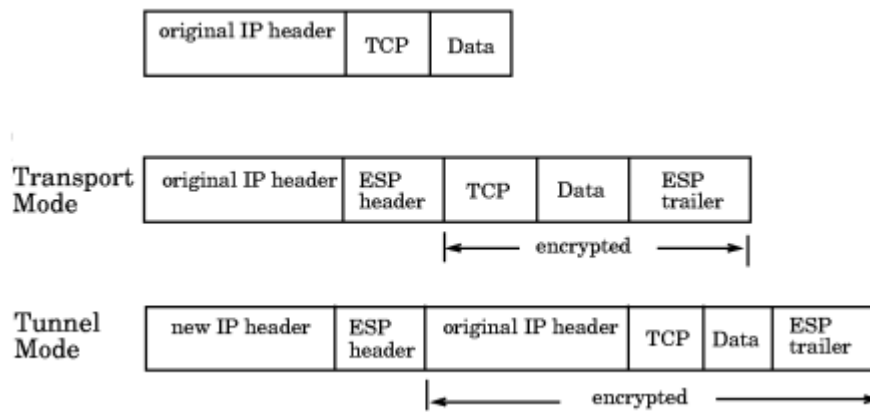
Authentication Header

AH can be used to authenticate a data stream, but it does not provide the 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 69](#).

FIGURE 69 AH header in transport mode and tunnel mode

Encapsulating Security Payload

ESP provides authentication, and it also provides privacy by encrypting the IP datagram. The use of an ESP header is similar to the use of an 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 the following figure.

FIGURE 70 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 it will 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 endpoints of an IP connection and to determine the 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. The following table lists the available encryption algorithms. A brief description is provided. If you need further information, refer to the RFC.

TABLE 21 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 (HMACs) check data integrity through a mathematical calculation on a message using a hash algorithm combined with a shared, secret key. The following table 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 it 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 22 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 (5 minutes), more than 12 keys may be generated to complete the communication.

The SA lifetime limits the length of time for which a key is used before it is replaced by a new key, thus limiting the amount of time for which 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 standard (DSS), and RSA digital signature:

- A preshared key (PSK) is a shared secret that is shared between two parties over a secure channel before use. 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 is also confirmed because the public key can decrypt only data that has been encrypted with the corresponding private key.

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 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.

Enabling the Ethernet IPsec policies

To access the **Ethernet IPsec Policies** dialog box, perform the following steps.

1. Open the **Switch Administration** window.
2. Select **Show Advanced Mode**.
3. Select the **Security Policies** tab.
4. Select **Ethernet IPsec**.

The **Ethernet IPsec Policies** dialog box 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 *Brocade Fabric OS Administration Guide* for details on IKE functionality.

To establish an IKE policy, perform the following steps.

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 an 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 an 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.

To create a security association, perform the following steps.

1. Select the **IPsec** tab.
The **IPsec Policies** windows 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 it 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 (SAs). 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

To create an SA proposal, perform the following steps.

1. Select the **SA Proposal** tab on the **IPsec Policies** window.
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 process:

- 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 that a data packet is processed using IPsec encryption, IKE authentication, or both, using encapsulation security protocol (ESP) processing or authentication header (AH) protocol processing.

To add an IPsec transform policy, perform the following steps.

1. Select the **Transforms** tab.

The **Transforms** window 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 must 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 unidirectional. 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** window 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 they 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. Use the **Protocol Name** selector to select a specific protocol.
9. Click **OK**.

Manually creating an SA

Part of manually creating a security association (SA) is to select an IPsec protection type. The choices are discard, bypass, and process:

- 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 that a data packet is processed using IPsec encryption, IKE authentication, or both, using encapsulation security protocol (ESP) processing or authentication header (AH) protocol processing.

To manually create an SA, perform the following steps.

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 they 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 **IPsec Mode**.

The choices are **Transport** or **Tunnel**. Refer to [Transport mode and tunnel mode](#) on page 211 if you are unfamiliar with Transport and Tunnel modes.

9. Select **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.

To edit an IKE or IPsec policy, perform the following steps.

1. Open the **Switch Administration** window.
2. Click **Show Advanced Mode**.
3. Select the **Security Policies** tab.
4. Select **Ethernet IKE** or **Ethernet IPsec**.
5. Select the policy that 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.

To delete an IKE or IPsec policy, perform the following steps.

1. Open the **Switch Administration** window.
2. Select **Show Advanced Mode**.
3. Select the **Security Policies** tab.
4. Select **Ethernet IKE** or **Ethernet IPsec**.
5. Select the policy or policies that 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

To establish and enable authentication policies for HBAs as the login to a fabric, perform the following steps.

1. Open the **Switch Administration** window.
2. Click **Show Advanced Mode**.
3. Select **Authentication** under the **Security Policies** tab.
The **Authentication Policy Settings** window displays.
4. Under **Configure Authentication Policy**, do the following:
 1. Select the **Authentication Type**. The choices are FCAP, DHCHAP, or both.
 2. Select the **Switch Authentication Policy Mode**. The choices are Passive, Active, On, or Off.
 3. Select the **Hash Type** used. The choices are SHA 1, SHA 256, MD5, SHA 1 and MD5, or all.
 4. Select the **DH-Group Type**. The choices are 0, 1, 2, 3, 4; 0 (DH Null option), 1 (1024-bit key), 2 (1280-bit key), 3 (1536-bit key), or 4 (2048-bit key).
 5. Use the **Device Authentication Policy Mode** selector to set the desired mode. The choices are On, Off, or Passive.
 6. Click **Apply**.
5. If your authentication method uses a shared secret, select the **Shared Secret Keys** tab.
The **Shared Secret Keys** window displays.
6. Select **Add**.
The **Add Shared Secret Keys** dialog box displays.
7. Enter or browse to select the switch/HBA WWN, name, or domain ID in the **Switch or HBA WWN/Name/Domain ID** field.
8. 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.
9. Enter the shared secret for the switch in the **Local Shared Secret** and **Confirm Local Shared Secret** fields.

10. Click **Add**.

An entry is added in the **Switch WWN** box.

11. Click **OK**.

12. Add more shared secrets, if needed.

Administering FICON CUP Fabrics

• FICON CUP fabrics overview.....	223
• Enabling port-based routing.....	224
• Enabling or disabling FICON Management Server mode.....	224
• FMS parameter overview.....	225
• Configuring FMS mode parameters.....	226
• Displaying code page information.....	226
• Viewing the control device state.....	226
• Allow / Prohibit Matrix configuration.....	227
• CUP logical path configuration.....	230
• Link Incident Registered Recipient configuration.....	231
• Displaying Request Node Identification Data	232

FICON CUP fabrics overview

Control Unit Port (CUP) is a protocol for managing FICON directors. Host-based management programs manage the switches using the 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 Network Advisor. (Refer to the *Brocade Network Advisor SAN User Manual* for more information about the Brocade Network Advisor.) The director allows control to be shared between host-based management programs and director consoles.

NOTE

While enabling FMS mode with online devices connected to ports with addresses of 0xFE or 0xFF, the following error message displays: `FMS mode enable failed due to port(s) with areas 0xFE or 0xFF is (are) connected to device(s) . You must disable the ports or remove the online devices from those ports that are mapped to the 0xFE or 0xFF address.`

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.
- Install a FICON CUP license on the Brocade switch.
- 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.

NOTE

If the switch does not have the FICON_CUP license installed, Web Tools prevents the enabling of FMS mode and displays the following error message: `Enabling FMS mode requires FICON CUP license installed on the switch. Contact your preferred storage vendor for more details.`

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 attached FICON devices must have a 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.

To enable port-based routing, perform the following steps.

1. Select a switch with attached FICON devices from the **Fabric Tree**.
2. Open the **Switch Administration** window as described in [Opening the Switch Administration window](#) on page 45.
3. Click the **FICON CUP** tab.
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.

NOTE

While enabling FMS mode with online devices connected to FE or FF, the following error will be shown: `FMS mode enable failed due to port(s) with areas 0xFE or 0xFF is (are) connected to device(s).`

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 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 60.

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.

To enable or disable FICON Management Server, perform the following steps.

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 45.
3. Click **Show Advanced Mode**.
4. Select the **FICON CUP** tab.

The FICON CUP tabbed page displays the **FICON Management Server** subpage. 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.

NOTE

High Integrity Fabric (HIF) must be enabled to enable the FMS mode.

FMS parameter overview

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 the following table.

TABLE 23 FMS mode parameters

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 it can be reset by Web Tools.
Active=Saved Mode	<p>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 activates its contents as the default configuration.</p> <p>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.</p> <p>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 it can be reset by Web Tools.</p> <p>Note: When FMS mode is enabled and the Active=Saved Mode parameter is disabled, you can enable and disable ports, but the setting is <i>not</i> persistent. When the Active=Saved Mode parameter is enabled, you can enable and disable ports and the setting <i>is</i> persistent.</p>
Alternate Control Prohibited	<p>Determines whether alternate managers are allowed to modify port connectivity.</p> <p>Enabling this mode prohibits alternate manager control of port connectivity; otherwise, alternate managers can manage port connectivity.</p> <p>This parameter is set as enabled by the hardware after system installation, and it can be reset by Web Tools.</p>
User Alert Mode	<p>Controls director console behavior for alerts.</p> <p>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.</p> <p>This parameter is always read-only in Web Tools. Each time that the switch is powered on, the parameter is reset to disabled.</p>
Director Clock Alert Mode	<p>Controls behavior for attempts to set the switch timestamp clock through the director console.</p> <p>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.</p> <p>This parameter is set as disabled by the hardware after system installation, and it can be reset by Web Tools.</p>
Host Control Prohibited	<p>Determines whether host programming allows modifying port connectivity.</p> <p>Enabling this mode prohibits host programming control of port connectivity; otherwise, host programming can manage port connectivity.</p> <p>This parameter is set as disabled by the hardware after system installation, and it can be reset by Web Tools.</p>

Configuring FMS mode parameters

To configure FMS mode parameters, perform the following steps.

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 45.
3. Select the **FICON CUP** tab.

The **FICON CUP** page displays the **FICON Management Server** subpage. 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.

To display code page information, perform the following steps.

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 45.
3. Select the **FICON CUP** tab.

The **FICON CUP** page displays the **FICON Management Server** subpage 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 state or a switched state. When it is neutral, the control device accepts commands from any channel that has established a logical 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 the 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.

If the following switch parameters are read or modified, a FICON CUP busy error can be generated:

- Mode Register
- Port Names (also called Port Address Name)
- Allow/Prohibit Matrix and Port Connectivity Attributes
- Switch enable/disable
- Switch name change

To access the FICON CUP tab, perform the following steps.

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 45.
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.

Allow / Prohibit Matrix configuration

In the Allow / Prohibit Matrix subpanel, you can manage the configuration files and the active configuration. All 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 Allow / Prohibit Matrix 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 a 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 [Figure 72](#) on page 229) 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 Allow / Prohibit Matrix configurations

To display a list of Allow / Prohibit Matrix configurations, perform the following steps.

1. Select a FICON-enabled switch from the **Fabric Tree**.
2. Select **Configure > Switch Admin**.
3. Click **Show Advanced Mode** to see all available tabs and options.
4. Select the **FICON CUP** tab.

The **FICON CUP** tabbed page displays the **FICON Management Server** subtabbed page in front. All attributes on this page are read-only until FMS mode is enabled.

5. Click the **Allow / Prohibit Matrix** subtab.

Modifying Allow / Prohibit Matrix configurations

In the **Allow / Prohibit Matrix Configuration** dialog box, swapped ports are indicated with the "(Swapped)" label.

FIGURE 71 Edit Allow / Prohibit Matrix dialog box swapped label

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)

To create a new Allow / Prohibit Matrix configuration or to edit an existing configuration, perform the following steps.

1. Display the Allow / Prohibit Matrix configuration list.
2. You can either create a new configuration or edit an existing configuration:
 - To create a new configuration, click **New**.

The **Allow / Prohibit Matrix Configuration** dialog box displays all ports and port names on the selected switch (similar to the dialog box shown in the following figure). 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 **Allow / Prohibit Matrix Configuration** dialog box displays the content of the selected configuration from the switch in a table format (as shown in the following figure).

3. Optional: Select the check box corresponding to a port that you want to block on the **Block** column. Repeat this step for all ports that you want to block. Select the **Block All** check box to block all ports.
4. Optional: Select the check box corresponding to a port that you want to prohibit on the **Prohibit** column. Repeat this step for all ports that you want to prohibit. Select the **Prohibit All** check box to prohibit all ports.

The cells in the matrix are updated with crossed-circle icons to identify prohibited ports.

FE and FF port are not shown in the **Allow / Prohibit Matrix Configuration** dialog box. The FE and FF port state displays only in the **Port Admin** page.

5. Optional: Click the individual cells corresponding to the combination of ports that you want to prohibit. You cannot prohibit a port to itself.

If you prohibit an 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 an Allow / Prohibit Matrix configuration](#) on page 230.
 - 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 72 Allow / Prohibit Matrix Configuration dialog box

Create Allow / Prohibit Matrix Configuration for switch WT.Elara.48

Code Page
Language used to exchange information with Host Programming:
(EBCDIC) USA/Canada -- 00037

Settings
☐ Prohibit All ☐ Block All

Port	Port Name	Block	Prohibit	Port Area	00	01	02	03	04	05	06	07	08
0		<input type="checkbox"/>	<input type="checkbox"/>	00									
1		<input type="checkbox"/>	<input type="checkbox"/>	01									
2		<input type="checkbox"/>	<input type="checkbox"/>	02									
3		<input type="checkbox"/>	<input type="checkbox"/>	03									
4		<input type="checkbox"/>	<input type="checkbox"/>	04									
5		<input type="checkbox"/>	<input type="checkbox"/>	05									
6		<input type="checkbox"/>	<input type="checkbox"/>	06									
7		<input type="checkbox"/>	<input type="checkbox"/>	07									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	08									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	09									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	0A									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	0B									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	0C									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	0D									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	0E									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	0F									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	10									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	11									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	12									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	13									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	14									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	15									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	16									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	17									
Unavailable		<input type="checkbox"/>	<input type="checkbox"/>	18									

Activate Save As... Cancel

Activating an Allow / Prohibit Matrix configuration

When you activate a saved Allow / Prohibit Matrix configuration on the switch, the preceding configuration (currently activated) is overwritten.

To activate an Allow / Prohibit Matrix configuration, perform the following steps.

1. Open the Allow / Prohibit Matrix configuration list.
2. Select the saved configuration from the list.
3. Click **Activate**.

The **Activate Allow / Prohibit Matrix 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 an Allow / Prohibit Matrix configuration

To copy an Allow / Prohibit Matrix configuration to a new configuration, perform the following steps.

1. Display the Allow / Prohibit Matrix configuration list.
2. Select a saved configuration or the active configuration from the list.
3. Click **Copy**.

The **Allow / Prohibit Matrix 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 an Allow / Prohibit Matrix configuration

To delete a saved Allow / Prohibit Matrix configuration.

1. Display the Allow / Prohibit Matrix configuration list.
2. Select the saved configuration from the list.
3. Click **Delete**.

The **Delete Allow / Prohibit Matrix 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

To display a list of CUP logical path configurations, perform the following steps.

1. Select a FICON-enabled switch from the **Fabric Tree**.
2. Select **Configure** > **Switch Admin**.
3. Click **Show Advanced Mode** to see all 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

To configure a CUP logical path, perform the following steps.

1. Select a FICON-enabled switch from the **Fabric Tree**.
2. Select **Configure** > **Switch Admin**.
3. Click **Show Advanced Mode** to see all 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 (LIR) on the source N_Port. The LIRR database is stored on the switch.

Viewing Link Incident Registered Recipient configurations

To display a list of Link Incident Registered Recipient (LIRR) configurations.

1. Select a FICON-enabled switch from the **Fabric Tree**.
2. Select **Configure** > **Switch Admin**.
3. Click **Show Advanced Mode** to see all 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

To configure the Link Incident Registered Recipients (LIRRs), perform the following steps.

1. Select a FICON-enabled switch from the **Fabric Tree**.
2. Select **Configure** > **Switch Admin**.
3. Click **Show Advanced Mode** to see all 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. Select a port from the list.
7. Click **Set Current**.
8. Click **Close**.
9. Optional: Reset the selected port using the **reset** button.

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 as shown in the following figure.

FIGURE 73 Switch RNID information

The screenshot shows the Brocade Fabric OS Web Tools interface for switch 128. The 'Name Server' tab is selected, displaying a table of port information. The 'RNID' section is highlighted, showing details for the attached FICON device.

Port	Capability	Speed	Link	Flow Control	Flow Control Mode	Flow Control Priority	Flow Control Priority Mode	Flow Control Priority Queue	Flow Control Priority Queue Mode	Flow Control Priority Queue Size	Flow Control Priority Queue Weight	Flow Control Priority Queue Delay	Flow Control Priority Queue Jitter	Flow Control Priority Queue Loss	Flow Control Priority Queue Drop	Flow Control Priority Queue Error	Flow Control Priority Queue Reset	Flow Control Priority Queue Status	Flow Control Priority Queue Type	Flow Control Priority Queue Value	Flow Control Priority Queue Weight	Flow Control Priority Queue Delay	Flow Control Priority Queue Jitter	Flow Control Priority Queue Loss	Flow Control Priority Queue Drop	Flow Control Priority Queue Error	Flow Control Priority Queue Reset	Flow Control Priority Queue Status	Flow Control Priority Queue Type	Flow Control Priority Queue Value
1	FICON	10 G	Up	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	
2	FICON	10 G	Up	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On		
3	FICON	10 G	Up	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On		
4	FICON	10 G	Up	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On		
5	FICON	10 G	Up	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On		
6	FICON	10 G	Up	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On		
7	FICON	10 G	Up	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On		
8	FICON	10 G	Up	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On	On		

The 'RNID' section displays the following information:

Field	Value
Type	SLKWRM
Model	X64
Tag	04ff
Sequence number	0DZH0339L025
Insistent Domain ID Mode	Disabled
Manufacturer	BRD
Manufacturer Plant	CA

RNID information for attached FICON devices and channel paths displays on the **Name Server** view. To view this information, click the **Name Server** tab. 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 rows:

- Type
- Model
- Tag
- Sequence Number
- Insistent Domain ID Mode
- Manufacturer
- Manufacturer Plant

Web Tools limitations

- General Web Tools limitations.....235

General Web Tools limitations

The following table lists general Web Tools limitations that apply to all browsers and switch platforms.

TABLE 24 Web Tools limitations

Area	Details
Blade Failure	If a blade fails on a bladed switch platform, 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	<p>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 (by normal HA failover, restart, or other causes).</p> <p>Workaround: If Switch Admin, HA, Name Server, or Zone Admin hang, close this window, and relaunch the module.</p>
Browser	<p>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.</p> <p>Workaround: If one of the Web Tools modules is hanging, close the current Web Tools and associated windows and relaunch Web Tools.</p>
Browser	When you launch 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 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.
Firmware Download	<p>There are multiple phases to firmware download and activation. When Web Tools reports that the firmware download completed successfully, this indicates that a basic sanity check, package retrieval, package unloading, and verification were successful.</p> <p>Web Tools forces a full package install.</p> <p>A restart is required to activate the newly downloaded firmware. This restart is done automatically; however, although the Web Tools windows continue to display during the restart, they are not available. Wait approximately 10 minutes to ensure that all application windows are restored. If Web Tools fails to respond after 20 minutes, you might need to close all Web Tools application windows and restart them or contact your system administrator for network assistance.</p> <p>The Web Tools loss of network connectivity during a failover or restart (initiated through firmwareDownload) varies for different configurations:</p> <p>Brocade G610: Loss of network connectivity is up to 2 minutes if the power-on self-test (POST) is disabled. If POST is enabled, loss of network connectivity can exceed 2 minutes.</p> <p>Brocade G620: Loss of network connectivity is up to 5 minutes if POST is disabled. If POST is enabled, loss of network connectivity can exceed 5 minutes.</p> <p>Brocade G630: Loss of network connectivity is up to 3 minutes if POST is disabled. If POST is enabled, loss of network connectivity can exceed 3 minutes.</p> <p>Brocade 6510, Brocade 6520, Brocade 7800, and Brocade 7840: Loss of network connectivity is up to 1 minute if POST is disabled. If POST is enabled, loss of network connectivity can exceed 1 minute.</p>

TABLE 24 Web Tools limitations (continued)

Area	Details
Firmware Downgrade	<p>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.</p> <p>After upgrading or downgrading the firmware, delete the application cache files.</p>
HTTP Timeout	<p>Occasionally, you might see the following message when you try to get data from a switch or to send a request to the switch:</p> <p>Failed to get switch response. Please verify the status of your last operation and try again if necessary.</p> <p>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.</p> <p>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 that you were working on to retrieve related data. If your request did not get through to the switch, resubmit it. Performing a refresh from Web Tools retrieves a copy of switch data at that moment; the data that you entered can be lost if it had not already been committed to the switch.</p>
In-band Management Support	Fabric OS 7.3.0 and later support Web Tools, SNMP polling, and SNMP traps only in IPv4 on the Brocade 7800 and FX8-24.
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 19.
Java Plug-in	<p>If you have a Web Tools session open and you open a second session using the File > New browser menu, unexpected behavior of the original Web Tools session results.</p> <p>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.</p> <p>Workaround: Open two independent browser sessions.</p>
Loss of Connection	<p>Occasionally, you might see the following message when you try to retrieve data from the switch or send a request to the switch:</p> <p>Switch Status Checking</p> <p>The switch is not currently accessible.</p> <p>The dialog box title may vary because it indicates which module is having the problem.</p> <p>This message 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:</p> <p>"You will have to resubmit your request after closing this message."</p> <p>If the temporary switch connection loss is caused by a switch hot code load or other similar operation, the Switch Explorer that you are currently running can be downloaded from a different firmware version than the new one. In this case, the following message displays:</p> <p>"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."</p> <p>You need to close Switch Explorer and relaunch Web Tools to reopen the connection.</p> <p>Launching Web Tools takes around 5 minutes when the machine where the client is running does not have access to the internet. This is due to the certificate revocation check for the web start applications.</p> <p>Workaround: Go to Java Control Panel > Advanced and select Do no check (not recommended) under Perform certificate revocation checks on to skip the certificate revocation check for web start applications.</p>
Non-FIPS Secure Mode HTTPS	HTTPS supports only the TLSv1 and SSLv3 protocols with the !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 10 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:

TABLE 24 Web Tools limitations (continued)

Area	Details
	<p>java.lang.OutOfMemoryError: Java heap space</p> <p>To avoid this problem, increase the default heap size in the Java Control Panel. Refer to Java plug-in configuration on page 21 for instructions.</p>
Performance Monitor	<p>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.</p> <p>Workaround: You might need to use the CLI to manually delete these counters. For example, if you detect that Web Tools owned resources (using perfshoweemonitor), but you have verified that no web users are actually using them, use the perfdeleemonitor or perfcleareemonitor commands to free the resources.</p>
Performance Monitor	<p>The Switch Throughput Utilization, Switch Percent Utilization, and Port Snapshot Error graphs display the faulty/ powered-off slot node in the Y-axis of the graph.</p> <p>Workaround: Launch any port selection dialog box, and load the graphs accordingly.</p>
Refresh Option in Browsers	<p>When a window requesting a user response is pushed into the background and a refresh is requested, a fatal Internet Explorer error might occur.</p> <p>Workaround: Restart the browser.</p>
Refresh Option in Browsers	<p>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.</p> <p>Workaround: Restart the browser.</p>
Refresh Option in Browsers	<p>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.</p> <p>Workaround: Click the Switch tab, and the Switch Administration header updates.</p>
Refresh Option in Browsers	<p>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 fast-booted.</p> <p>Workaround: In order to display the correct switch name in the CLI prompt after a switch name update using Web Tools or SNMP, fast-boot the switch.</p>
Refresh Option in Browsers	<p>After a switch is enabled or disabled, you must wait at least 25 to 30 seconds for the fabric to reconfigure and for FSPF route calculations to complete before requesting routing information. If accessed too early, routing information is not shown.</p> <p>Workaround: After a switch is enabled or disabled, wait at least 25 to 30 seconds before further action.</p>
Refresh Option in Browsers	<p>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.</p> <p>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.</p>
Refresh Option in Browsers	<p>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.</p> <p>Workaround: If this happens, you should close the browser window and restart it.</p>
Switch Explorer Closure	<p>If a session times out or you exit or close the 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.</p>
Switch View	<p>Occasionally, switches might display the port icons correctly, but be missing one or more control button icons.</p> <p>Workaround: Close the Switch View of the switch and reopen it.</p>
Windows Operating Systems	<p>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.</p> <p>Workaround: Verify the Ethernet connection to the switch by pinging the logical switch IP address.</p>