

Brocade Fabric OS Troubleshooting and Diagnostics User Guide, 8.2.1

Supporting Fabric OS 8.2.1

Copyright © 2018 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 Inc. and/or its subsidiaries.

Brocade, a Broadcom Inc. 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 <https://www.broadcom.com/support/fibre-channel-networking/tools/oscd>.

Contents

Introduction.....	7
About This Guide.....	7
What's new in this document for Fabric OS 8.2.1.....	7
Supported Hardware and Software.....	7
Brocade Gen 5 (16Gb/s) Fixed-Port Switches.....	7
Brocade Gen 5 (16Gb/s) Directors.....	8
Brocade Gen 6 (32Gb/s) Fixed-Port Switches.....	8
Brocade Gen 6 (32Gb/s) Directors.....	8
Contacting Brocade Technical Support.....	8
Document Feedback.....	9
Overview.....	10
Network Time Protocol.....	10
Most Common Problem Areas.....	10
Questions for common symptoms.....	11
Gathering Information for Your Switch Support Provider.....	13
Setting up your switch for FTP.....	14
Using The <code>supportsave</code> Command.....	14
Capturing Output From a Console.....	15
Capturing Command Output.....	16
Building a Case for Your Switch Support Provider.....	16
Gathering Basic Information.....	16
Gathering Detailed Problem Information.....	17
Gathering Additional Information.....	18
General Troubleshooting.....	19
License Issues.....	19
Time Issues.....	19
Frame Dropping.....	20
Switch message logs.....	20
Switch Boot	21
Rolling Reboot Detection.....	21
FC-FC routing connectivity.....	23
Generating and routing an ECHO.....	23
Superping.....	25
Routing and statistical information.....	28
Performance issues.....	29
Connectivity.....	30
Port initialization and FCP auto-discovery process.....	30
Link issues.....	31
Connection problems.....	32
Checking the physical connection.....	32
Checking the logical connection.....	32
Checking the Name Server	34
Link failures.....	34
Determining a successful speed negotiation.....	35
Checking for a loop initialization failure	36

Checking for a point-to-point initialization failure.....	36
Correcting a port that has come up in the wrong mode	37
Marginal links.....	37
Troubleshooting a marginal link.....	38
Device login issues on Fabric switches.....	40
Pinpointing problems with device logins.....	40
Device login issues on Access Gateway.....	42
Media-related issues.....	43
Testing the external transmit and receive path of a port.....	43
Testing the internal components of a switch.....	43
Testing components to and from the HBA.....	43
Segmented fabrics.....	44
Reconciling fabric parameters individually.....	45
Downloading a correct configuration.....	45
Reconciling a domain ID conflict.....	45
Reconciling incompatible software features.....	47
Configuration.....	48
Configuration upload and download issues.....	48
Gathering additional information.....	50
Brocade configuration form.....	51
Firmware Download Errors.....	52
Blade troubleshooting tips.....	52
Firmware download issues.....	53
Troubleshooting with the firmwareDownload command.....	55
Gathering additional information.....	55
USB error handling.....	55
Considerations for downgrading firmware.....	56
Preinstallation messages.....	56
Blade types.....	58
Firmware versions.....	58
Security	60
User account management.....	60
Password issues.....	60
Password recovery options.....	60
Device authentication	61
Protocol and certificate management	61
Gathering additional information.....	62
SNMP issues.....	62
FIPS	62
Virtual Fabrics.....	64
General Virtual Fabrics Troubleshooting.....	64
Fabric identification issues.....	65
Logical Fabric issues.....	65
Base switch issues.....	65
Logical switch issues.....	66
Switch configuration blade compatibility.....	67
Gathering additional information.....	67
ISL Trunking	68

Trunking Link Issues.....	68
Buffer credit issues.....	69
Getting out of buffer-limited mode	69
Zoning.....	70
Zoning Corrective Action.....	70
Verifying a fabric merge problem.....	70
Verifying a TI zone problem.....	70
Segmented fabrics.....	71
Zone conflicts.....	72
Resolving zoning conflicts.....	73
Correcting a fabric merge problem quickly	73
Changing the default zone access.....	73
Editing zone configuration members.....	74
Reordering the zone member list	74
Checking for Fibre Channel connectivity problems.....	74
Checking for zoning problems.....	75
Gathering additional information.....	76
Diagnostic Features.....	77
Fabric OS diagnostics.....	77
Diagnostic information.....	77
Power-on self-test.....	78
Disabling POST.....	79
Enabling POST.....	79
Switch status.....	80
Viewing the overall status of the switch.....	80
Displaying switch information.....	81
Displaying the uptime for a switch.....	83
Using the spinFab and portTest commands.....	83
Debugging spinFab errors.....	84
Clearing the error counters.....	85
Enabling a port.....	85
Disabling a port.....	85
Port information.....	86
Viewing the status of a port	86
Displaying the port statistics.....	86
Displaying a summary of port errors for a switch.....	87
Equipment status.....	89
Checking the temperature, fan, and power supply.....	89
Checking the status of the fans.....	89
Checking the status of a power supply.....	89
Checking temperature status.....	90
System message log.....	90
Displaying the system message log with no page breaks.....	91
Displaying the system message log one message at a time.....	91
Clearing the system message log	91
Port log.....	91
Viewing the port log.....	91
Syslogd configuration.....	93
Configuring the host.....	93

Configuring the switch.....	93
Automatic trace dump transfers.....	94
Specifying a remote server.....	95
Enabling the automatic transfer of trace dumps.....	95
Setting up periodic checking of the remote server.....	95
Saving comprehensive diagnostic files to the server.....	95
Multiple trace dump files support.....	96
Auto FTP support.....	96
Trace dump support.....	96
Brocade ClearLink Diagnostic Port.....	97
Enhanced support for 32Gbps QSFPs.....	97
Limitations and Considerations.....	97
Example output.....	98
Supported platforms for D_Ports.....	99
Licensing requirements for D_Ports.....	101
Understanding D_Ports.....	101
D_Port configuration modes and testing.....	102
General limitations and considerations for D_Port tests.....	104
Topology 1: ISLs.....	106
Topology 2: ICLs.....	107
Topology 3: Access Gateways.....	107
Saving port mappings on an Access Gateway.....	109
Topology 4: HBA to switch.....	109
Using a D_Port in static-static mode between switches.....	110
Enabling a D_Port in static mode.....	110
Using D_Ports in dynamic mode.....	112
Preprovisioning D_Ports.....	113
Using D_Port mode between switches and HBAs.....	114
Enabling a D_Port in static mode between a switch and an HBA.....	114
Using non-Brocade HBAs for D_Port testing.....	115
BCU D_Port commands.....	116
Host Bus Adapter limitations and considerations for D_Ports.....	116
Confirming SFP and link status with an HBA.....	118
Using a D_Port in on-demand mode	118
Using the fabriclog command.....	119
Calculating buffers for long-distance cables.....	120
Support for audit logs.....	120
Using D_Port show commands.....	120
Switch Type and Blade ID.....	124
Hexadecimal Conversion.....	126
Hexadecimal overview.....	126
Example conversion of the hexadecimal triplet 0x616000.....	126
Decimal-to-hexadecimal conversion table.....	127
Revision History.....	129
FOS-821-TD-UG101; September 28, 2018.....	129
FOS-821-TD-UG100; August 28, 2018.....	129

Introduction

• About This Guide.....	7
• What's new in this document for Fabric OS 8.2.1.....	7
• Supported Hardware and Software.....	7
• Contacting Brocade Technical Support.....	8
• Document Feedback.....	9

About This Guide

This book is a companion guide to be used in conjunction with the *Brocade Fabric OS Administration Guide*. Although it provides many common troubleshooting tips and techniques, it does not teach troubleshooting methodology. Troubleshooting should begin at the center of the SAN—the fabric. Because switches are located between the hosts and the storage devices and have visibility into both sides of the storage network, starting with them can help narrow the search. After eliminating the possibility of a fault within the fabric, see if the problem is on the storage side or the host side, and continue a more detailed diagnosis from there. Using this approach can quickly pinpoint and isolate problems. For example, if a host cannot detect a storage device, run the `switchshow` command to determine if the storage device is logically connected to the switch. If not, focus first on the switch directly connected to storage. Use your vendor-supplied storage diagnostic tools to better understand why the storage device is not visible to the switch. If the storage device can be detected by the switch, and the host still cannot detect the storage device, then there is still a problem between the host and the switch.

What's new in this document for Fabric OS 8.2.1

In addition to general changes to improve clarity and comprehension, the following changes have been made to this document to support this release:

- Revised the publication number.
- Revised the "Rolling Reboot Detection" topic.
- Added "Revision History" to the document.
- Updated the document template.

Supported Hardware and Software

Although many different software and hardware configurations are tested and supported by Brocade for Fabric OS 8.2.1, documenting all possible configurations and scenarios is beyond the scope of this document.

The following hardware platforms are supported by Brocade Fabric OS 8.2.1.

Brocade Gen 5 (16Gb/s) 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 (16Gb/s) 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 (32Gb/s) Fixed-Port Switches

- Brocade G610 Switch
- Brocade G620 Switch
- Brocade G630 Switch
- Brocade 7810 Extension Switch

Brocade Gen 6 (32Gb/s) Directors

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

Contacting Brocade Technical Support

For product support information and the latest information on contacting the Technical Assistance Center, go to <https://www.broadcom.com/support/fibre-channel-networking/>. If you have purchased Brocade product support directly from Brocade, use one of the following methods to contact the Brocade Technical Assistance Center 24x7.

Online	Telephone
For nonurgent issues, the preferred method is to go to MyBrocade (my.brocade.com) and then go to one of the following sites: • My Cases • Software Downloads • Licensing tools • Knowledge Base	Required for Severity 1-Critical and Severity 2-High issues: • North America: 1-800-752-8061 (Toll-free) • International: 1-669-234-1001 (Not toll-free) Toll-free numbers are available in many countries and are listed at https://www.broadcom.com/support/fibre-channel-networking/.

If you purchased Brocade product support from a Brocade OEM/solution provider, contact your OEM/solution provider for all your product support needs.

- OEM/solution providers are trained and certified by Brocade to support Brocade products.
- Brocade provides backline support for issues that cannot be resolved by the OEM/solution provider.
- Brocade Supplemental Support augments your existing OEM support contract, providing direct access to Brocade expertise. For more information, contact Brocade or your OEM.
- For questions regarding service levels and response times, contact your OEM/solution provider.

Document Feedback

Quality is our first concern. We have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission or if 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, publication number, topic heading, page number, and as much detail as possible.

Overview

• Network Time Protocol.....	10
• Most Common Problem Areas.....	10
• Questions for common symptoms.....	11
• Gathering Information for Your Switch Support Provider.....	13
• Building a Case for Your Switch Support Provider.....	16

Network Time Protocol

One of the most frustrating parts of troubleshooting is trying to synchronize a switch's message logs and port logs with those of other switches in the fabric. If you do not have Network Time Protocol (NTP) set up on your switches, trying to synchronize log files to track a problem is more difficult.

Most Common Problem Areas

Table 1 identifies the most common problem areas that arise within SANs, and it identifies tools to use to resolve them.

TABLE 1 Common Troubleshooting Problems and Tools

Problem Area	Investigate	Tools
Fabric	• Missing devices • Marginal links (unstable connections) • Incorrect zoning configurations • Incorrect switch configurations	• Switch LEDs • Switch commands for diagnostics (for example, switchshow or nsallshow) • Web- or GUI-based monitoring and management software tools
Storage Devices	• Physical issues between the switch and devices • Incorrect storage software configurations	• Device LEDs • Storage diagnostic tools • Switch commands for diagnostics (for example, switchshow or nsallshow)
Hosts	• Physical issues between the switch and devices • Lower or incompatible HBA firmware • Incorrect device driver installation • Incorrect device driver configuration	• Device LEDs • Host operating system diagnostic tools • Device driver diagnostic tools • Switch commands for diagnostics (for example, switchshow or nsallshow) Also, make sure to use the latest HBA firmware recommended by the switch supplier or on the HBA supplier's website.
Storage Management Applications	• Incorrect installation and configuration of the storage devices that the software references For example, if using a volume-management application, check for: • Incorrect volume installation • Incorrect volume configuration	• Application-specific tools and resources

Questions for common symptoms

You first must determine what the problem is. Some symptoms are obvious, such as the switch rebooting without any user intervention; whereas other symptoms are more obscure, such as your storage having intermittent connectivity to a particular host. Whatever the symptom is, you must gather information from the devices that are directly involved with the symptom.

The following table lists common symptoms and possible areas to check. You may notice that an intermittent connectivity problem has many variables to research, such as the type of connection between the two devices, how the connection is behaving, and the port type involved.

TABLE 2 Common symptoms

Symptom	Areas to check	Chapter or document
Blade is faulty.	Firmware or application download Hardware connections Power-on self test (POST)	General Troubleshooting on page 19 Firmware Download Errors on page 52 Virtual Fabrics on page 64
Blade is stuck in the "LOADING" state.	Firmware or application download	Firmware Download Errors on page 52
Configuration upload or download fails.	FTP or SCP server or USB availability	Configuration on page 48
Connectivity is intermittent.	Links Trunking Buffer credits FCIP tunnel	Connectivity on page 30 ISL Trunking on page 68 <i>Brocade Fabric OS Extension Configuration Guide</i>
E_Port does not come online.	Licensing Fabric parameters Zoning Hardware or link problems	Overview on page 10 General Troubleshooting on page 19 Connectivity on page 30 Virtual Fabrics on page 64
EX_Port does not form.	Links	Connectivity on page 30 Virtual Fabrics on page 64
Fabric merge fails.	Fabric segmentation	General Troubleshooting on page 19 Connectivity on page 30 Virtual Fabrics on page 64
Fabric segments.	Licensing Zoning Virtual Fabrics Fabric parameters Fabric merge failed	Overview on page 10 General Troubleshooting on page 19 Connectivity on page 30 Virtual Fabrics on page 64
FCIP tunnel bounces.	FCIP tunnel, including the network between FCIP tunnel endpoints	<i>Brocade Fabric OS Extension User Guide</i>
FCIP tunnel does not come online.	FCIP tunnel, including the network between FCIP tunnel endpoints	<i>Brocade Fabric OS Extension User Guide</i>
FCIP tunnel does not form.	Licensing Fabric parameters	General Troubleshooting on page 19 <i>Brocade Fabric OS Extension User Guide</i>
FCIP tunnel is sluggish.	FCIP tunnel, including the network between FCIP tunnel endpoints	<i>Brocade Fabric OS Extension User Guide</i>

TABLE 2 Common symptoms (continued)

Symptom	Areas to check	Chapter or document
FCR is slowing down.	FCR LSAN tags	General Troubleshooting on page 19
Feature is not working.	Licensing	General Troubleshooting on page 19
FICON switch does not talk to hosts.	FICON settings	<i>Brocade Fabric OS FICON User Guide</i>
Firmware download fails.	FTP or SCP server or USB availability Firmware version compatibility Unsupported features enabled Firmware versions on switch	Firmware Download Errors on page 52 Virtual Fabrics on page 64
Host application times out.	FCR LSAN tags Marginal links	General Troubleshooting on page 19 Connectivity on page 30
LEDs are flashing.	Links	Connectivity on page 30
LEDs are steady.	Links	Connectivity on page 30
License has issues.	Licensing	General Troubleshooting on page 19
Link is marginal.	Links	Connectivity on page 30
LSAN is slow or times out.	LSAN tagging	General Troubleshooting on page 19
No connectivity exists between the host and storage.	Cables SCSI timeout errors SCSI retry errors Zoning	Overview on page 10 Connectivity on page 30 ISL Trunking on page 68 <i>Brocade Fabric OS Extension User Guide</i>
No connectivity exists between switches.	Licensing Fabric parameters Segmentation Virtual Fabrics Zoning (if applicable) Incompatible firmware versions	Overview on page 10 General Troubleshooting on page 19 Connectivity on page 30 Virtual Fabrics on page 64
No light on LEDs.	Links	Connectivity on page 30
Performance problems.	Links FCR LSAN tags FCIP tunnels	General Troubleshooting on page 19 Connectivity on page 30 <i>Brocade Fabric OS Extension User Guide</i>
Port cannot be moved.	Virtual Fabrics	Virtual Fabrics on page 64
SCSI experiences retry errors.	Buffer credits FCIP tunnel bandwidth	<i>Brocade Fabric OS Extension User Guide</i>
SCSI experiences timeout errors.	Links HBA Buffer credits FCIP tunnel bandwidth	Connectivity on page 30 ISL Trunking on page 68 <i>Brocade Fabric OS Extension User Guide</i>
Switch constantly reboots.	Rolling reboot detection FIPS	Security on page 60
Switch is unable to join the fabric.	Security policies	Overview on page 10

TABLE 2 Common symptoms (continued)

Symptom	Areas to check	Chapter or document
	Zoning Fabric parameters Switch firmware Virtual Fabrics	Connectivity on page 30 Virtual Fabrics on page 64
Switch reboots during configuration upload or download.	Configuration file discrepancy	Configuration on page 48
Syslog messages are generated.	Hardware SNMP management station	General Troubleshooting on page 19 Security on page 60
Trunk bounces.	Cables on the same port group SFPs Trunked ports	Security on page 60
Trunk does not form.	Licensing Cables on the same port group SFPs Trunked ports Zoning E_Port QoS configuration mismatch Virtual Fabrics (switch ports on either end being in different logical partitions)	Overview on page 10 General Troubleshooting on page 19 Connectivity on page 30 ISL Trunking on page 68
User forgot password.	Password recovery	Security on page 60
User is unable to change switch settings.	RBAC settings Account settings	Security on page 60
Virtual Fabric does not form.	FIDs	Virtual Fabrics on page 64
Zone configuration mismatch exists.	Effective configuration	Zoning on page 70
Zone content mismatch exists.	Effective configuration	Zoning on page 70
Zone type mismatch exists.	Effective configuration	Zoning on page 70

Gathering Information for Your Switch Support Provider

To ensure the maximum availability of diagnostic-related information, consider the following important tips in advance of problems that may occur:

- Always use the **supportftp** command to set up parameters in support of automatic FTP transfers. This ensures that various dumps, FFDC files, and trace files are copied to the defined FTP site and are available for support teams when needed. In addition, you can use **supportftp -t** to ping the remote FTP server to verify its continued availability.
- Consider using syslog on all switches, especially director-class switches. The RASLogs are kept separately on each command processor (CP); syslog provides a means of supplying a chronological error log that can expedite diagnosis. The CPs can fail over during certain problem events, and having an integrated external log can be very helpful.

- If the installation permits, attaching a remote terminal server with logging enabled to the serial console ports of each CP of the director can provide additional advanced information to the support team when made available along with the **supportsave** files.

If you are troubleshooting a production system, you must gather data quickly. As soon as a problem is observed, perform the following tasks. For more information about these commands and their operands, refer to the *Brocade Fabric OS Command Reference*.

1. Enter the **supportsave** command to save RASLog, TRACE, supportshow, core file, FFDC data, and other support information from the switch, chassis, blades, and logical switches.
2. Gather console output and logs.

NOTE

To issue the **supportsave** command on the chassis, you must log in to the switch using an admin account that has chassis-access permissions.

Setting up your switch for FTP

You can use the **supportftp** command to automatically transfer trace dumps to a specified FTP site. This helps minimize trace dump overwrites on the local switch.

To enable automatic trace dump transfers to a specified FTP site, complete the following steps.

1. Connect to the switch, and log in using an account with admin permissions.
2. Enter **supportftp -S** to see the current FTP parameters.
3. Enter **supportftp -s** and respond to the prompts to set the parameters.

```
device:admin> supportftp -s
Host IP Addr[1080:000:000:000:000:017A]:
User Name[njoe]: userFoo
Password[*****]: <hidden>
Remote Dir[support]:
supportftp: parameters changed
```

You can use **supportftp -e** to enable automatic trace dump transfers to your FTP site. This helps minimize trace dump overwrites on the local switch. You can also use **supportftp -t** to periodically ping the FTP server and ensure that it is ready when needed for automatic transfers.

NOTE

Refer to [Automatic trace dump transfers](#) on page 94 for more information on setting up for automatic transfer of diagnostic files as part of the standard switch configuration.

Using The **supportsave** Command

The **supportsave** command uses the default switch name to replace the chassis name regardless of whether the chassis name has been changed to a nonfactory setting. If Virtual Fabrics is enabled, the **supportsave** command uses the default switch name for each logical fabric.

1. Connect to the switch, and log in using an account with admin permissions.

2. Enter the appropriate `supportsave` command based on your needs:
 - If you are saving to an FTP or SCP server, enter `supportsave` with the following options.

NOTE

If `supportsave` is invoked without operands, it goes into interactive mode.

- **-c:** This option uses the FTP parameters saved by the `supportftp` command. This operand is optional; if omitted, specify the FTP parameters through command-line options or interactively. To display the current FTP parameters, run `supportftp` (on a dual-CP system, run `supportftp` on the active CP).

NOTE

The **-c** option works only if you used the `supportftp -s` command to set up the FTP server.

- **-n:** This option runs the command without prompting for confirmation. If this operand is omitted, you are prompted for confirmation.
- On platforms that support USB devices, you can use your Brocade USB device to save the support files. To use your USB device, enter `supportsave` with the following options:
 - **-U:** This option saves support data to an attached USB device. When using this option, a target directory must be specified with the **-d** option.
 - **-d *remote_dir*:** This option specifies the remote directory to which the file is to be transferred. When saving to a USB device, the remote directory is created in the `/support` directory of the USB device by default.

Refer to the *Brocade Fabric OS Command Reference* for commands to mount and view the contents of the USB storage.

Changing The `supportsave` timeout value

While running the `supportsave` command, you may encounter a timeout. A timeout occurs if the system is in a busy state due to the CPU or is I/O bound due to a high level of port traffic or file access requests. A timeout can also occur on very large machine configurations or when the machine is under heavy usage. If a timeout occurs, an SS-1004 message (SS-1004: One or more modules timed out during supportsave. Please retry supportsave with -t option to collect all logs.) is sent to both the console and the RASLog to report the error. To avoid the timeout error, rerun the `supportsave` command with the **-t** option.

1. Connect to the switch and log in using an account with admin permissions.
2. Enter `supportsave -t value`, and specify a multiplier value from 1 through 5.

The following example increases the `supportsave` timeout value to twice the default timeout setting.

```
device:admin> supportsave -t 2
```

Capturing Output From a Console

Some information (such as boot information) is output directly to the console only. To capture this information, you must connect directly to the switch through its management interface using either a serial cable or an RJ-45 connector that is specifically used for Ethernet connectivity to the management network.

1. Connect directly to the switch using a terminal utility.
2. Set the utility to capture output from the screen.

Some utilities require this step to be performed before opening a session. Check with your utility vendor for instructions.

3. Log in to the switch using an account with admin permissions.

4. Enter the command or start the process whose console output you want to capture.

Capturing Command Output

1. Connect to the switch through a Telnet or SSH utility.
2. Log in using an account with admin permissions.
3. Set the Telnet or SSH utility to capture output from the screen.

Some Telnet or SSH utilities require this step to be performed before opening a session. Check with your Telnet or SSH utility vendor for instructions.

4. Enter the command or start the process to capture the required data on the console.

Building a Case for Your Switch Support Provider

The questions listed in [Gathering Basic Information](#) on page 16 should be printed and answered in their entirety and be ready to send to your switch support provider when you contact them. Having this information immediately available expedites the information-gathering process that is necessary to begin determining the problem and finding a solution.

Gathering Basic Information

1. What is the switch's current Fabric OS level?

To determine the switch's Fabric OS level, enter the **firmwareshow** command and write down the information.

2. What is the switch model?

To determine the switch model, enter the **switchshow** command and write down the value in the `Switch Type` field. Cross-reference this value with the chart located in the [Switch Type and Blade ID](#) on page 124 sections.

3. Is the switch operational? Yes or no.
4. Are one or more blades showing a "Faulty" state in the **slotshow** command output? Yes or no.

5. Impact assessment and urgency:

- Is the switch down? Yes or no.
- Is it a standalone switch? Yes or no.
- Are VE, VEX, or EX ports connected to the chassis? Yes or no.

Enter **switchshow** to determine the answer.

- How large is the fabric?

Enter **nsallshow** to determine the answer.

- Do you have encryption blades or switches installed in the fabric? Yes or no.
- Do you have Virtual Fabrics enabled in the fabric? Yes or no.

Enter **lscfg --show** to determine the answer.

- Do you have IPsec installed on the switch's Ethernet interface? Yes or no.

Enter **ipseconfig --show** to determine the answer.

- Do you have In-band Management installed on the switch's Gigabit Ethernet ports? Yes or no.

Enter **portshow iproute ge x** to determine the answer.

- Are you using NPIV? Yes or no.

Use the **switchshow** command to determine the answer.

- Are security policies turned on in the fabric? If so, what are they? Gather the output from the following commands:

- **authutil --show**
- **fddcfg --showall**
- **fipscfg --show**
- **ipfilter --show**
- **secauthsecret --show**
- **secpolicyshow**

6. Is the fabric redundant? If yes, what is the MPIO software? (List vendor and version.)
7. If you have a redundant fabric, did a failover occur? To verify this, you need to examine the RASLogs on both CPs and look for messages related to **hafailover**, for example, HAM-1004.
8. Is POST enabled on the switch? Use the **diagpost** command to verify if POST is enabled.
9. Which CP blade is active? (This is applicable only to chassis-based devices such as Brocade DCX 8510 Backbones and Brocade X6 Directors.) Use the **hashow** command in conjunction with the RASLogs to determine which is the active CP and which is the standby CP. They will reverse roles in a failover, and their logs are separate.

Gathering Detailed Problem Information

Obtain as much of the following information as possible before contacting the SAN technical support vendor.

Document the sequence of events by answering the following questions:

- When did the problem occur?
- Is this a new installation?
- How long has the problem been occurring?
- Are specific devices affected? If so, what are their World Wide Number (WWN) names?
- What happened before the problem?

- Is the problem reproducible? If so, what are the steps to reproduce the problem?
- What configuration was in place when the problem occurred?
- A description of the problem with the switch or the fault with the fabric.
- The last actions or changes made to the system environment:
 - Settings
 - **supportsave** output
- Host information:
 - OS version and patch level
 - HBA type
 - HBA firmware version
 - HBA driver version
 - Configuration settings
- Storage information:
 - Disk/tape type
 - Disk/tape firmware level
 - Controller type
 - Controller firmware level
 - Configuration settings
 - Storage software (EMC Control Center, Veritas SPC, and so on)
- If this is a Brocade DCX 8510 or X6 platform, are the CPs in sync? Yes or no.
Enter `hashow` to determine the answer.
- List the last actions or changes made to the switch, the fabric, and the SAN or metaSAN, as well as when they occurred.
- In [Table 3](#), list the environmental changes made to the network.

TABLE 3 Environmental changes

Type of Change	Date when change occurred

Gathering Additional Information

The following features require you to gather additional information. This additional information is necessary in order for your switch support provider to effectively and efficiently troubleshoot your issue. Refer to the specified chapter or document for the commands used for the data you must capture:

- Configurations: see [Connectivity](#) on page 30
- Firmware downloading: see [Firmware Download Errors](#) on page 52
- Trunking: see [ISL Trunking](#) on page 68
- Zoning: see [Zoning](#) on page 70
- FCIP tunnels: see *Brocade Fabric OS Extension Configuration Guide*
- FICON: see the *Brocade FICON Administration Guide*

General Troubleshooting

• License Issues.....	19
• Time Issues.....	19
• Frame Dropping.....	20
• Switch message logs.....	20
• Switch Boot	21
• FC-FC routing connectivity.....	23

License Issues

Some features require licenses to work properly. To view a list of features and their associated licenses, refer to the *Brocade Fabric OS Software Licensing User Guide*. Licenses are created using a switch's license identifier so you cannot apply one license to different switches. Before calling your switch support provider, use the **licenseshow** command to verify that you have the correct licenses installed.

TABLE 4 Issues related to licences

Symptom	Probable cause and recommended action
A feature is not working.	Refer to the <i>Brocade Fabric OS Software Licensing Guide</i> to determine if the appropriate licenses are installed on the local switch and any connecting switches. To determine installed licenses, complete the following steps. 1. Connect to the switch and log in using an account with admin permissions. 2. Enter the licenseshow command. A list of the currently installed licenses on the switch is displayed.
Unknown. No licenses are installed.	In chassis-based systems, the license keys are stored on the CP blade based on the WWN. So, if you move a CP blade from one chassis to another chassis, and the new chassis already has license keys installed on it, the following occurs: 1. Entering licenseshow displays the following: <pre>device:admin> licenseshow No licenses installed</pre> 2. If the CP blade is moved to a chassis as the standby CP, the active CP will sync over the active CP's license keys and the standby CP will work properly using the active CP's license keys. 3. If the CP blade is moved to a chassis as the active CP, the license keys are shown as "unknown" and the CP does not work because the WWN of the new chassis will not match the license keys that exist on the active CP. To resolve this issue, do one of the following: • Generate new license keys on the active CP based on the WWN of the chassis where the active CP is located. • If you have a configdb file that was generated using the configupload command for this chassis before the new active CP was added, you can run the configdownload command against the configdb file, which has the license keys to restore the active CP. The configdownload command downloads the license keys to the new active CP.

Time Issues

The following table covers time-related issues.

TABLE 5 Issues Related to Time

Symptom	Probable Cause and Recommended Action
Time is not in sync.	Cause: NTP is not set up on the switches in your fabric. Solution: Set up NTP on your switches in all fabrics in your SAN and metaSAN. For more information on setting up NTP, refer to the <i>Brocade Fabric OS Administration Guide</i>.

Frame Dropping

When a frame is unable to reach its destination due to a timeout, it is discarded. You can use Frame Viewer to determine the flows that contained the dropped frames, which can help you determine which applications might be impacted. Using Frame Viewer, you can see exactly what time the frames were dropped. (The timestamps are accurate to within one second.)

You can view and filter up to 20 discarded frames per chip per second for 1200 seconds using a number of options with the **framelog** command.

TABLE 6 Issues Related to Frame Dropping

Symptom	Probable cause and recommended action
Frames are being dropped.	Frames are timing out. Viewing frames: 1. Connect to the switch and log in using an account with admin permissions. 2. Enter framelog --show to determine which frames are being dropped and when.

Switch message logs

Switch message logs (RAS logs) contain information on events that happen on the switch or in the fabric. These logs are an effective tool in understanding what is happening in your fabric or on your switch. RAS logs are independent on director-class switches. Weekly review of the RAS logs is necessary to prevent minor problems from becoming larger or to catch problems at an early stage. There are two sets of logs. The **ipaddrshow** command provides the IP addresses of the CP0 and CP1 control processor blades and associated RAS logs.

The following common problems can occur with your system message log.

TABLE 7 Issues Related to Switch Message Logs (RAS Logs)

Symptom	Probable Cause and Recommended Action
Inaccurate information appears in the system message log.	In rare instances, events gathered by the Track Changes feature can report inaccurate information to the system message log. For example, a user enters a correct user name and password, but the login is rejected because the maximum number of users has been reached. However, when looking at the system message log, the login is reported as successful. If the maximum number of switch users is reached, the switch still performs correctly in that it rejects the login of additional users, even if they enter the correct user name and password. However, in this limited example, the Track Changes feature reports this event inaccurately to the system message log; it appears as though the login was successful. This scenario occurs only when the maximum number of users has been reached; otherwise, the login information displayed in the system message log reflects reality. Refer to the <i>Brocade Fabric OS Administration Guide</i> for information regarding enabling and disabling Track Changes (TC).

TABLE 7 Issues Related to Switch Message Logs (RAS Logs) (continued)

Symptom	Probable Cause and Recommended Action
MQ errors appear in the switch log.	An MQ error is a message queue error. Identify an MQ error message by looking for the two letters MQ followed by a number in the error message: <pre>2004/08/24-10:04:42, [MQ-1004], 218,, ERROR, ras007, mqRead, queue = raslog- test- string0123456-raslog, queue I D = 1, type = 2</pre> MQ errors can result in devices dropping from the switch's Name Server or can prevent a switch from joining the fabric. MQ errors are rare and difficult to troubleshoot; resolve them by working with the switch supplier. When encountering an MQ error, issue the <code>supportsave</code> command to capture debug information about the switch; then, forward the <code>supportsave</code> data to the switch supplier for further investigation.
I2C bus errors appear in the switch log.	I2C bus errors generally indicate defective hardware or poorly seated devices or blades; the specific item is listed in the error message. Refer to the <i>Brocade Fabric OS Message Reference Manual</i> for information specific to the error that was received. Some Chip-Port (CPT) and Environmental Monitor (EM) messages contain I2C-related information. If the I2C message does not indicate the specific hardware that may be failing, begin debugging the hardware, as this is the most likely cause of the errors.
Core file or FFDC warning messages appear on the serial console or in the system log.	Issue the <code>supportsave</code> command. The messages can be dismissed by entering <code>supportsave -R</code> after all data is confirmed to be collected properly. Error example: <pre>*** CORE FILES WARNING (10/22/08 - 05:00:01) *** 3416 KBytes in 1 file(s) use "supportsave" command to upload</pre>

Switch Boot

The following table covers device booting issues.

TABLE 8 Issues Related to Device Booting

Symptom	Probable Cause and Recommended Action
Enterprise-class platform reboots after initial bootup.	This issue can occur during an enterprise-class platform bootup with two CPs. If any failure occurs on the active CP before the standby CP is fully functional and has obtained HA sync, the standby CP may not be able to take on the active role to perform failover successfully. In this case, both CPs reboot to recover from the failure.

Rolling Reboot Detection

A rolling reboot occurs when a switch or enterprise-class platform has continuously experienced unexpected reboots of one or both control processor blades (CPs). This behavior is continuous until the rolling reboot is detected by the system. Once the Rolling Reboot Detection (RRD) occurs, the switch is put into a stable state so that only minimal **supportsave** output need be collected and sent to your service support provider for analysis. Not every type of reboot reason activates the Rolling Reboot Detection feature. For example, issuing the **reboot** command multiple times in itself does not trigger rolling reboot detection, but quickly power-cycling the chassis repeatedly will trigger the RRD process. After five events within the RRD window, the RRD process will first attempt to disable the switch persistently. If a subsequent reboot still occurs, the RRD process will then disable Fabric OS in order to allow improved diagnostic access.

The following message is displayed when the RRD feature is activated and the switch is persistently disabled.

```
*****
* Fabric OS has detected frequent switch reboot condition.      *
* The switch has been disabled as a preventive action. If      *
* the switch reboots again Rolling Reboot Detection will      *
* halt the switch from loading Fabric OS.                      *
```

```

*
*
* Please use 'switchcfgpersistentenable' to enable the
* switch, once the problem is resolved.
*****

```

ATTENTION

If a rolling reboot is caused by a Linux kernel panic, the RRD feature is not activated.

Reboot classification

There are two types of reboots that occur on a switch and enterprise-class platform: expected and unexpected. Expected reboots occur when the reboots are initialized by commands; these types of reboots are ignored by the Rolling Reboot Detection (RRD) feature. They include the following commands:

- **reboot**
- **hafailover**
- **fastboot**
- **firmwaredownload**

The RRD feature is activated and halts rebooting when an unexpected reboot reason is shown continuously in the reboot history within a certain period of time. For example, the switch reboots five times in 300 seconds. The period of time depends on the switch, but typically the period is 10 minutes. The following reboots are considered unexpected reboots:

- Reset
 - A reset reboot may be caused by one of the following:
 - Power-cycling the switch or CP
 - Linux `reboot` command
 - Hardware watchdog timeout
 - Heartbeat-loss-related reboot
- Software fault: kernel panic
 - If the system detects an internal fatal error from which it cannot safely recover, it outputs an error message to the console, dumps a stack trace for debugging, and then performs an automatic reboot.
 - After a kernel panic, the system may not have enough time to write the reboot reason, causing the reboot reason to be empty. This is treated as a reset case.
- Software fault
 - Software watchdog
 - ASSERT
- Software recovery failure

This is an HA bootup-related issue and happens when a switch is unable to recover to a stable state. The HASM log contains more details and specific information on this type of failure, such as one of the following:

- Failover recovery failed: Occurs when failover recovery fails and the CP must reboot.
- Failover when standby CP unready: Occurs when the active CP must fail over, but the standby CP is not ready to take over mastership.
- Failover when LS trans incomplete: Occurs when a logical switch transaction is incomplete.
- Software bootup failure
 - System bring up timed out: The CP failed to come up within the time allotted.
 - LS configuration timed out and failed: The logical switch configuration failed and timed out.

After RRD is activated, admin-level permission is required to log in. Enter the **supportshow** or **supportsave** command to collect a limited amount of data to resolve the issue.

Restrictions

The following restrictions apply to the RRD feature:

- RRD works only on CFOS-based systems and is not available on AP blades.
- If FIPS mode is enabled, the RRD feature works in *record-only* mode.
- RRD works only during the 30 minutes immediately after the switch boots. If the switch does not reboot for 30 minutes, RRD is deactivated.

Collecting limited supportsave output about the Rolling Reboot Detection

1. Log in to the switch using an admin account.
A user account with admin privileges can collect limited **supportsave** output.
2. After you see the message in the following example, press **Enter**.
3. Enter the **supportsave** command to go into interactive mode. Alternatively, if you are using a USB device, enter **usbstorage -e**.
4. Respond to the prompts.
5. Once the **supportsave** command had completed, contact your service support provider to provide them with the data.

The following message is an example of the screen.

```
Fabric OS Version 8.2.0
switch login: admin
Password: <hidden text>
*****
*
* Fabric OS has detected frequent switch reboot condition.
* Following actions can be taken to recover the switch:
* - take off or replace the bad blades.
* - use supportsave to collect supportsave data.
*
*
*****
Please change passwords for switch default accounts now.
Use Control-C to exit or press 'Enter' key to proceed.
```

FC-FC routing connectivity

This section describes tools you can use to troubleshoot Fibre Channel routing connectivity and performance.

Generating and routing an ECHO

The FC-FC Routing Service enables you to route the ECHO generated when an **fcping** command is issued on a switch, providing **fcping** capability between two devices in different fabrics across the FC router.

The **fcping** command sends a Fibre Channel ELS ECHO request to a pair of ports. It performs a zone check between the source and destination. In addition, two Fibre Channel Extended Link Service (ELS) requests are generated. The first ELS request is from the domain controller to the source port identifier. The second ELS request is from the domain controller to the destination port identifiers. The ELS ECHO request elicits an ELS ECHO response from a port identifier in the fabric and validates link connectivity.

Use the **fcping** command to validate link connectivity to a single device or between a pair of devices.

ATTENTION

There are some devices that do not support the ELS ECHO request. In these cases, the device either does not respond to the request or sends an ELS reject. When a device does not respond to the ELS request, further debugging is required; however, do not assume that the device is not connected.

On the edge Fabric OS switch, make sure that the source and destination devices are properly configured in the LSAN zone before entering the **fcping** command. This command performs the following functions:

- Checks the zoning configuration for the two ports specified.
- Generates an ELS ECHO request to the source port specified and validates the response.
- Generates an ELS ECHO request to the destination port specified and validates the response.

```
device:admin> fcping 0x020800 22:00:00:04:cf:75:63:85
Source:      0x020800
Destination: 22:00:00:04:cf:75:63:85
Zone Check:  Zoned
Pinging 0x020800 with 12 bytes of data:
received reply from 0x020800: 12 bytes time:1159 usec
received reply from 0x020800: 12 bytes time:1006 usec
received reply from 0x020800: 12 bytes time:1008 usec
received reply from 0x020800: 12 bytes time:1038 usec
received reply from 0x020800: 12 bytes time:1010 usec
5 frames sent, 5 frames received, 0 frames rejected, 0 frames timeout
Round-trip min/avg/max = 1006/1044/1159 usec
```

Regardless of the device's zoning configuration, the **fcping** command sends the ELS frame to the destination port. A destination device can take any one of the following actions:

- Send an ELS Accept to the ELS request.
- Send an ELS Reject to the ELS request.
- Ignore the ELS request.

For details about the **fcping** command, refer to the *Brocade Fabric OS Command Reference*.

Example of one device that accepts the request and another device that rejects the request

```
device:admin> fcping 10:00:00:00:c9:29:0e:c4 21:00:00:20:37:25:ad:05
Source: 10:00:00:00:c9:29:0e:c4
Destination: 21:00:00:20:37:25:ad:05
Zone Check: Not Zoned
Pinging 10:00:00:00:c9:29:0e:c4 [0x20800] with 12 bytes of data:
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1162 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1013 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1442 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1052 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1012 usec
5 frames sent, 5 frames received, 0 frames rejected, 0 frames timeout
Round-trip min/avg/max = 1012/1136/1442 usec
Pinging 21:00:00:20:37:25:ad:05 [0x211e8] with 12 bytes of data:
Request rejected by 21:00:00:20:37:25:ad:05: Command not supported: time: 1159 usec
Request rejected by 21:00:00:20:37:25:ad:05: Command not supported: time: 1006 usec
Request rejected by 21:00:00:20:37:25:ad:05: Command not supported: time: 1008 usec
Request rejected by 21:00:00:20:37:25:ad:05: Command not supported: time: 1038 usec
Request rejected by 21:00:00:20:37:25:ad:05: Command not supported: time: 1010 usec
5 frames sent, 0 frames received, 5 frames rejected, 0 frames timeout
Round-trip min/avg/max = 0/0/0 usec
```

Example using **fcping** with a single destination (in this example, the destination is a device node WWN)

```
device:admin> fcping 20:00:00:00:c9:3f:7c:b8
Destination: 20:00:00:00:c9:3f:7c:b8
Pinging 20:00:00:00:c9:3f:7c:b8 [0x370501] with 12 bytes of data:
```



```

received reply from 20:00:00:00:c9:3f:7c:b8: 12 bytes time:825 usec
received reply from 20:00:00:00:c9:3f:7c:b8: 12 bytes time:713 usec
received reply from 20:00:00:00:c9:3f:7c:b8: 12 bytes time:714 usec
received reply from 20:00:00:00:c9:3f:7c:b8: 12 bytes time:741 usec
received reply from 20:00:00:00:c9:3f:7c:b8: 12 bytes time:880 usec
5 frames sent, 5 frames received, 0 frames rejected, 0 frames timeout
Round-trip min/avg/max = 713/774/880 usec

```

Superping


Superping refers to the **fcping** --allpaths command, which is a diagnostic tool used to test all least-cost ISLs between a source and destination switch. When you run the command, you are provided with a list of all available least-cost paths from a source domain to a destination device. Superping isolates links with potential failures so that you can investigate these ISLs to determine the exact links.

Superping works by sending ECHO frames to a destination device and outputting the status of each ISL it traverses whether or not the response from the destination device is received. Each ECHO frame can choose any path from multiple available paths in the fabric to reach the destination device. This utility allows you to do the following:

- Run a sanity test that exercises all the ISLs and internal links in different paths that route to the destination device.
- Determines the least-cost path to aid in designing fabric redundancy.
- Determines the specific ISLs and internal links with failures.
- Exercises all ISL links in the base fabric for a logical fabric configuration.

The number of actual paths covered when using the superping feature depends on two other parameters that you can specify optionally. When you enter **fcping** --allpaths without any other options, Superping covers all ISLs in the routes between source to destination, as shown in [Figure 1](#).

FIGURE 1 Superping and fcPing paths



In the following example, Superping is invoked using the **fcping --allpaths** command to destination domain 165. The following example displays each hop in (Domain1/Index1-> Domain2/Index2) format. To reach destination domain 165 from source domain 3 there are two unique end-to-end paths. In the first path, the frame traverses from egress port index 205 on source domain 3 to ingress port index 25 on domain 207. On domain 207, the frame traverses from egress port index 42 to ingress port index 3 in domain 101. On domain 101, the frame goes from egress port index 16 to ingress port index 99 on domain 165.

```
device:admin> fcping -allpaths 165
Pinging(size:12 bytes) destination domain 165 through all paths
PATH SWITCH1--> SWITCH2--> SWITCH3 SWITCH4 STATUS
-----
1. (3/EMB, 3/205) [128] (207/25,207/42) [128] (101/3,101/16) [128] (165/99,165/0) [128] SUCCESS
2. (3/EMB, 3/204) [128] (207/27,207/42) [128] (101/3,101/16) [128] (165/99,165/0) [128] SUCCESS
```

Superping can isolate links with failures so that you can further investigate these ISLs to determine the exact links giving the errors.

NOTE

Superping provides an indication if all ISLs are covered. If all the ISLs are not covered, you can increase the coverage count and maximum retries to transmit, so that complete coverage of all ISLs is achieved.

Consider the following example in which a few errors are recorded on ISLs 3/205 to 2/25, 3/204 to 2/27, 2/42-->101/3, and 2/1 to 101/8. The potential faulty link is internal port 0/284 on domain 2 with the maximum of 100 percent failure.

ISL COVERAGE			
SNO	ISL		STATUS
1	3/123[128]--> 165/96[128]		SUCCESS (5/5)
2	3/205[128]--> 2/25[128]		FAILURE (7/50)
3	3/204[128]--> 2/27[128]		FAILURE (11/50)
4	165/99[128]--> 101/16[128]		SUCCESS (5/5)
6	2/42[128]--> 101/3[128]		FAILURE (10/67)
7	2/1[128]--> 101/8[128]		FAILURE (8/33)

INTERNAL PORT COVERAGE			
SNO	DOMAIN	INTRNL_PORT	STATUS
1	2[128]	0/272	SUCCESS (40/40)
2	2[128]	0/276	SUCCESS (44/44)
3	2[128]	0/280	SUCCESS (30/30)
4	2[128]	0/284	FAILURE (20/20) <== 100% failure

When an ECHO frame is dropped, all the ISLs in the path are marked as failed. It is not possible to determine the exact ISL link that dropped the frame. Because of this, all the ISLs in the path record some failures. The ISL with the actual error has the maximum percentage of failures, as this ISL, when selected in any possible path, causes the ECHO frame to be dropped and accumulates a higher failure percentage.

Restrictions

- Fabric reconfiguration cannot occur while using the Superping feature. It is assumed that the fabric is stable before the **fcping --allpaths** command is executed.
- The control path for interswitch communication should be available, even if the data path for device-to-device communication may have resource starvation.
- When executed in a fabric with trunk ports, only the trunk master index is output to the user (for example, individual coverage statistics for each trunk-member are not available).
- All switches must have Fabric OS 6.3.0 or later.
- Superping requires that the Fibre Channel ECHO ELS frame be supported by end-devices.
- In TI Zones, when failover is disabled and Superping is executed on a destination device included in the TI Zone, then Superping displays failures on all ISLs that are not part of the TI Zone. Also, when Superping is executed on a device that is not present in a TI Zone, failures are shown on all ISLs that are part of any TI Zone.
- This feature is not supported in interopMode 2 or 3.
- In frame redirection configurations, where there is a physical host, physical target, virtual initiator, and virtual target; Superping only identifies the path from the physical host to the physical target regardless of whether the data path consists of the path from the physical target to the virtual target through the virtual initiator.

Routing and statistical information

The **pathinfo** command displays routing and statistical information from a source port index on the local switch to a destination port index on another switch. This routing information describes the full path that a data stream travels between these ports, including all intermediate switches.

ATTENTION

Using the **pathinfo** command when exchange-based routing is turned on can provide different paths with each attempt.

The routing and statistics information are provided by every switch along the path, based on the current routing-table information and statistics calculated continuously in real time. Each switch represents one hop.

Use the **pathinfo** command to display routing information from a source port on the local switch to a destination port on another switch. The command output describes the exact data path between these ports, including all intermediate switches.

When using the **pathinfo** command in Fabric OS 6.3.0 and later across fabrics connected through an FC router, the command represents backbone information as a single hop. The command captures details about the FC router to which ingress and egress EX_Ports are connected, but it hides the details about the path the frame traverses from the ingress EX_Ports to the egress EX_Ports in the backbone.

To use **pathinfo** across remote fabrics, you must specify both the fabric ID (FID) and the domain ID of the remote switch. Optionally, you can specify the source PID and destination PID. You cannot use the **pathinfo** command to obtain source port information across remote FCR fabrics. When obtaining path information across remote fabrics, the destination switch must be identified by its domain ID. Identifying the switch by name or WWN is not accepted.

To display basic path information to a specific domain in command line mode:

```
device:admin> pathinfo 5
```

Hop	In Port	Domain ID (Name)	Out Port	BW	Cost
0	2	1 (sw0)	6	4G	500
1	23	2 (sw0)	8	4G	500
2	4	3 (sw0)	3	4G	500
3	12	4 (sw0)	18	4G	10000
4	4	7 (switch_3)	0	4G	500
5	26	5 (switch_3)	E	-	-

To display basic and extended statistics in interactive mode:

```
device:admin> pathinfo
Max hops: (1..127) [25]
Fabric Id: (1..128) [-1]
Domain|Wwn|Name: [] 8
Source port: (0..15) [-1]
Destination port: (0..255) [-1]
Source pid: (0x0..0xffff00) [ffffff] 0x061600
Destination pid: (0x0..0xffff00) [0] 0x01f001
Basic stats (yes, y, no, n): [no] y
Extended stats (yes, y, no, n): [no] y
Trace reverse path (yes, y, no, n): [no]
Source route (yes, y, no, n): [no]
Timeout: (1000..30000) [10000]
Target port is Embedded
```

Hop	In Port	Domain ID (Name)	Out Port	BW	Cost
0	2	1 (sw0)	6	4G	500
1	23	2 (sw0)	8	4G	500
2	4	3 (sw0)	3	4G	500
3	2	4 (sw0)	24	4G	10000
4	3	7 (switch_3)	2	4G	500
5	27	5 (switch_3)	24	-	-
Reverse path					
6	24	5 (switch_3)	27	4G	500
7	2	7 (switch_3)	3	4G	500

```

8          24          4 (sw0)          2          4G          500
9          3          3 (sw0)          4          4G          10000
10         8          2 (sw0)         23          4G          500
11         6          1 (sw0)          2          -          -
(output truncated)

```

For more information on the **pathinfo** command, refer to the *Brocade Fabric OS Command Reference*.

Performance issues

The following table covers performance issues.

TABLE 9 Issues related to performance

Symptom	Probable cause and recommended action
General slow-down in FCR performance and scalability.	Issue: As LSAN zone databases get bigger, it takes more switch resources to process them. Solution: Use the enforce tag feature to prevent a backbone switch from accepting unwanted LSAN zone databases into its local database.
Host application times out.	Issue: The FCR tends to take a long time (more than 5 seconds) to present and set up paths for the proxy devices. Certain hosts are able to do discovery much faster and, as a result, they end up timing out. Solution: Use the speed tag feature to always present the target proxy to the host. This helps sensitive hosts to do a quick discovery without timing out or causing an application failure.

Connectivity

• Port initialization and FCP auto-discovery process.....	30
• Link issues.....	31
• Connection problems.....	32
• Link failures.....	34
• Marginal links.....	37
• Device login issues on Fabric switches.....	40
• Device login issues on Access Gateway.....	42
• Media-related issues.....	43
• Segmented fabrics.....	44


Port initialization and FCP auto-discovery process

The steps in the port initialization process represent a protocol used to discover the type of connected device and establish the port type and port speed. The possible port types are as follows:

- U_Port—Universal FC port. The base Fibre Channel port type and all unidentified, or uninitiated ports are listed as U_Ports.
- L_Port/FL_Port—Fabric Loop port. Connects public loop devices.
- G_Port—Generic port. Acts as a transition port for non-loop fabric-capable devices.
- E_Port—Expansion port. Assigned to ISL links.
- F_Port—Fabric port. Assigned to fabric-capable devices.
- EX_Port—A type of E_Port. It connects a Fibre Channel router to an edge fabric. From the point of view of a switch in an edge fabric, an EX_Port appears as a normal E_Port. It follows applicable Fibre Channel standards as other E_Ports. However, the router terminates EX_Ports rather than allowing different fabrics to merge as would happen on a switch with regular E_Ports.
- M_Port—A mirror port. A mirror port allows you to configure a switch port to connect to a port to mirror a specific source port and destination port traffic passing through any switch port. This is only supported between F_Ports.
- VE_Port—A virtual E_Port. A Gigabit Ethernet switch port configured for an FCIP tunnel is called a VE_Port (virtual E_Port). However, with a VEX_Port at the other end, it does not propagate fabric services or routing topology information from one edge fabric to another.
- VEX_Port—A virtual EX_Port. It connects a Fibre Channel router to an edge fabric. From the point of view of a switch in an edge fabric, a VEX_Port appears as a normal VE_Port. It follows the same Fibre Channel protocol as other VE_Ports. However, the router terminates VEX_Ports rather than allowing different fabrics to merge as would happen on a switch with regular VE_Ports.

Figure 2 shows the process behind port initialization. Understanding this process can help you determine where a problem resides. For example, if your switch cannot form an E_Port, you understand that the process never got to that point or does not recognize the switch as an E_Port. Possible solutions would be to look at licensing and port configuration. Verify that the correct licensing is installed or that the port is not configured as a loop port, a G_Port, or the port speed is not set.

FIGURE 2 Simple port initialization process



The FCP auto-discovery process enables private storage devices that accept the process login (PRLI) to communicate in a fabric.

If device probing is enabled, the embedded port commences with port login (PLOGI) and attempts a PRLI into the device to retrieve information to enter into the Name Server. This enables private devices that do not perform a fabric login (FLOGI), but accept PRLI, to be entered in the Name Server and receive full fabric citizenship.

A fabric-capable device registers information with the Name Server during a FLOGI. These devices typically register information with the Name Server before querying for a device list. The embedded port will still conduct PLOGI and attempt PRLI with these devices.

To display the contents of a switch's Name Server, use the **nsShow** or **nsAllShow** command. For more information about these Name Server commands, refer to the *Brocade Fabric OS Command Reference*.

Link issues

The following table covers link-related issues, as indicated by port-status LEDs.

TABLE 10 Issues related to port-status LEDs

Symptom	Probable cause and recommended action
Port LEDs are flashing.	Depending on the rate of the flash and the color of the port LED, this could mean several things. To determine what is happening on either your port status LED or power status LED, refer to the hardware installation guide or reference manual for that switch. There is a table that describes the purpose of the LEDs and explains the current behavior as well as suggested resolutions.

TABLE 10 Issues related to port-status LEDs (continued)

Symptom	Probable cause and recommended action
Port LEDs are steady.	The color of the port LED is important in this instance. To determine what is happening on either your port status LED or power status LED, refer to the hardware installation guide or reference manual for that switch. There is a table that describes the purpose of the LEDs and explains the current behavior as well as suggested resolutions.
No light from the port LEDs.	If there is no light coming from the port LED, then no signal is being detected. Check your cable and SFP to determine the physical fault.

Connection problems

Determine if the problem is the target or the host, then continue to divide the suspected problem-path in half until you can pinpoint the problem. One of the most common solutions is zoning. Verify that the host and target are in the same zone. For more information on zoning, refer to [Zoning](#) on page 70.

Checking the physical connection

- Check the cables running to and from the host and storage to the switch.
This path includes the patch panel. Verify that none of the cables are damaged, including indentations or bent cable.
- Check the SFP on the HBAs and switches.
Verify that they are known to be in good working condition. You can do this by swapping the current SFP with a known working SFP.
- Clean the optics.
There are many kits on the market for cleaning fiber optics. You want to find a product that does not leave residue either from a lint-free wipe or from the solvent.

Checking the logical connection

1. Enter the **switchShow** command.
2. Review the output from the command and determine if the device successfully logged in to the switch.
 - A device that is logically connected to the switch is registered as an F_Port, L_Port, E_Port, EX_Port, VE_Port, VEX_Port, or N_Port.
 - A device that is *not* logically connected to the switch is registered as a G_Port or U_Port, if NPIV is not on the switch.
3. Enter the **slotShow -m** command to verify that all blades are enabled and not faulty, disabled, or in some other non-available state.

4. Perform the appropriate actions based on how your missing device is connected:
 - If the missing device is logically connected, proceed to the next troubleshooting procedure ([Checking the Name Server](#) on page 34).
 - If the missing device is *not* logically connected, check the device and everything on that side of the data path. Also refer to [Link failures](#) on page 34 for additional information.

Checking the path includes verifying the following for the host:

- The host OS is configured correctly.
- The third-party vendor multi-pathing input/output (MPIO) software, if it is being used, is configured correctly.
- The HBA and storage device and the driver and firmware are compatible with the switch based on the compatibility matrix.
- The driver settings and binaries are up-to-date.
- The device Basic Input Output System (BIOS) settings are correct.
- The HBA configuration is correct according to manufacturer's specifications.
- The SFPs in the HBA are compatible with the host's HBA.
- The SFP on the switch is compatible with the switch.
- The switch settings related to the host are configured correctly.

Checking the path includes the following for the target:

- The driver settings and binaries are up-to-date.
- The device Basic Input Output System (BIOS) settings are correct.
- The HBA configuration is correct according to the manufacturer's specifications.
- The SFPs in the HBA are compatible with the target HBA.
- The switch settings related to the target are configured correctly.

Refer to [Checking for a loop initialization failure](#) on page 36 as the next potential trouble spot.

Checking the Name Server

1. Enter the **nsshow** command on the switch to determine if the device is attached:

```
switch:admin> nsshow
The Local Name Server has 9 entries {
Type Pid      COS PortName      NodeName      TTL(sec)
*N  021a00;    2,3;20:00:00:e0:69:f0:07:c6;10:00:00:e0:69:f0:07:c6; 895
    Fabric Port Name: 20:0a:00:60:69:10:8d:fd
NL  051edc;    3;21:00:00:20:37:d9:77:96;20:00:00:20:37:d9:77:96; na
    FC4s: FCP [SEAGATE ST318304FC 0005]
    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL  051ee0;    3;21:00:00:20:37:d9:73:0f;20:00:00:20:37:d9:73:0f; na
    FC4s: FCP [SEAGATE ST318304FC 0005]
    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL  051ee1;    3;21:00:00:20:37:d9:76:b3;20:00:00:20:37:d9:76:b3; na
    FC4s: FCP [SEAGATE ST318304FC 0005]
    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL  051ee2;    3;21:00:00:20:37:d9:77:5a;20:00:00:20:37:d9:77:5a; na
    FC4s: FCP [SEAGATE ST318304FC 0005]
    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL  051ee4;    3;21:00:00:20:37:d9:74:d7;20:00:00:20:37:d9:74:d7; na
    FC4s: FCP [SEAGATE ST318304FC 0005]
    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL  051ee8;    3;21:00:00:20:37:d9:6f:eb;20:00:00:20:37:d9:6f:eb; na
    FC4s: FCP [SEAGATE ST318304FC 0005]
    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL  051eef;    3;21:00:00:20:37:d9:77:45;20:00:00:20:37:d9:77:45; na
    FC4s: FCP [SEAGATE ST318304FC 0005]
    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
N   051f00;    2,3;50:06:04:82:bc:01:9a:0c;50:06:04:82:bc:01:9a:0c; na
    FC4s: FCP [EMC SYMMETRIX 5267]
    Fabric Port Name: 20:0f:00:60:69:10:9b:5b
```

2. Look for the device in the Name Server (NS) list, which lists the nodes connected to that switch. This allows you to determine if a particular node is accessible on the network.
 - If the device is not present in the NS list, the problem is between the device and the switch. There may be a time-out communication problem between edge devices and the name server, or there may be a login issue. First check the edge device documentation to determine if there is a time-out setting or parameter that can be reconfigured. Also, check the port log for NS registration information and FCP probing failures (using the **fcpprobeshow** command). If these queries do not help solve the problem, contact the support organization for the product that appears to be inaccessible.
 - If the device is listed in the NS, the problem is between the storage device and the host. There may be a zoning mismatch or a host/storage issue. Proceed to [Zoning](#) on page 70.
3. Enter the **portloginshow** command to check the port login status.
4. Enter the **fcpprobeshow** command to display the FCP probing information for the devices attached to the specified F_Port or L_Port. This information includes the number of successful logins and SCSI INQUIRY commands sent over this port and a list of the attached devices.
5. Check the port log to determine whether or not the device sent the FLOGI frame to the switch, and the switch probed the device.

Link failures

A link failure occurs when a server, storage, or switch device is connected to a switch, but the link between the devices does not come up. This prevents the devices from communicating to or through the switch.

Use the **switchShow** command to find the reason for link failures. If the **switchShow** command or LEDs indicate that the link has not come up properly, use one or more of the following procedures.

The port negotiates the link speed with the opposite side. The negotiation usually completes in one or two seconds; however, sometimes the speed negotiation fails.

Determining a successful speed negotiation

NOTE

Skip this procedure if the port speed is set to a static speed through the **portCfgSpeed** command.

1. Enter the **portcfgshow** command to display the port speed settings of all the ports.
2. Enter the **switchshow** command to determine if the port has module light.
3. Enter the **portcfgspeed** command to change the port speed to 4, 8, 10, 16, 32 Gbps, depending on what speed can be used by both devices. This should correct the negotiation by setting to one speed.
4. Enter the **portlogshow** or **portlogdump** command.
5. Check the events area of the output:

The following example shows the **portlogdump** output. The output of the **portlogshow** command is very similar.

```
device:admin> portlogdump
time      task      event  port  cmd  args
-----
14:38:51.976  SPEE    sn      200   NC   00000001,00000000,00000001
14:39:39.227  SPEE    sn      200   NC   00000002,00000000,00000001
```

- In the event column, "sn" indicates a speed negotiation.
- In the cmd column, "NC" indicates the negotiation has completed.

If these fields do not appear, proceed to Step 6.

6. If the fields in Step 5 do not appear, correct the negotiation by entering the **portcfgspeed** command with a speed option. Valid values are:
 - 0 - Auto Negotiate
 - 4 - 4Gbps
 - 8 - 8Gbps
 - 10 - 10Gbps
 - 16 - 16Gbps
 - 32 - 32Gbps
 - ax - Auto Negotiate plus enhanced retries

Checking for a loop initialization failure

1. Verify the port is an L_Port.
 - a) Enter the **switchshow** command.
 - b) Check the last field of the output to verify that the switch port indicates an L_Port. If a loop device is connected to the switch, the switch port must be initialized as an L_Port.
 - c) Check to ensure that the state is online; otherwise, check for link failures.

The following **switchshow** example shows an online L_Port.

```
device:admin> switchshow
Area Port Media Speed State      Proto
=====
(output truncated)
66   66   --    N8     No_Module
67   67   id    AN     No_Sync
68   68   id    N2     Online    L-Port  13 public
```

2. Verify that loop initialization occurred if the port to which the loop device is attached does not negotiate as an L_Port.
 - a) Enter either **portlogshow** or **portlogdump** to display the port log for all ports on the switch; or if you are looking for a specific port, enter **portlogdumpport**.
 - b) Check argument number four for the loop initialization soft assigned (LISA) frame "0x11050100".

```
switch:admin> portlogdumpport 4
time          task  event  port  cmd  args
-----
11:40:02.078  PORT  Rx3    23    20    22000000,00000000,ffffffff,11050100 Received LISA frame
```

The LISA frame indicates that the loop initialization is complete.

3. Point-to-point initialization sometimes causes trouble with older HBAs. You can skip point-to-point initialization by using the **portcfglport** command.

The switch will then change to point-to-point initialization after the LISA phase of the loop initialization.

Checking for a point-to-point initialization failure

1. Enter **switchshow** to confirm that the port is active and has a module that is synchronized.
If a fabric device or another switch is connected to the switch, the switch port must be online.
2. Enter either the **portlogshow** or the **portlogdump** command.
Verify the event area for the event entry is "pstate". The command entry "AC" indicates that the port has completed point-to-point initialization.

```
device:admin> portlogdumpport 4
time          task  event  port  cmd  args
-----
11:38:21.726  INTR   pstate  4      AC
```

3. Skip over the loop initialization phase.

After becoming an active port, the port becomes an F_Port or an E_Port depending on the device on the opposite side. If the opposite device is a host or target device, the port becomes an F_Port. If the opposite device is another switch, the port becomes an E_Port.

If there is a problem with the host or target device, enter **portcfggport** to force the port to try to come up as point-to-point only.

Correcting a port that has come up in the wrong mode

1. Enter **switchshow**.
2. Check the output from the **switchshow** command and follow the suggested actions in [Table 11](#).

TABLE 11 Suggested actions based on switchshow command output

Output	Suggested action
Disabled	If the port is disabled because persistent disable or security reasons, attempt to resolve the issue and then enter the portenable or, if persistently disabled, portcfgpersistentenable command.
Bypassed	The port may be testing.
Loopback	The port may be testing.
E_Port	If the opposite side is not another switch, the link has come up in a wrong mode. Check the output from the portlogshow or portlogdump commands and identify the link initialization stage where the initialization procedure went wrong.
F_Port	If the opposite side of the link is a private loop device or a switch, the link has come up in a wrong mode. Check the output from portlogshow or portlogdump commands.
G_Port	The port has not come up as an E_Port or F_Port. Check the output from portlogshow or portlogdump commands and identify the link initialization stage where the initialization procedure went wrong.
L_Port	If the opposite side is not a loop device, the link has come up in a wrong mode. Check the output from portlogshow or portlogdump commands and identify the link initialization stage where the initialization procedure went wrong.

NOTE

If you are unable to read a port log dump, contact your switch support provider for assistance.

Marginal links

A marginal link involves the connection between the switch and the edge device. Isolating the exact cause of a marginal link involves analyzing and testing many of the components that make up the link (including the switch port, switch SFP, cable, edge device, or edge device SFP). Troubleshooting a marginal link can involve inspecting the error counters described in [Troubleshooting a marginal link](#) on page 38 or running diagnostics on a link, a port, or an end-to-end path.

you can use the **portloopbacktest** command to verify the functional operation of a path on a switch. This test sends frames from a given port's transmitter and loops them back into the same port's receiver. The loopback is done at the parallel loopback path. The path traversed in this test does not include the media or the fiber cable.

Only one frame is transmitted and received at any given time. An external cable is not required to run this test. The port LEDs will flicker green rapidly while the test is running.

[Table 12](#) shows the different loopback modes you can use when using **portloopbacktest** to test a marginal link.

TABLE 12 Loopback modes

Loopback mode	Description
1	Port Loopback (loopback plugs)
2	External Serializer/Deserializer (SerDes) loopback
5	Internal (parallel) loopback (indicates no external equipment)
7	Back-end bypass and port loopback
9	Back-end bypass and internal loopback
11	Back-end bypass and external loopback

NOTE

Mode-8 is not supported on (Gen 6) platforms except Brocade 7810 Extension Switch. Instead mode-11 has been introduced. Mode-7 is supported only in Brocade DCXchassis with FX8-24. For Mode-7 test to go through, you should have optical loopback plugs in the FX8-24 blade's FE ports.

Troubleshooting a marginal link

Marginal links can cause intermittent problems in fabrics and reduce performance levels. To locate marginal links, complete the following steps.

1. Run the standard D_Port tests.
2. Enter **porterrshow**.

The following is a typical example of the output of this command.

```
device:admin> porterrshow
      frames  enc  crc  crc   too  too  bad  enc  disc  link  loss  loss  frjt  fbsy
      tx   rx   in   err g_eof shrt long eof  out c3   fail sync sig
=====
0:  665k 7.0k   0   0   0   0   0   0   6   0   0   1   2   0   0
1:   0   0   0   0   0   0   0   0   0   0   0   0   2   0   0
2:   0   0   0   0   0   0   0   0   0   0   0   0   1   0   0
3:   0   0   0   0   0   0   0   0   0   0   0   0   1   0   0
4:   0   0   0   0   0   0   0   0   0   0   0   0   1   0   0
5:   0   0   0   0   0   0   0   0   0   0   0   0   1   0   0
6:   0   0   0   0   0   0   0   0   0   0   0   0   1   0   0
7:   0   0   0   0   0   0   0   0   0   0   0   0   1   0   0
8:   78   60   0   0   0   0   0   0   7   0   0   3   6   0   0
9:   12    4   0   0   0   0   0   0   3   0   0   1   2   0   0
10:  0   0   0   0   0   0   0   0   0   0   0   0   1   0   0
11:  0   0   0   0   0   0   0   0   0   0   0   0   1   0   0
12:  0   0   0   0   0   0   0   0   0   0   0   0   1   0   0
13:  0   0   0   0   0   0   0   0   0   0   0   0   1   0   0
14:  0   0   0   0   0   0   0   0   0   0   0   0   1   0   0
15:  0   0   0   0   0   0   0   0   0   0   0   0   1   0   0
16: 665k 7.4k   0   0   0   0   0   0   6   0   0   1   2   0   0
(output truncated)
```

3. Examine the output to determine if there are a relatively high number of errors (such as CRC errors or ENC_OUT errors), or if there are a steadily increasing number of errors to confirm a marginal link. You should sample the data every five minutes until you see the counters increment.
 - The frames tx and frames rx values are the number of frames being transmitted and received.
 - The crc_err counter records frames with CRC errors. If this counter goes up, then the physical path should be inspected. Check the cables to and from the switch, patch panel, and other devices. Check the SFP by swapping it with a known good working SFP.

If you see this issue on a Gen 5 8 Gbps blade, use the **portcfgfillword** command to reduce EMI.

- The crc_g_eof counter records frames with CRC errors and a good EOF. The first port detecting a CRC error marks the frame with a bad EOF and passes the frame on to its destination. Subsequent ports in the path also detect the CRC error and the crc_err counter increments on these ports. However, because the first port marked the frame with a bad EOF, the good EOF counter on the subsequent ports does not increment. The marginal link associated with the port with an increasing good EOF counter is the marginal link and the source of the errors.
- The enc_out counter records errors that occur outside the frame and usually indicate a bad primitive. To determine if you are having a cable problem, take snapshots of the port errors using **porterrshow** every 5 to 10 minutes. If you notice the crc_err counter go up, you have either a bad or damaged cable or device in the path.

NOTE

ICLs will see enc_out errors when ports on one side of the link are disabled.

- The disc_c3 counter records discarded class 3 errors, which means that the switch is holding onto the frame longer than the hold time allows. One problem this could be related to is ISL oversubscription.
4. If you suspect a marginal link, isolate the areas by moving the suspected marginal port cable to a different port on the switch. Reseating the SFPs may also cure marginal port problems.

If the problem stops or goes away, either the switch port or the SFP is marginal (proceed to Step 6).

If the problem does not stop or go away, refer to Step 5.

5. Perform the following steps to rule out cabling issues:
 - a) Insert a new cable in the suspected marginal port.
 - b) Enter the **porterrshow** command to determine if a problem still exists.
 - If the **porterrshow** output displays a normal number of generated errors, the issue is solved.
 - If the **porterrshow** output still displays a high number of generated errors, follow the troubleshooting procedures for the host or storage device in [Device login issues on Fabric switches](#) on page 40.
6. If there are no cabling issues, run **portloopbacktest** on the marginal port. You will need an adapter to run the loopback test for the SFP. Alternatively, you can run the test on the marginal port using the loopback mode "lb=5".

Use the modes shown in [Marginal links](#) on page 37 to test the port. Refer to the *Brocade Fabric OS Command Reference* for additional information on this command.
7. Check the results of the loopback test and proceed as follows:
 - If the loopback test failed, the port is bad. Block the port, replace the port blade, or replace the switch as needed.
 - If the loopback test did not fail, the SFP is bad, and should be replaced.

Device login issues on Fabric switches

A correct login is when the port type matches the device type that is plugged in. In the following example, it shows that the device connected to Port 1 is a fabric point-to-point device and it is correctly logged in an F_Port.

```
switch:admin> switchshow
switchName:      brcd_G620
switchType:      162.0
switchState:     Online
switchMode:      Native
switchRole:      Subordinate
switchDomain:     1
switchId:        fffc01
switchWwn:       10:00:00:00:00:00:00:00
zoning:          OFF
switchBeacon:    OFF
FC Router:       OFF
FC Router BB Fabric ID: 1
Area Port Media Speed State          Proto
=====
  0  0  --   N32   No_Module
  1  1  --   N16   Online          FC  F-Port  10:00:00:05:1e:8f:c1:31
  2  2  --   N32   No_Module
  3  3  --   N32   No_Module
(output truncated)
 61 61  --   N32   No_Module
 62 62  --   N32   No_Module
 63 63  --   N32   No_Module
 64 64  id   N32   Online          E-Port  10:00:00:05:1e:34:d0:05 "1_d1" (Trunk master)
 65 65  --   N32   No_Module
 66 66  --   N32   No_Module
 67 67  id   AN    No_Sync
 68 68  id   N16   Online          L-Port  13 public
 69 69  --   N32   No_Module
 70 70  --   N32   No_Module
 71 71  id   N16   Online          L-Port  13 public
 72 72  --   N32   No_Module
 73 73  --   N32   No_Module
 74 74  --   N32   No_Module
 75 75  --   N32   No_Module
 76 76  id   N16   Online          E-Port  10:00:00:05:1e:34:d0:05 "1_d1" (upstream) (Trunk master)
 77 77  id   N16   Online          F-Port  10:00:00:06:2b:0f:6c:1f
 78 78  --   N32   No_Module
 79 79  id   N16   Online          E-Port  10:00:00:05:1e:34:d0:05 "1_d1" (Trunk master)
```

Pinpointing problems with device logins

1. Log in to the switch as admin.
2. Enter the **switchShow** command and check the login state.

3. Enter the **portCfgShow** command to see if the port is configured correctly.

In some cases, you may find that the port has been locked as an L_Port and the device attached is a fabric point-to-point device such as a host or switch. This would be an incorrect configuration for the device and therefore the device cannot log in to the switch.

To correct this type of problem, unlock the port using the **portCfgLport [[slot]/port] 0** command. The **portCfgDefault** command may cause other settings to be changed.

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

4. Enter the **portErrShow** command; then, check for errors that can cause login problems. A steadily increasing number of errors can indicate a problem. Track errors by sampling the port errors every five or ten minutes until you see the problem occur again.
5. Enter the **portFlagsShow** command; then, check to see how a port has logged in and where a login failed (if a failure occurred):

```
switch:admin> portflagsshow
Port SNMP Physical Flags
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
0 Offline In_Sync PRESENT U_PORT LED
1 Online In_Sync PRESENT ACTIVE F_PORT G_PORT U_PORT LOGICAL_ONLINE LOGIN NOELP LED ACCEPT
2 Offline No_Light PRESENT U_PORT LED
3 Offline No_Module PRESENT U_PORT LED
4 Offline No_Module PRESENT U_PORT LED
5 Offline No_Light PRESENT U_PORT LED
6 Offline No_Module PRESENT U_PORT LED
7 Offline No_Module PRESENT U_PORT LED
8 Offline No_Light PRESENT U_PORT LED
9 Offline No_Light PRESENT U_PORT LED
10 Offline No_Module PRESENT U_PORT LED
11 Offline No_Module PRESENT U_PORT LED
12 Offline No_Module PRESENT U_PORT LED
13 Offline No_Module PRESENT U_PORT LED
14 Online In_Sync PRESENT ACTIVE F_PORT G_PORT U_PORT LOGICAL_ONLINE LOGIN NOELP LED ACCEPT
15 Online In_Sync PRESENT ACTIVE E_PORT G_PORT U_PORT SEGMENTED LOGIN LED
```

6. Enter the **portLogDumpPort** command; then, view the device-to-switch communication.

```
switch:admin> portlogdumpport 8 | more
time          task          event    port cmd  args
-----
Thu Nov  6 16:52:39 2008
16:52:39.066 PORT          scn       8      1  00010004,4302000f,02000000
16:52:39.066 PORT          scn       8      2  ce3dfab0,d9672800,00000002
16:52:39.066 PORT          scn       8      2  ce3dfab0,d9672800,00000080
16:52:39.066 PORT          scn       8      5  00000000,00000000,00000002
16:52:39.066 PORT          scn       8      1  00010004,4302000f,00000002
16:52:39.066 PORT          scn       8      1  00010004,4302000f,02000000
16:52:39.071 PORT          ioctl     88010004 1,0 * 4
16:52:42.311 SPEE          sn        8      WS  00000000,00000000,00000000
16:52:42.558 SPEE          sn        8      NM  00000000,00000000,00000000
16:52:42.558 SPEE          sn        8      NF  00000000,00000000,00000000
16:52:42.558 SPEE          sn        8      NC  00000001,00000000,00000000
16:52:42.559 LOOP          loopscn   8      LIP  8002
16:52:42.559 LOOP          loopscn   8      LIP  f7f7
16:52:42.572 LOOP          loopscn   8      LIM  0
16:52:42.572 PORT          Tx3       8      12  22000000,00000000,ffffff,11010000
16:52:42.572 PORT          Rx3       8      12  22000000,00000000,ffffff,11010000
16:52:42.572 PORT          Tx3       8      20  22000000,00000000,ffffff,11020000
16:52:42.572 PORT          Rx3       8      20  22000000,00000000,ffffff,11020000
16:52:42.572 PORT          Tx3       8      20  22000000,00000000,ffffff,11030000
16:52:42.572 PORT          Rx3       8      20  22000000,00000000,ffffff,11030000
```

NOTE

Refer to [Port log](#) on page 91 for overview information about **portLogDump**

Device login issues on Access Gateway

Hosts might have problems logging into the fabric through an Access Gateway under the following conditions:

- Hosts are connected to an Access Gateway.
- F-ports on Access Gateway have NPIV logins.
- Different hosts login and logout of the same Access Gateway F_Port.
- Access Gateway Persistent AL_PA feature is enabled.
- Devices are connected to Access Gateway port configured as N_Ports.
- More than 126 logins attempted via an F_Port configured for default number of logins (126).
- **porttrunkarea** is configured on edge switch but the trunking license is missing on the Access Gateway device.
- Link level issues between edge switch and access gateway
- Other access gateway policy specific features like port grouping, advance device security

To resolve this issue, follow these steps:

1. Identify all the affected F_Ports with duplicate ALPA entries and do the following for each port.
2. Enter **ag --printalpamap port_#** to print the ALPA map.
3. Enter **portdisable port_#** to disable all the affected F_Ports with duplicate ALPA entries.
4. Enter **ag --clearalpamap port_#** to clear the ALPA map.
5. Enter **portenable port_#** to reenabale the specified port.

Media-related issues

This section provides procedures that help pinpoint any media-related issues, such as bad cables and SFPs, in the fabric. The tests listed in [Table 13](#) are a combination of structural and functional tests that can be used to provide an overview of the hardware components and help identify media-related issues.

- *Structural* tests perform basic testing of the switch circuit. If a structural test fails, replace the main board or port blade.
- *Functional* tests verify the intended operational behavior of the switch by running frames through ports or bypass circuitry.

TABLE 13 Component test descriptions

Test name	Checks
portTest	Used to isolate problems to a single replaceable element and isolate problems to near-end terminal equipment, far-end terminal equipment, or transmission line. Diagnostics can be executed every day or on demand.
spinFab	Tests switch-to-switch ISL cabling and trunk group operations.

The following procedures are for checking switch-specific components.

Testing the external transmit and receive path of a port

1. Connect to the switch and log in as admin.
2. Connect external loopback cables to the port that needs to be tested.
3. Disable the switch/chassis using `chassisdisable` command followed by `portloopbacktest -lb_mode 1 -ports <bladeportnumber>`.

Testing the internal components of a switch

1. Connect to the switch and log in as admin
2. Disable the switch/chassis using `chassisdisable` command followed by `portloopbacktest -lb_mode 2` command where 2 is the operand that causes the test to run on the internal switch components (Serdes level loopback).

Testing components to and from the HBA

To test components to and from an HBA, complete the following steps.

1. Connect to the switch and log in as admin.

- Enter the **porttest** command (refer to the *Brocade Fabric OS Command Reference* for information on the command options).

Refer to [Table 14](#) for a list of additional tests that can be used to determine the switch components that are not functioning properly. Refer to the *Brocade Fabric OS Command Reference* for additional command information.

The **bcu fcdiag --linkbeacon** command can be used to beacon a target port on the switch. These commands work only in Brocade-branded or qLogic BR-series adapters.

TABLE 14 Switch component tests

Test	Function
portbeacon	Sets port beaconing mode.
portloopbacktest	Performs a functional test of port N-to-N path. Verifies the functional components of the switch.
turboramtest	Verifies that the on chip SRAM located in the ASIC is using the Turbo-Ram BIST circuitry. This allows the BIST controller to perform the SRAM write and read operations at a much faster rate.

Segmented fabrics

Fabric segmentation is generally caused by one of the following conditions:

- Incompatible fabric parameters (refer to [Reconciling fabric parameters individually](#) on page 45)
- Incompatible zoning configuration (refer to [Zoning](#) on page 70)
- Domain ID conflict (refer to [Reconciling fabric parameters individually](#) on page 45)
- Fabric ID conflict (refer to [Virtual Fabrics](#) on page 64)
- Incompatible security policies
- Incorrect fabric mode
- Incorrect policy distribution
- Incompatible software features

There are a number of settings that control the overall behavior and operation of the fabric. Some of these values, such as the domain ID, are assigned automatically by the fabric and can differ from one switch to another in the fabric. Other parameters, such as the BB credit, can be changed for specific applications or operating environments, but must be the same among all switches to allow the formation of a fabric.

In general, the following fabric parameters must be identical on each switch for fabrics to merge. However, the following table summarizes the scenarios when the fabrics are not segmented (but merged) even when these fabric parameters are not identical.

TABLE 15 Segmented fabric scenarios

Fabric Parameter	LISL	XISL	ISL (Fabric OS 7.1.0 or later)	ISL (Fabric OS 7.3.0 or earlier)
Domain ID	Not segmented (merged)	Segmented	Segmented	Segmented
R_A_TOV	Segmented	Segmented	Segmented	Segmented
E_D_TOV	Segmented	Segmented	Segmented	Segmented
Data field size	Segmented	Segmented	Segmented	Segmented
Sequence level switching	Not segmented (merged)	Segmented	Segmented	Segmented
Disable device probing	Not segmented (merged)	Not segmented (merged)	Segmented	Not segmented (merged)
Suppress class F traffic	Segmented	Segmented	Segmented	Segmented
Per-frame route priority	Not segmented (merged)	Segmented	Segmented	Segmented
Long-distance fabric	Not segmented (merged)	Segmented	Segmented	Segmented

TABLE 15 Segmented fabric scenarios (continued)

Fabric Parameter	LISL	XISL	ISL (Fabric OS 7.1.0 or later)	ISL (Fabric OS 7.3.0 or earlier)
Virtual Channel parameters	Not segmented (merged)	Not segmented (merged)	Segmented	Not segmented (merged)

NOTE

The long-distance fabric parameter is not needed to be identical on Brocade 6505, 6510, 6520, and G620 switches, Brocade DCX 8510 Backbones, and Brocade X6 Directors.

Reconciling fabric parameters individually

1. Log in to one of the segmented switches as admin.
2. Enter the **configShow -pattern "fabric.ops"** command.
3. Log in to another switch in the same fabric as admin.
4. Enter the **configShow -pattern "fabric.ops"** command.
5. Compare the two switch configurations line by line and look for differences. Do this by comparing the two Telnet windows or by printing the **configShow -pattern "fabric.ops"** output. Also, verify that the fabric parameter settings (refer to the list of fabric parameters in [Segmented fabrics](#) on page 44) are the same for *both* switches.
6. Connect to the segmented switch after the discrepancy is identified.
7. Disable the switch by entering the **switchDisable** command.
8. Enter the **configure** command to edit the appropriate fabric parameters for the segmented switch.
9. Enable the switch by entering the **switchEnable** command.

Alternatively, you can reconcile fabric parameters by entering the **configUpload** command for each switch and upload a known-good configuration file. If you do this option, the two switches must be the same model.

Downloading a correct configuration

You can restore a segmented fabric by downloading a previously saved correct backup configuration to the switch. Downloading in this manner reconciles any discrepancy in the fabric parameters and allows the segmented switch to rejoin the main fabric. For details on uploading and downloading configurations, refer to the Brocade Fabric OS Administration Guide.

Reconciling a domain ID conflict

If a domain ID conflict appears, the conflict is only reported at the point where the two fabrics are physically connected. However, there may be several conflicting domain IDs, which appear as soon as the initial conflict is resolved.

Typically, the fabric automatically resolves domain conflicts during fabric merges or builds unless Insistent Domain ID (IDID) is configured. If IDID is enabled, switches that cannot be programmed with a unique domain ID are segmented out. Check each switch that has IDID configured and make sure their domain IDs are unique within the configuration.

Repeat the following procedure until all domain ID conflicts are resolved.

1. Enter **fabricshow** on a switch from one of the fabrics.
2. In a separate Telnet window, enter **fabricshow** on a switch from the second fabric.

3. Compare the **fabricshow** outputs from the two fabrics. Note the number of domain ID conflicts; there may be several duplicate domain IDs that must be changed.

Determine which switches have domain overlap and change the domain IDs for each of those switches.

4. Choose the fabric on which to change the duplicate domain ID; connect to the conflicting switch in that fabric.
5. Enter **switchdisable**.
6. Enter **configure**.
7. When the **Fabric Parameters** prompt displays, enter **y** and press **Enter**.
8. When the **Domain** prompt displays, enter the new domain ID number and press **Enter**.
9. Press **Enter** for all remaining prompts to accept their default settings.



CAUTION

Do not press Ctrl+C as this will exit the configuration without saving it.

10. Enter **switchenable**.

This enables the joining switch to obtain a new domain ID as part of the process of coming online. The fabric principal switch allocates the next available domain ID to the new switch during this process.

11. Repeat step 4 through step 10 if additional switches have conflicting domain IDs.

The following example illustrates setting the domain ID.

```
device:admin> switchdisable
device:admin> configure
Configure...
Fabric parameters (yes, y, no, n): [no] y
Domain: (1..239) [1] 89
WWN Based persistent PID (yes, y, no, n): [no]
Allow XISL Use (yes, y, no, n): [yes]
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000..5000) [2000]
WAN_TOV: (0..30000) [0]
MAX_HOPS: (7..19) [7]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1) [0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1) [0]
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]
BB credit: (1..27) [16]
Disable FID Check (yes, y, no, n): [no]
Insistent Domain ID Mode (yes, y, no, n): [no]
Virtual Channel parameters (yes, y, no, n): [no]
F-Port login parameters (yes, y, no, n): [no]
Zoning Operation parameters (yes, y, no, n): [no]
RSCN Transmission Mode (yes, y, no, n): [no]
Arbitrated Loop parameters (yes, y, no, n): [no]
System services (yes, y, no, n): [no]
Portlog events enable (yes, y, no, n): [no]
ssl attributes (yes, y, no, n): [no]
rpcd attributes (yes, y, no, n): [no]
webtools attributes (yes, y, no, n): [no]
WARNING: The domain ID will be changed. The port level zoning may be affected
```

Reconciling incompatible software features

Earlier releases of software may not be supported in new versions of Fabric OS. This may be due to a software feature changing or new services being supported. If you suspect that you are trying to introduce a switch into a fabric that has an older version of code, check the release notes to verify that any features on that switch are supported in the fabric with the newer code.

When the Management Server (MS) Platform services are enabled on a switch running Fabric OS 7.0.0 and later and you try to merge this switch into a fabric that does not have this feature enabled, the switch does not merge and a segmentation occurs. To resolve this, either turn the MS Platform services off or enable them on every switch in the fabric.

In Fabric OS 7.0.0 and later, an ESC frame is used to exchange fabric parameters to detect Enhanced TI Zones, interoperability mode, and Virtual Fabric FID conflicts. If at any point during the ESC frame exchange, a link with incompatible parameters is detected, the switch running Fabric OS 7.0.0 and later does not join into the existing fabric. To fix this issue, refer to the *Brocade Fabric OS Administration Guide* for more information on that specific software feature.

Configuration

- Configuration upload and download issues.....48
- Brocade configuration form..... 51

Configuration upload and download issues

It is important to maintain consistent configuration settings on all switches in the same fabric because inconsistent parameters (such as inconsistent PID formats) can cause fabric segmentation. As part of standard configuration maintenance procedures, it is recommended that you back up all important configuration data for every switch on a host computer server for emergency reference and possible download purposes.

If the configuration download fails, you must examine the following aspects:

- Does it download the correct file for the switch?
- Is the switch logically partitioned and was the correct FID used?
- Was the switch changed recently with logical switch management such that FIDs may no longer correspond?

NOTE

For information about Virtual Fabrics using Fabric OS 6.3.0 or later, refer to the *Brocade Fabric OS Administration Guide*.

The following table covers configuration upload and download issues.

TABLE 16 Issues related to configuration uploading and downloading

Symptom	Probable cause and recommended action
The configuration upload fails.	If the configuration upload fails, It may be because of one or more of the following reasons: • The FTP or SCP server's host name is not known to the switch. Verify with your network administrator that the switch has access to the FTP server. • The USB path is not correct. If your platform supports a USB memory device, verify that it is connected and running. Verify that the path name is correct by using the usbstorage -l command. Example of usbstorage -l command <pre>device:admin> usbstorage -l firmwarekey\ 0B 2016 Aug 15 15:13 support\ 106MB 2016 Aug 24 05:36 support1034\ 105MB 2007 Aug 23 06:11 config\ 0B 2016 Aug 15 15:13 firmware\ 380MB 2016 Aug 15 15:13 FW_v6.0.0\ 380MB 2016 Aug 15 15:13 Available space on usbstorage 74%</pre> • The FTP or SCP server's IP address cannot be contacted. Verify that you can connect to the FTP server. Use your local PC to connect to the FTP server or ping the FTP server. Example of a successful ping <pre>C:\> ping 192.001.001.001 Pinging 192.001.001.001 with 32 bytes of data: Reply from 192.001.001.001: bytes=32 time=5ms TTL=61 Ping statistics for 192.001.001.001: Packets: Sent = 4, Received = 4, Lost = 0 (0%loss),</pre>

TABLE 16 Issues related to configuration uploading and downloading (continued)

Symptom	Probable cause and recommended action
	Approximate round trip times in milli-seconds: Minimum = 4ms, Maximum = 5ms, Average = 4ms If your ping is successful from your computer, but you cannot reach it from inside your data center, there could be a block on the firewall to not allow FTP connections from inside the data center. Contact your network administrator to determine if this is the cause and to resolve it by opening the port up on both inbound and outbound UDP and TCP traffic. Example of a failed ping <pre>C:\> ping 192.001.001.001 Pinging 192.001.001.001 with 32 bytes of data: Request timed out. Request timed out. Request timed out. Request timed out. Ping statistics for 192.001.001.001: Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),</pre> If your ping has failed, then you should verify the following: - The ports are open on the firewall. - The FTP server is up and running. • You have configured Admin Domain on the switch. A message such the following is displayed: <pre>2016/11/08/-05:17:34, [ZONE-1069], 28, FID 128, WARNING, SW6510_35_111, Admin Domains are not supported in Fabric OS 8.1.0 (Event:AD_Update, AD4).</pre> • You do not have configuration upload permission on the switch. There may be some restrictions based on Role-Based Access Control settings. For more information on these types of restrictions, refer to the <i>Brocade Fabric OS Administration Guide</i>. • You do not have permission to write to the directory on the FTP or SCP server. Example of a failed login to the FTP server For an unsuccessful login, the output should be similar to the following example: <pre>C:\> ftp 192.001.001.001 Connected to 192.001.001.001 220 Welcome to Services FTP service. User (10.10.10.10:(none)): userFoo 331 Please specify the password. Password: <hidden> 530 Login incorrect. Login failed.</pre> If your login to the FTP or SCP server has failed, verify that the user name and password combination you used is correct. • On a Virtual Fabrics-enabled switch, you do not have the chassis role permission set on your user account. You will need to change this permission or use an account that does have this permission and then retry the upload. Implement one change at a time, then issue the command again. By implementing one change at a time, you are able to determine what works and what does not work. Knowing which change corrected the problems will help you avoid this problem in the future.
The configuration download fails	If the configuration download fails, It may be because of one or more of the following reasons: • The FTP or SCP server's host name is not known to the switch. Verify with your network administrator that the switch has access to the FTP server. • The USB path is incorrect.

TABLE 16 Issues related to configuration uploading and downloading (continued)

Symptom	Probable cause and recommended action
	If your platform supports a USB memory device, verify that it is connected and running. Verify that the path name is correct. It should be the relative path from <code>/usb/usbstorage/brocade/configdownload</code> or use absolute path. NOTE Root access is required to see the mentioned path. - The FTP or SCP server's IP address cannot be contacted. Verify that you can connect to the FTP server. Use your local PC to connect to the FTP server or ping the FTP server. - There was a reason to disable the switch. Note, however, that you must disable the switch for some configuration downloads. For more information on how to perform a configuration download without disabling a switch, refer to the <i>Brocade Fabric OS Administration Guide</i>. - You do not have permission on the host to perform configuration download. There may be some restrictions if you are using Role-Based Access Control. For more information on these types of restrictions, refer to the <i>Brocade Fabric OS Administration Guide</i>. - The configuration file you are trying to download does not exist on the host. - The configuration file you are trying to download is not a switch configuration file. - If you selected the (default) FTP protocol, the FTP server is not running on the host. - The configuration file that you are trying to download uses incorrect syntax. - The user name and password combination you used is incorrect.
The switch reboots during the configuration download.	When you run configdownload -vf, having both CPs reboot is normal. This option downloads the Virtual Fabrics-related data. After the reboot, you can continue following the instructions. Otherwise, try the command again without the -vf operand to download the regular configuration data. This does not cause a reboot. However, a notice that "A switch reboot is required for the changes to take effect." is displayed, and you will need to reboot the device. NOTE There is no reliable mechanism to determine which parameters may or may not require a reboot. To ensure that all configuration changes are applied correctly, Brocade strongly recommends that you always reboot the switch after a configuration download.
Configuration did not seem to change after the configuration download process finished.	Verify that the switch was rebooted by checking the system log. If you are doing this on a enterprise-class platform, verify that both CPs rebooted by checking the system log. If any error occurs during the download, such as an error about a particular key, it is important to enter configdefault and then run the configdownload command again.

Gathering additional information

Be sure to capture the output from the commands you are issuing both from the switch and from your computer when you are analyzing the problem.

Send this and all logs to your switch support provider.

Messages captured in the logs

Configuration download generates both RASlog and Audit log messages resulting from execution of the **configDownload** command.

The following messages are written to the logs:

- configDownload completed successfully ... (RASlog and Audit log)

- configUpload completed successfully ... (RASlog)
- configDownload not permitted ... (Audit log)
- configUpload not permitted ... (RASlog)
- (Warning) Downloading configuration without disabling the switch was unsuccessful. (Audit log)

Brocade configuration form

Print and use the following table as a hard copy reference to record the configuration information for the various blades, switches, and chassis.

TABLE 17 Brocade configuration and connection

Brocade configuration settings	Value
IP address	
Gateway address	
Chassis configuration option	
Management connections	
Serial cable tag	
Ethernet cable tag	
Configuration information	
Domain ID	
Switch name	
Ethernet IP address	
Ethernet subnet mask	
Total number of local devices (nsshow)	
Total number of devices in fabric (nsallshow)	
Total number of switches in the fabric (fabricshow)	

Firmware Download Errors

• Blade troubleshooting tips.....	52
• Firmware download issues.....	53
• Troubleshooting with the firmwareDownload command.....	55
• USB error handling.....	55
• Considerations for downgrading firmware.....	56

Blade troubleshooting tips

This chapter refers to the following specific types of blades inserted into the Brocade DCX 8510 Backbones and X6 Directors:

- FC blades or port blades contain only Fibre Channel ports: Brocade FC16-32/48/64, FC32-48, and FC 32-64.
- AP blades such as the FX8-24 contain extra processors, some with specialized ports, and have the Fabric OS firmware downloaded to them automatically in the chassis.
- CP blades have a control processor (CP) used to control the entire switch; they can be inserted only into slots 6 and 7 on the Brocade DCX 8510-8, slots 4 and 5 on the Brocade DCX 8510-4, and slots 1 and 2 on the Brocade X6-4 and Brocade X6-8.
- CR16-8 and CR16-4 core blades provide ICL functionality between two Brocade DCX 8510 Backbones. CR16-8 blades can be inserted only into slots 5 and 8 on the Brocade DCX 8510-8. CR16-4 blades can be inserted only into slots 3 and 6 on the Brocade DCX 8510-4.

Typically, issues detected during firmware download to AP blades do not require recovery actions on your part.

If you experience frequent failovers between CPs that have different versions of firmware, then you may notice multiple blade firmware downloads and a longer startup time.

- CR32-8 and CR32-4 core blades provide ICL functionality between two Brocade X6 Directors. CR32-8 blades can be inserted only into slots 7 and 8 on the Brocade X6-8. CR32-4 blades can be inserted only into slots 5 and 6 on the Brocade X6-4.

Blade fault auto recovery

If a blade is faulty, it should attempt once for an auto recovery. The auto recovery gets triggered for the specified slot faults provided auto recovery feature is enabled .only one attempt is made for auto recovery and If it does not happen successfully , you have to recover the blade. The auto recovery is disabled by default. Enable the auto recovery feature using the **configurechassis** command.

```
device:admin> configurechassis
Configure...
  cfgload attributes (yes, y, no, n): [no]
  Custom attributes (yes, y, no, n): [no]
  system attributes (yes, y, no, n): [no] y
  system.blade.bladeFaultOnHwErrMsk: (0x0..0x7fffffff) [0x0]
  system.cpuLoad: (10..121) [121]
  system.i2cTurboCnfg: (0..2) [1]
  system.Enable.bladeAutoRecovery (yes, y, no, n): [yes] no
  fos attributes (yes, y, no, n): [no]
```

The following table covers blade-related issues.

TABLE 18 Issues related to blades

Symptom	Probable cause and recommended action
The blade is faulty. (Use the slotshow command to confirm this.)	If the port or application blade is faulty, enter the slotpoweroff and slotpoweron commands for the port or application blade. If the port or application blade still appears to be faulty, remove it and re-insert it into the chassis.
The AP blade is stuck in the "LOADING" state. (Use the slotshow command to confirm this.)	If the blade remains in the loading state for a significant period of time, the firmware download times out. Remove the blade and re-insert it. When it boots up, autoleveling is triggered and the firmware download is attempted again.

Firmware download issues



CAUTION

After you start the firmware download process, do not enter any disruptive commands (such as reboot) that interrupt the process. The entire firmware download and commit process can take up to 30 minutes. If there is a problem, wait for the time-out (30 minutes for network problems) before entering the **firmwaredownload** command again. Disrupting the process can render the switch inoperable and require you to seek help from your switch service provider. Do not disconnect the switch from power during the process because the switch could become inoperable when it reboots.

The following table describes common firmware download issues and their recommended actions.

TABLE 19 Issues related to firmware downloading

Symptom	Probable cause and recommended action
Firmware download times out.	This can be caused by an excessively slow network. If it takes more than 30 minutes to download firmware on a switch, or on each CP in a director, the firmware download process times out. If a timeout occurs on a switch, the firmware download process synchronizes the two partitions on the switch by starting a firmware commit operation. If a timeout occurs in a director, the firmware download process synchronizes the firmware on the two partitions on the CP blades by starting a firmware commit operation on each CP. Wait at least 15 minutes for the commit operation to complete, then use the firmwareshow command to verify the partitions are synchronized. In some older versions of firmware, the firmware commit operation may not be started automatically on the switch (or on the standby CP in the director). In this case, you can enter the firmwarecommit command manually on the switch (or on the standby CP in the director) to synchronize the partitions. After the firmware commit operation completes, re-issue the firmwaredownload command to upgrade the system. To avoid timeouts, ensure that you are using the highest connection speeds supported on your device for downloads. Use the 1 GbE management port connection if supported.
Server is inaccessible or firmware path is invalid.	- The FTP or SCP server's host name is not known to the switch. Verify with your network administrator that the switch has access to the FTP server. Verify the path to the FTP or SCP server is accessible from the switch. For more information on checking your FTP or SCP server, refer to Configuration on page 48. - The USB path is not correct. If your platform supports a USB memory device, verify that it is connected and running using the usbstorage -e command. If USB memory is connected, you are notified, and commands such as usbstorage -l will work if the device is running. If device is running, verify that the path name is correct using usbstorage -l. NOTE The entire firmware directory must be loaded onto the USB device under the "firmware" directory prior to using the USB device for firmware download.

TABLE 19 Issues related to firmware downloading (continued)

Symptom	Probable cause and recommended action
	Here is an example of the usbstorage -l command output: <pre>switch:admin> usbstorage -l firmware\ 1585MB 2016 Oct 22 17:46 v8.1.0_main_bld14\ 1585MB 2016 Apr 22 14:16 config\ 0B 2016 Oct 22 17:46 support\ 0B 2016 Oct 22 17:46 firmwarekey\ 0B 2016 Oct 22 17:46 Available space on USB storage 16%</pre> Here is an example of an error message: <pre>switch:admin> firmwaredownload Download from USB [No]: y Firmware filename: v8.2.0 Server IP: 127.1.1.1, Protocol IPv4 Checking system settings for firmwaredownload.. Failed to access usb://127.1.1.1/usb/usbstorage/Brocade/firmware/ v8.2.0/release.plist Cannot access the firmware on USB device. Please check the firmware path.</pre>
Cannot download the requested firmware.	The firmware you are trying to download on the switch is incompatible. Check the firmware version against the switch type. If the firmware is incompatible, retrieve the correct firmware version and try again. Example of error message <pre>switch:admin> firmwaredownload Server Name or IP Address: 192.000.000.000 User Name: userFoo File Name: /users/home/userFoo/firmware/v8.1.0 Network Protocol(1-auto-select, 2-FTP, 3-SCP) [1]: 2 Password: <hidden> Server IP: 192.000.000.000, Protocol IPv4 Checking system settings for firmwaredownload... Cannot download the requested firmware because the firmware doesn't support this platform. Please enter another firmware path.</pre>
Cannot download on a switch with Interop mode turned on.	On single CP system, the Interop fabric does not support Coordinated HotCode Load. Enter firmwaredownload -o. Using -o bypasses the checking of Coordinated HotCode Load (HCL). On single-CP systems in interop fabrics, the HCL protocol is used to ensure data traffic is not disrupted during firmware upgrades. This option allows a firmware download to continue even if HCL is not supported in the fabric or the protocol fails. Using this option may cause traffic disruption for some switches in the fabric.
You receive a "firmware download is already in progress" message.	The firmware download process has already been started and it is in progress. Wait until it completes. You can use the firmwaredownloadstatus and firmwaredownloadshow commands to monitor its progress. If the problem persists, contact your switch support provider. This error commonly occurs when the application blades are upgrading or a commit operation is in progress. Example of a firmware download already in progress: <pre>switch:admin> firmwaredownload Server Name or IP Address: 192.000.000.000 User Name: userFoo File Name: /users/home/userFoo/firmware/v8.2.0 Network Protocol(1-auto-select, 2-FTP, 3-SCP) [1]: 2 Password: <hidden> Server IP: 192.000.000.000, Protocol IPv4 Checking system settings for firmwaredownload... Sanity check failed because firmwaredownload is already in progress.</pre>

Troubleshooting with the `firmwareDownload` command

ATTENTION

Do not run mixed firmware versions on CPs.

A network diagnostic script and preinstallation check is a part of the firmware download procedure. The script and preinstallation check performs troubleshooting and automatically checks for any blocking conditions. If the firmware download fails, refer to the *Brocade Fabric OS Message Reference* for details about error messages. Also refer to [Considerations for downgrading firmware](#) on page 56.

If a firmware download fails in a director, the **`firmwaredownload`** command synchronizes the firmware on the two partitions of each CP by starting a firmware commit operation. Wait at least 15 minutes for this commit operation to complete before attempting another firmware download. Use the **`firmwaredownloadstatus`** command to determine the commit status if needed.

If the firmware download fails in a director or enterprise-class platform, the CPs may end up with different versions of firmware and are unable to achieve HA synchronization between the two blades. In such cases, issue the **`firmwaredownload -s`** command on the standby CP; the single mode (**`-s`**) option allows you to upgrade the firmware on the standby CP to match the firmware version running on the active CP. This situation can also occur during a CP replacement in which the replacement CP has a lower firmware version. HA should be disabled and the firmware should be updated on the replacement CP before any other operations are considered.

Then reissue the **`firmwaredownload`** command to download the desired firmware version to both CPs. For example, if CP0 is running Fabric OS 7.3.0 on the primary and secondary partitions, and CP1 is running Fabric OS 7.4.0 on the primary and secondary partitions, then synchronize them by issuing the **`firmwaredownload`** command. Alternatively, use the **`firmwaresync`** command on the active CP to copy the currently active firmware to the standby CP.

NOTE

Some of the messages include error codes (as shown in the following example). These error codes are for internal use only and you can disregard them. Port configuration with EX ports enabled along with trunking for port(s) 63, use the `portCfgEXPort`, `portCfgVEXPort`, and `portCfgTrunkPort` commands to remedy this. Verify blade is ENABLED. (error 3)

Gathering additional information

You should follow these best practices for firmware download before you start the procedure:

- Keep all session logs.
- Enter the **`supportSave`** or the **`supportShow`** command before and after entering the **`firmwareDownload`** command.
- If a problem persists, package together all of the information (the Telnet session logs and serial console logs, output from the **`supportSave`** command) for your switch support provider. Make sure you identify what information was gathered before and after issuing the **`firmwareDownload`** command.

USB error handling

[Table 20](#) outlines how the USB device handles errors under specific scenarios and details what actions you should take after the error occurs.

TABLE 20 USB error handling

Scenario under which download fails	Error handling	Action
An access error occurs during firmware download because the removal of the USB device, or USB device hardware failure, and so on.	Firmware download times out and commit is started to repair the partitions of the CPUs that are affected.	None.
USB device is not enabled.	Firmware download fails with an error message	Enable the USB device using the usbstorage -e command and retry firmware download.

Considerations for downgrading firmware

The pre-installation check of the **firmwaredownload** command detects all of the blocking conditions that can prevent a successful downgrade, and warns you about all these conditions. The error messages displayed by the **firmwaredownload** command states the blocking conditions and the corresponding commands to correct them. You must address all of these blocking conditions before proceeding. Refer to the *Brocade Fabric OS Administration Guide* for more information regarding individual features and commands.

To avoid failure of a firmware downgrade, verify the firmware you are downgrading to supports all the blades in the chassis, and that the switch, blades, or chassis supports all the features you are currently using. If not, you must disable or remove those features that are not supported.

Preinstallation messages

NOTE

The system messages in this section are for illustration purposes only. They do not represent the entire range of possible error messages appropriate to a wide variety of installation scenarios.

The system messages in this section may be displayed if an exception is encountered during firmware download. For example, the following message is displayed when you upgrade from Fabric OS 8.0.x to 8.2.1:

```
Cannot upgrade directly to 8.2. Please upgrade to 8.0 first and then upgrade to 8.2.
```

The following examples show feature-related messages that you may see if you were upgrading to Fabric OS 8.2.1 or downgrading from Fabric OS 8.2.1:

```
Cannot upgrade directly to 8.2. Please upgrade to 8.0 first and then upgrade to 8.2.
```

```
Cannot downgrade to 7.4 or lower. Please downgrade to 8.0 first and
then download the desired firmware version.
```

```
Non disruptive firmwaredownload is not supported when firmwaredownload with two
versions apart. Please try to use "firmwaredownload" with single mode option enabled.
```

```
Non disruptive firmwaredownload is not supported when firmwaredownload with two versions
apart. Disruptive firmware download is disallowed when auto-boot option is enabled.
Please disable auto-boot and try again.
```

```
FC8-32E blade is not supported by the targeted firmware. Please remove the blade
before upgrading.
```

```
FC8-48E blade is not supported by the targeted firmware. Please remove the blade
before upgrading.
```

```
Firmware upgrade to Fabric OS 8.0.1 or higher is not allowed when there are more than
4 chassis interconnected through Inter-Chassis Links (ICLs) and the Enterprise ICL (EICL)
license is not installed in the system. Note that even with an EICL license installed,
only 10 chassis are allowed to interconnect through ICLs. You can either install an EICL
license, or you must disable the additional ICL links before performing a firmware upgrade.
```


Logical Switches with LS instance > 7 are not supported in FOS8.0. Please delete the Logical switch with "LS instance >7" before downgrade. LS instance can be verified using the command "lscfg --show -instance."

Downgrade is not allowed because some FCIP features are enabled and are not supported on the selected version. Please address these unsupported features before downgrading.

Certificates and/or CSRs related to Extension feature are present in the switch that are not supported in the older versions. Please use "seccertmgmt delete -all extn" to delete them, before proceeding with downgrade.

Downgrade is not allowed since snmpv3 user password encryption is enabled. Please use "snmpconfig --set snmpv3 -disable passwd_encryption" to disable it.

Firmware downgrade is not allowed as some of the LSAN Zones or Devices cannot be accommodated in previous Fabric OS v8.1.0 version. Use "fcrlsancount" to configure LSAN count as 3000 or 5000. Use "lsanzoneshow --maxcapacity" to view LSAN Zones and Devices that need to be removed. Use "lsanzoneshow --remove" to generate configuration script for removal of the same.

Flow Mirror flow(s) exist. Please delete Flow Mirror flow(s) and proceed with downgrade.

Downgrade is not allowed because switch is in AG mode and Mirror port configuration is enabled on some of the ports. Please use "portcfgshow" to get port list and use "portcfg mirrorport <port_no> --disable" to disable it before downgrading.

Firmware downgrade to Fabric OS 8.0.x is not allowed because Alias Peer Zones are configured. Before downgrading, remove all alias members from all peer zones.

Firmware downgrade to Fabric OS 8.0.x is not allowed because Enhanced Zone Object Names are configured. Before downgrading, remove any zone objects containing enhanced names or modify zone object names such that they are not numeric-starting and do not contain special characters ('-', '\$', '^').

Downgrade is not allowed because discard frame logging is enabled for frame type other than timeout du unrout. Please disable those discard frame types first. {Usage: framelog --disable -type [type1miss | type2miss | type6miss]}.

Manually configured ipv6 gateway is present. Please remove the manually configured gateway ip before downgrade.

Downgrade is not allowed because ssh rekey is configured. Please remove configuration by executing "sshutil rekeyinterval 0".

Downgrade is not allowed due to the presence of FICON logical switch. Please delete the FICON logical switch using the command "lscfg --delete [-FID]" and retry.

Download is not allowed because ldap role map to root is configured. Please check and unmap the configuration by executing "ldapcfg --unmaprole <LDAP rolename>".

Non-disruptive firmware downgrade is not supported due to registered Application Server entries, see "appserver --show -domain <local domain ID>". Either disable registered devices or issue "firmwaredownload" with single mode option enabled.

Downgrade is not allowed because some Extension features are enabled and are not supported on the selected version. Please address these unsupported features before downgrading.

Downgrade is not allowed because System hash type is not set to MD5, or all user passwords are not hashed using MD5. Please configure the system hash using "passwdcfg --hash md5" and change all account passwords before proceeding. Also verify all local passwords are configured with the MD5 hash using the command "passwdcfg --showhash -all".

Downgrade is not allowed because access time range is configured for one or more user accounts. Please delete the access time configuration of the user account using "userConfig --change <username> -at 00:00-00:00"

or delete the user account using "userConfig --delete <username>". Please ensure that access time is not configured for any accounts by executing "userConfig --show -a" command before proceeding further.

Dynamic portname format is not set to default in one or more partitions. Please run "portname -d -default" in the corresponding partitions to set default portname format.

Downgrade is not allowed because one or more flows are defined on VE ports. Please delete those flows using the command "flow --delete <flow name>".

Please disable DHCPV6 option. As it is not supported in firmware < 8.0+.

vTap and QOS compatibility mode is ON. Please set this config option to OFF (configureChassis CLI) and proceed with downgrade.

Blade types

Where blades are incompatible with a firmware download, they must be removed or powered off before a firmware download begins, as noted in the following messages.

TABLE 21 Issues related to blade types

Message	Probable cause and recommended action
FC8-32E blade is not supported by the targeted firmware. Please remove the blade before upgrading.	An attempt was made to upgrade a system running Fabric OS 7.4.0 to Fabric OS 8.0.1 or downgrade a system from Fabric OS 8.1.0 to 8.0.1 with Brocade FC8-32E or FC8-48E blades installed. These blades are not supported by Fabric OS 8.0.1, so the firmware download operation will fail. If either of these blades are installed after upgrade or downgrade, the slot containing the blade will fault. NOTE These messages will not display if you are upgrading directly from Fabric OS 7.4.0 to 8.1.0 using firmwaredownload -s. Use the slotshow command to display which slots these blades occupy. Physically remove the blades from the chassis, and then retry the firmware download operation.
FC8-48E blade is not supported by the targeted firmware. Please remove the blade before upgrading.	

Firmware versions

The system messages in this section refer to differences between the current firmware and the firmware you are applying to the switch.

The following information should be considered regarding firmware upgrades and downgrades for this release:

- Disruptive direct upgrades and downgrades from Fabric OS 8.1.x to 8.2.0 are supported, but nondisruptive direct upgrades and downgrades between Fabric OS 7.4.0 or a previous release to 8.2.0 are not.
- Nondisruptive upgrades from Fabric OS 8.0.0 or 8.0.x to 8.2.1 is a two-step process: first upgrade to 8.1.x, and then upgrade to 8.2.1.
- Direct upgrades from Fabric OS 8.0.0 to 8.2.1 are possible using **firmwaredownload -s**. Note that the **-s** option is disruptive to switch traffic.

Refer to the current *Brocade Fabric OS Upgrade Guide* for details on firmware download options. The "Supported upgrade paths" section of this guide provides details on upgrading between Fabric OS 7.2.x and 8.2.1

NOTE

The firmware upgrade and downgrade guidelines described in these sections are generally applicable for most releases. Users are advised to always refer to the Fabric OS release notes for the most current upgrade and downgrade information.

The following table covers firmware version-related issues.

TABLE 22 Issues related to firmware versions

Message	Probable cause and recommended action
Cannot upgrade directly to 8.1. Please upgrade to 7.4 first and then upgrade to 8.1.	An attempt was made to upgrade to Fabric OS 8.1.0 from earlier than 7.4.0. If the switch is running Fabric OS 7.4.0 or earlier, you are not allowed to upgrade directly to 8.1.0. Upgrade your switch to Fabric OS version 7.4.0 before upgrading to 8.1.0.
Cannot downgrade to 7.3 or lower. Please downgrade to 7.4 first and then download the desired firmware version.	An attempt was made to downgrade from Fabric OS 8.1.0 to 7.3.0 or earlier. If the switch is running Fabric OS 8.1.0, you are not allowed to downgrade directly to 7.3.0 or lower. Downgrade your switch to Fabric OS 7.4.0 first, and then downgrade to the desired firmware version. You must use the disruptive firmwaredownload -s command for direct downgrades.
Non disruptive firmware download is not supported for firmwaredownload with two versions apart. Disruptive firmware download is disallowed when auto-boot option is enabled. Please disable auto-boot and try again.	An attempt was made to upgrade or downgrade between Fabric OS 8.1.0 and 7.4.0 using the auto-boot option (-b). In other words, firmwaredownload -sb was entered. Download firmware again using firmwaredownload -st to disable auto-boot mode and allow a disruptive download.

Security

• User account management.....	60
• Password issues.....	60
• Device authentication	61
• Protocol and certificate management	61
• SNMP issues.....	62
• FIPS	62

User account management

The most common error when managing user accounts is not setting up Role-Based Access Control (RBAC). Errors such as a user not being able to run a command or modify switch settings are usually related to what role the user has been assigned. If a user is unable to modify switch settings, check the RBAC permissions for that user.

Password issues

The following table describes various ways to recover forgotten passwords.

TABLE 23 Issues related to passwords

Symptom	Probable cause and recommended action
User forgot password	If you know the root password and your device ships with a root account, you can use this procedure to recover the password for the default accounts of user and admin. If you do not know the root password, you must contact your service support provider to recover admin passwords. To recover a password, complete the following steps. 1. Open a CLI connection (serial or Telnet) to the switch. 2. Log in as root. 3. Enter the command for the type of password that was lost: device:admin> passwd user device:admin> passwd admin 4. Enter the requested information at the prompts.
Unable to log in as root	To recover your root password, contact your switch service provider. Not all devices ship with a root account; if your device did not ship with a root account, one cannot be created.
Unable to log into the boot PROM	To recover a lost boot PROM password, contact your switch service provider. You must have previously set a recovery string to recover the boot PROM password. This does not work on lost or forgotten passwords in the account database.

Password recovery options

The following table describes the options available when one or more types of passwords are lost.

TABLE 24 Password recovery options

Topic	Solution
If all the passwords are forgotten, what is the password recovery mechanism? Are these procedures non-disruptive recovery procedures?	Contact your switch service provider. A non-disruptive procedure is available.
If a user has only the root password, what is the password recovery mechanism?	Use the passwd command to set other passwords. Use the passwddefault command to set all passwords to default.
How to recover boot PROM password?	Contact your switch service provider and provide the recovery string.
How to recover a user, or admin password?	Refer to Password issues on page 60 for more information on recovering these passwords.

Device authentication

The following table covers device authentication issues.

TABLE 25 Issues related to device authentication

Symptom	Probable cause and recommended action
Switch is unable to authenticate device.	Issue: When the device authentication policy is set to ON, the switch expects a FLOGI with the FC-SP bit set. If this bit is not set, the switch rejects the FLOGI with reason LS_LOGICAL_ERROR (0x03), in the switch log with the explanation of "Authentication Required"(0x48), and disables the port. Solution: Set the device authentication policy mode on the switch to ON.
Switch is unable to form an F_Port.	Regardless of the device authentication policy mode on the switch, the F_Port is disabled if the DH-CHAP protocol fails to authenticate. If the HBA sets the FC-SP bit during FLOGI and the switch sends a FLOGI accept with FC-SP bit set, then the switch expects the HBA to start the AUTH_NEGOTIATE. From this point on until the AUTH_NEGOTIATE is completed, all ELS and CT frames, except the AUTH_NEGOTIATE ELS frame, are blocked by the switch. During this time, the Fibre Channel driver rejects all other ELS frames. The F_Port does not form until the AUTH_NEGOTIATE is completed. It is the HBA's responsibility to send an Authentication Negotiation ELS frame after receiving the FLOGI accept frame with the FC-SP bit set.

Protocol and certificate management

This section provides information and procedures for troubleshooting standard Fabric OS security features such as protocol and certificate management.

Troubleshooting certificates. If you receive messages in the browser or in a pop-up window when logging in to the target switch using HTTPS, refer to [Table 26](#) for recommended actions you can take to correct the problem.

The following table covers SSL messages and related actions.

TABLE 26 SSL messages and actions

Message	Action
The page cannot be displayed	Issue: The SSL certificate is not installed correctly or HTTPS is not enabled correctly. Solution: Make sure that the certificate has not expired, that HTTPS is enabled, and that certificate file names are configured correctly.
The security certificate was issued by a company you have not chosen to trust.	The certificate is not installed in the browser. Install it as described in the <i>Brocade Fabric OS Administration Guide</i> .
The security certificate has expired or is not yet valid	Issue: Either the certificate file is corrupted or it needs to be updated.

TABLE 26 SSL messages and actions (continued)

Message	Action
	Solution: Click View Certificate to verify the certificate content. If it is corrupted or out of date, obtain and install a new certificate.
The name on the security certificate is invalid or does not match the name of the site file	Issue: The certificate is not installed correctly in the Java Plug-in. Solution: Install it as described in the <i>Brocade Fabric OS Administration Guide</i> .
This page contains both secure and nonsecure items. Do you want to display the nonsecure items?	Issue: The page contains both secure and nonsecure items. Solution: Click No in this pop-up window. The session opens with a closed lock icon on the lower-right corner of the browser, indicating an encrypted connection.

Gathering additional information

For security-related issues, use the following guidelines to gather additional data for your switch support provider.

- Use the **supportSave -n** command.
- If not sure about the problem area, use the **supportSave -n** to collect data from all switches in the fabric.
- If you think it may be related to E_Port authentication, use the **supportSave -n** command to collect data from both switches of the affected E_Port.
- If you think this is a policy-related issue, FCS switch, or other security server-related issue, then use **supportSave -n** to collect data from the Primary FCS switch and all affected switches.
- If login-related, then also include the following information:
 - Does login problem appear on a Serial, CP IP, or Switch IP address connection?
 - Is it CPO or CP1?
 - Is the CP in active or standby?
 - Is it the first time login after **firmwareDownload** and reboot?

SNMP issues

This table describes symptoms with associated causes and recommended actions for SNMP-related issues.

TABLE 27 Issues related to SNMP

Symptom	Probable cause and recommended action
SNMP management station server is unable to receive traps from fabric.	There are several causes related to this generic issue. You must verify the following: • There are no port filters in the firewalls between the fabric and the SNMP management station. • If your SNMP management station is a dual-homed server, check that the routing tables are set up correctly for your network. If you continue to have problems, run supportsave, and gather that output along with a MIB browser snapshot with the problem (like Adventnet screen snapshot) for an MIB variable and contact your switch support provider.

FIPS

This table describes symptoms with associated causes and recommended actions for problems related to FIPS.

TABLE 28 Issues related to FIPS

Symptom	Probable cause and recommended action
Prior to Fabric OS 8.2.1, When FIPS is turned on, the switch constantly reboots.	When FIPS is turned on, the switch runs conditional tests each time it is rebooted. These tests run random number generators and are executed to verify the randomness of the random number generator. The conditional tests are executed each time prior to using the random number provided by the random number generator. The results of all self-tests, for both power-up and conditional, are recorded in the system log or are output to the local console. This includes logging both passing and failing results. If the tests fail on your switch, it may go into a reboot loop because the cryptography of the switch can no longer be validated or trusted. In this case the switch would have to be recovered by a reinstall of Fabric OS.

Virtual Fabrics

• General Virtual Fabrics Troubleshooting.....	64
• Fabric identification issues.....	65
• Logical Fabric issues.....	65
• Base switch issues.....	65
• Logical switch issues.....	66
• Switch configuration blade compatibility.....	67
• Gathering additional information.....	67

General Virtual Fabrics Troubleshooting

All of the following constraints apply when the Virtual Fabrics feature is enabled:

- The base fabric works only in Brocade native mode, not in an interoperable mode.
- The base switch does not have any devices. The base fabric can have devices in remote Layer 2 switches; traffic between those devices is supported.
- Only 12 base switches should be configured per Virtual Fabric in Fabric OS 8.0.1 and later.
- Since an individual base switch cannot have devices, the entire base fabric does not have any devices as well. The devices can be in other logical partitions that are configured on the same switch or remote switches.
- A nonbase switch in a Virtual Fabrics-capable chassis must not be part of a fabric that serves as a base fabric for some other logical fabric traffic. Although software does not detect or prevent users from deploying such a configuration, such a configuration is not supported.
- ICL ports can be in the base or default switch only. If the extended ISL (XISL) is turned off, you can connect ICLs to other logical switches.
- A default switch can be configured as a base switch in fixed-port switches, but not in a Brocade backbone. Fabric IDs of default switches cannot be manually changed.
- The default switch can participate in a logical fabric using (XISLs). In Brocade backbones, the default switch does not participate in a logical fabric and is a purely Layer 2 logical switch.
- EX_Ports and VEX_Ports are supported in the base switch. EX_Ports cannot be part of any other switch other than the base switch.
- EX_Ports and VEX_Ports cannot connect to a fabric that has a logical switch with the Allow XISL use mode on. The port is disabled with the reason "Conflict: XISL capability domain."
- External device sharing is supported through EX_Ports. Internal device sharing (sharing a device in a logical fabric with other fabrics without having an EX_Port) is not supported.
- A logical fabric cannot have EX_Ports that use XISLs and cannot serve as a backbone to any EX_Port traffic. Similarly, the default switch cannot be part of a fabric that serves as a backbone to any EX_Port traffic.
- Traffic Isolation (TL) zones with no failover option are not supported in logical fabrics. TI zones defined in the base fabric for logical fabric traffic must allow failover.

NOTE

The **configure** command has a "Disable FID check" fabric parameter option, which can be used to disable the FID check for FICON logical switches.

Fabric identification issues

The following table covers fabric identification issues.

TABLE 29 Issues related to fabric identification

Symptom	Probable cause and recommended action
E_Ports directly connecting two logical switches do not form, or are disabled.	The FIDs on each of the logical switches must be the same. Enter lscfg --show to view the current FIDs on the chassis and then enter lscfg --change -newfid to change the FID.
Invalid FID.	FIDs for switches may be from 1 through 128 as long as they are not already in use, except for EX_Ports, which are only assigned FIDs from 1 through 127. Enter lscfg --show to verify whether the FID is in use. If it is, use another FID.
The FID is currently in use.	You may not create two or more logical switches with the same FID. Use the lscfg --show and fc fabrics show commands to view FIDs in use

Logical Fabric issues

The following table covers logical fabric issues.

TABLE 30 Issues related to logical fabrics

Symptom	Probable cause and recommended action
Logical port <port_number> disabled.	This message indicates an LISL was disabled because of some protocol conflict or security or policy violation. This can result in possible traffic issues. You should resolve the cause of the conflict and re-enable the LISL by entering lfcfg --lislenable .
The switch with domain <domain> with firmware version <fw_version> has joined the FID <fid> fabric and may not be compatible with XISL use.	This message indicates the specified switch in the logical fabric that is using XISLs is running an incompatible firmware version and must be upgraded to Fabric OS 6.2.0 or later.
Logical port segmented	Ensure that the fabwide policy, zoning, and defzone settings are consistent on the logical switches intending to be joined.

Base switch issues

All logical switches in a fabric should have the same base switch attribute. If a base switch is connected to a non-base switch, then you must take the appropriate action to resolve the conflict.

The following table covers base switch-related issues.

TABLE 31 Issues related to base switches

Symptom	Probable cause and recommended action
EX_Port is disabled with reason "Export in non base switch".	An EX_Port must be in the base switch. 1. Enter lscfg --create -b to create a base switch. 2. Enter lscfg --config -slot -port to move the port to the base switch. 3. If the port is not intended to be used as an EX_Port, enter portcfgdefault to reset the port to its default configuration.

TABLE 31 Issues related to base switches (continued)

Symptom	Probable cause and recommended action
An EX_ or VEX_Port is disabled with reason "Conflict: XISL capability domain".	Use the configure command to set "Allow XISL use" to OFF on all logical switches of the connecting edge fabric.
E_Ports connecting two logical switches are disabled.	If a base switch is directly connected to a non-base switch, all E_Ports to that logical switch are disabled.
Fabric ID and base switch are conflicted.	If there is a Fabric ID conflict and a base switch conflict that exists between two switches, the Fabric ID conflict is detected first. Enter lscfg --change-newfid to change the FID.
A base switch already exists on this system.	Only one base switch is allowed on a platform. To remove the current base switch and then create a new one, complete the following steps. Enter lscfg --delete . Enter lscfg --create -b .

Logical switch issues



CAUTION

When a logical switch is created, all configuration for the logical switch is set to factory defaults. When a logical switch is deleted, all configuration for the logical switch is deleted permanently and is not recoverable.

The following table covers logical switch issues.

TABLE 32 Issues related to logical switches

Symptom	Probable cause and recommended action
The indicated slot is empty.	When lscfg was run, an empty slot was specified. Re-issue the command with the appropriate slot number.
Validation of switch configuration changes is not supported on this platform.	This platform is unknown to the logical switch subsystem.
Given slot number is not valid on this platform.	You are specifying a slot number that is not valid on the platform, for example, slot 0 on a Brocade X6-8 or slot 12 on a Brocade X6-4.
Slot must be enabled to configure ports.	You may only attempt to configure ports on enabled blades (blades may be faulted).
Unable to determine slot type.	The slot type is not known to the logical switch. Verify the slot and try again.
There are no ports on this slot	There are no configurable ports on the slot indicated by the lscfg command. Verify the ports and try again.
Unable to remove ports from their current switch.	When moving ports to a switch, you must first remove them from the switch in which they reside. This error message is displayed if this step fails.
A non-GE blade is within the slot range.	You are attempting to configure a GigabitEthernet (GE) port on a slot that does not contain GE ports.
A port or ports is already in the current switch.	You may not move a port to the same switch.
The maximum number of switches for this platform has been reached.	Each platform that supports Virtual Fabrics has a maximum number of logical switches that may be supported. The platform has reached this limit. Director platforms support eight logical switches and fixed-port switches support four logical switches when running Fabric OS 8.0.1 or later.
Unable to create the switch.	There was an error while creating the switch.
A port or ports cannot be moved to the requested switch because it	The area limit would be exceeded if the lscfg command were allowed.

TABLE 32 Issues related to logical switches (continued)

Symptom	Probable cause and recommended action
would exceed the 256 area limit for this switch.	
A port or ports cannot be moved to the requested switch because it may only exist in a base or default switch.	You are attempting to move ports on a core blade into a non-default or non-base switch.

Switch configuration blade compatibility

The following table covers switch configuration blade compatibility issues.

TABLE 33 Issues related to switch configuration blade compatibility

Symptom	Probable cause and recommended action																																																							
A slot in the chassis displays a FAULTY(91) in the slotshow output.	When an enterprise-class platform is coming up or when a blade is inserted, the switch configuration is checked based on the blade type. If the configuration does not match with the blade type, the blade is faulted. This is displayed as FAULTY(91) in the output of the slotshow command. Enter lscfg -restore_slot_to_default to correct the problem. Once the configuration discrepancy has been fixed, enter slotpoweroff followed by slotpoweron to recover. The following example shows this procedure. <pre>device:admin> lscfg -restore_slot_to_default 1 device:admin>slotpoweroff 1 device:admin>slotpoweron 1</pre> <table><thead><tr><th>Slot</th><th>Blade Type</th><th>ID</th><th>Model Name</th><th>Status</th></tr></thead><tbody><tr><td>2</td><td>SW BLADE</td><td>97</td><td>FC16-32</td><td>ENABLED</td></tr><tr><td>3</td><td>SW BLADE</td><td>96</td><td>FC16-48</td><td>ENABLED</td></tr><tr><td>4</td><td>SW BLADE</td><td>96</td><td>FC16-48</td><td>ENABLED</td></tr><tr><td>5</td><td>CORE BLADE</td><td>98</td><td>CR16-8</td><td>ENABLED</td></tr><tr><td>6</td><td>CP BLADE</td><td>50</td><td>CP8</td><td>ENABLED</td></tr><tr><td>7</td><td>CP BLADE</td><td>50</td><td>CP8</td><td>ENABLED</td></tr><tr><td>8</td><td>CORE BLADE</td><td>98</td><td>CR16-8</td><td>ENABLED</td></tr><tr><td>10</td><td>SW BLADE</td><td>96</td><td>FC16-48</td><td>ENABLED</td></tr><tr><td>11</td><td>AP BLADE</td><td>75</td><td>FX8-24</td><td>ENABLED</td></tr><tr><td>12</td><td>SW BLADE</td><td>97</td><td>FC16-32</td><td>ENABLED</td></tr></tbody></table>	Slot	Blade Type	ID	Model Name	Status	2	SW BLADE	97	FC16-32	ENABLED	3	SW BLADE	96	FC16-48	ENABLED	4	SW BLADE	96	FC16-48	ENABLED	5	CORE BLADE	98	CR16-8	ENABLED	6	CP BLADE	50	CP8	ENABLED	7	CP BLADE	50	CP8	ENABLED	8	CORE BLADE	98	CR16-8	ENABLED	10	SW BLADE	96	FC16-48	ENABLED	11	AP BLADE	75	FX8-24	ENABLED	12	SW BLADE	97	FC16-32	ENABLED
Slot	Blade Type	ID	Model Name	Status																																																				
2	SW BLADE	97	FC16-32	ENABLED																																																				
3	SW BLADE	96	FC16-48	ENABLED																																																				
4	SW BLADE	96	FC16-48	ENABLED																																																				
5	CORE BLADE	98	CR16-8	ENABLED																																																				
6	CP BLADE	50	CP8	ENABLED																																																				
7	CP BLADE	50	CP8	ENABLED																																																				
8	CORE BLADE	98	CR16-8	ENABLED																																																				
10	SW BLADE	96	FC16-48	ENABLED																																																				
11	AP BLADE	75	FX8-24	ENABLED																																																				
12	SW BLADE	97	FC16-32	ENABLED																																																				

Gathering additional information

For Virtual Fabrics-related issues, use the following guidelines to gather additional data for your switch support provider:

- Use the **supportSave** command.
- If not sure about the problem area, use the **supportSave** command on all chassis and logical switches in the fabric.
- If you think it may be related to E_Port authentication, then use the **supportSave -n** command on both switches or logical switches of the affected E_Port.

ISL Trunking

• Trunking Link Issues.....	68
• Buffer credit issues.....	69

Trunking Link Issues

This section describes trunking link issues that can arise and recommended actions to correct the problems.

You can use the **islshow** command to find the incompatible and segmented ISL links. To know the reason for incompatibility and segmentation, use the **switchshow** command on both ends of the ISL link.

The following table covers trunking link-related issues.

TABLE 34 Issues related to trunking links

Symptom	Probable cause and recommended action
A link that is part of an ISL trunk failed.	Use the trunkdebug command to troubleshoot the problem, as shown in the following procedure. 1. Connect to the switch and log in as admin. 2. Enter trunkdebug port_1 port_2. • port1 - The area number or index of port 1. • port2 - The area number or index of port 2. Use the switchshow command to determine the area or index numbers for the port values. Both operands are required. This example shows that port 3 is not configured as an E_Port. <pre>device:admin> trunkdebug 126, 127 port 126 is not E/EX port port 127 is not E/EX port</pre> This example shows that port 3 is configured as an E_Port. <pre>device:admin> trunkdebug 100, 101 port 100 and 101 connect to the switch 10:00:00:05:1e:34:02:45</pre> The trunkdebug command displays the possible reason that two ports cannot be trunked. Possible reasons are: • The switch does not support trunking. • A trunking license is required. • Trunking is not supported in switch interoperability mode. • Port trunking is disabled. • The port is not an E_Port. • The port is not 2 Gbps, 4 Gbps, 8 Gbps, 10 Gbps, or 16 Gbps. • The port connects to a switch other than the one you want. To correct this issue, connect additional ISLs to the switch with which you want to communicate. • The ports are not the same speed or they are not set to an invalid speed. Manually set port speeds to a speed supported on both sides of the trunk. • The ports are not set to the same long distance mode. Set the long distance mode to the same setting on all ports on both sides of the trunk. • Local or remote ports are not in the same port group.

TABLE 34 Issues related to trunking links (continued)

Symptom	Probable cause and recommended action
	Move all ISLs to the same port group. The port groups begin at port 0 and are in groups of four or eight, depending on the switch model. Until this is done, the ISLs do not trunk. The difference in the cable length among trunked links is greater than the allowed difference.

Buffer credit issues

This table describes the issues related to a trunk bouncing online and offline or hosts not being able to talk to a storage device.

TABLE 35 Issues related to buffer credits

Symptom	Probable cause and recommended action
Trunk goes offline and online (bounces).	A port that is disabled at one end because of buffer underallocation causes all disabled ports at the other end to become enabled. Some of these enabled ports become disabled because of a lack of buffers, which in turn triggers ports to be enabled once again at the other end. While the system is stabilizing the buffer allocation, it warns that ports are disabled because of a lack of buffers, but it does not send a message to the console when buffers are enabled. The system requires a few passes to stabilize the buffer allocation. Ultimately, the number of ports for which buffers are available come up and stabilize. You should wait for stabilization and then proceed with correcting the buffer allocation situation.

Getting out of buffer-limited mode

Occurs on LD_Ports.

1. Change the LD port speed to a lower speed (of non-buffer limited ports).
2. Change the LD port's estimated distance to a shorter distance (of non-buffer limited ports).
3. Change LD back to L0 (of non-buffer limited ports).
4. If you are in buffer-limited mode on the LD port, then increase the estimated distance.
5. Enable any of these changes on the buffer-limited port or switch by issuing the commands **portDisable** and **portEnable**.

Zoning

• Zoning Corrective Action.....	70
• Segmented fabrics.....	71
• Zone conflicts.....	72
• Gathering additional information.....	76

Zoning Corrective Action

The following overview provides a basic starting point to troubleshoot your zoning problem.

1. Verify that you have a zone problem.
2. Determine the nature of the zone conflict.
3. Take the appropriate steps to correct the zone conflict.

To correct a merge conflict without disrupting the fabric, first verify that it was a fabric merge problem, then edit the zone configuration members, and then reorder the zone member list if necessary.

The newly changed zone configuration is not effective until you issue the **cfgEnable** command. This should be done during a maintenance window because it may cause disruption in large fabrics.

Verifying a fabric merge problem

1. Enter the **switchShow** command to validate that the segmentation is due to a zone issue.
2. Review [Segmented fabrics](#) on page 71 to view the different types of zone discrepancies and determine what might be causing the conflict.

Verifying a TI zone problem

Use the **zone --show** command to display information about Traffic Isolation (TI) zones. This command displays the following information for each zone:

- Zone name
- E_Port members
- N_Port members
- Configured status (the latest status, which may or may not have been activated by **cfgEnable**)
- Enabled status (the status that has been activated by **cfgEnable**)

If you enter the **cfgShow** command to display information about all zones, the TI zones appear in the defined zone configuration only and do not appear in the effective zone configuration.

1. Connect to the switch and log in as admin.

- Enter the **zone --show** command where *name* is the name of the zone to be displayed. If the name is omitted, the command displays information about all TI zones in the defined configuration.

```
zone --show [name]
```

To display information about the TI zone purplezone, use the following command:

```
switch:admin> zone --show purplezone
Defined TI zone configuration:
TI Zone Name:   redzone:
Port List:      1,2; 1,3; 3,3; 4,5
Configured Status: Activated / Failover-Enabled
Enabled Status: Activated / Failover-Enabled
```

To display information about all TI zones in the defined configuration, use the following command:

```
switch:admin> zone --show
Defined TI zone configuration:
TI Zone Name:   greenzone:
Port List:      2,2; 3,3; 5,3; 4,11;
Configured Status: Activated / Failover-Enabled
Enabled Status: Activated / Failover-Enabled
TI Zone Name:   purplezone:
Port List:      1,2; 1,3; 3,3; 4,5;
Configured Status: Activated / Failover-Enabled
Enabled Status: Deactivated / Failover-Enabled
TI Zone Name:   bluezone:
Port List:      9,2; 9,3; 8,3; 8,5;
Configured Status: Deactivated / Failover-Disabled
Enabled Status: Activated / Failover-Enabled
```

- Run the **zoneshow --validate** command to display any errors in the zone configuration such as misspelled WWN names.

Correct the zone configuration errors before proceeding further.

Segmented fabrics

This section discusses fabric segmentation. Fabric segmentation occurs when two or more switches are joined by ISLs and do not communicate to each other. Each switch appears as a separate fabric when you use the **fabricshow** command.

TABLE 36 Issues related to fabric segmentation

Symptom	Probable cause and recommended action
Port gets segmented due to a zone conflict.	Occurs even when one of the following is true: - Zoning is enabled in both fabrics, and the zone configurations are different in each fabric. - The name of a zone object in one fabric is also used for a different type of zone object in the other fabric. A zone object is any device in a zone. - The definition in one fabric is different from the definition of a zone object with the same name in the other fabric. To resolve the conflict, ensure that the following are true: - The effective cfg (zone set) on each end of the segmented ISL is identical. - Any zone object with the same name has the same entries in the same sequence.
Fabric segmentation is caused by a "configuration mismatch."	Occurs when zoning is enabled in both fabrics and the zone configurations are different in each fabric.
Fabric segmentation is caused by a "type mismatch."	Occurs when the name of a zone object in one fabric is also used for a different type of zone object in the other fabric. A zone object is any device in a zone.
Fabric segmentation is caused by a "content mismatch."	Occurs when the definition in one fabric is different from the definition of a zone object with the same name in the other fabric.

Zone conflicts

Zone conflicts can be resolved by saving a configuration file with the **configupload** command, examining the zoning information in the file, and performing a cut-and-paste operation so that the configuration information matches in the fabrics being merged.

After examining the configuration file, you can choose to resolve zone conflicts by entering **cfgdisable** followed by **cfgclear** on the incorrectly configured segmented fabric and then entering **cfgsave** followed by **portdisable** and **portenable** on one of the ISL ports that connects the fabrics. This causes a merge, making the fabric consistent with the correct configuration.

ATTENTION

Be careful using the **cfgclear** command because it deletes the defined configuration.

Table 37 summarizes commands that are useful for debugging zoning issues.

TABLE 37 Commands for debugging zoning

Command	Function
aliadd	Use to add members to a zone alias.
alcreate	Use to create a zone alias.
aldelete	Use to delete a zone alias.
alremove	Use to remove members from a zone alias.
alishow	Use to display zone aliases.
cfgadd	Use to add zone members to a zone configuration.
cfgcreate	Use to create a zone configuration.
cfgclear	Use to clear the zoning database.
cfgremove	Use to remove zone members from a zone configuration.
cfgshow	Displays zoning configuration.
cfgsize	Use to display the zone database size.
cfgdelete	Use to delete the zone configuration.
cfgdisable	Disables the active (effective) configuration
cfgenable	Use to enable and activate (make effective) the specified configuration.
cfgsave	Use to save changes to the zone configuration database.
cfgtransabort	Use to abort the current zoning transaction without committing it.
cfgtransshow	Use to display the ID of the current zoning transaction.
defzone	Sets the default zone access mode to No Access, initializes a zoning transaction (if one is not already in progress), and creates the reserved zoning objects.
licenseshow	Displays current license keys and associated (licensed) products.
switchshow	Displays currently enabled configuration and any E_Port segmentations resulting from zone conflicts.
zoneadd	Use to add a member to an existing zone.
zonecreate	Use to create a zone. Before a zone becomes active, the cfgsave and cfgenable commands must be used.
zoneddelete	Use to delete zones.
zonehelp	Displays help information for zone commands.
zoneremove	Use to remove zone member aliases from zone database.
zonestow	Displays zone information.

For more information about setting up zoning on your switch, refer to the *Brocade Fabric OS Administration Guide*.

ATTENTION

The **cfgclear** command is a disruptive procedure.

Resolving zoning conflicts

When merging two fabrics, multiple zoning CLI sessions can be launched on the same switch or on different switches. This section describes these situations and how they are automatically resolved.

- **Dual-CLI sessions from the same switch** If you start a zone transaction from CLI-Session1 and then try to perform a zone modification from CLI-Session2, the CLI-Session2 zone transaction is not allowed due to CLI-Session2 not being the owner of the open transaction. If CLI-Session1 logs out, this ends the open transaction and aborts any current zone modifications. CLI-Session2 is then able to perform zone modifications. The zone transaction locking mechanism works on a single switch from the CLI perspective and there is no dangling transaction.
- **Dual-CLI sessions from different switches** If you start a CLI zone transaction on Switch1 and start another CLI zone transaction on Switch2, when committing the zone transaction from Switch1, the open zone transaction from Switch2 is aborted by Switch1. The following message is posted on Switch2 at the time of zone commit from Switch1:

```
2012/03/09-21:45:26, [ZONE-1027], 3285, FID 128, INFO, sw0, Zoning transaction aborted Zone Config
update Received
```

Correcting a fabric merge problem quickly

Brocade recommends that if there are devices attached to the switches on which you need to perform these recovery steps, ensure that these switches are in DefZone NoAccess so as not to cause any RSCN storms and bog down the NameServer.

1. Determine which switches have the incorrect zoning configuration; then log in to the switches as admin.
2. Enter the **switchDisable** command on all problem switches.
3. Enter the **cfgDisable** command on each switch.
4. Enter the **cfgClear** command on each switch.
5. Enter the **cfgSave** command on each switch to commit the change.

ATTENTION

The **cfgClear** command clears the zoning database on the switch where the command is run.

6. Enter the **switchEnable** command on each switch once the zoning configuration has been cleared.

This forces the zones to merge and populates the switches with the correct zoning database. The fabrics then merge.

Changing the default zone access

A switch is not allowed to merge with the another switch that does not have an active effective configuration and if the default zone modes do not match. If both switches have an empty configuration but mismatched default zone access, you must pick a default zone (not necessarily "all access") and make all switches match before retrying the merge. If there is an effective configuration, you must change the default zone setting of the new switch to the default zone setting of the existing fabric switch. When the default zone "no access" option is enabled and the active configuration is disabled by using the **cfgDisable** command, a special hidden configuration with no members is activated. This configuration does not allow the switch to merge with switches that have mismatching default zone access modes.

1. Connect to the switch and log in using an account with admin permissions.
2. Display the current setting using the **defZone --show** command.

3. If your default zone is set to "no access," use the **defZone --allaccess** command to change the default zone.
4. Enter the **cfgSave** command to save the current configuration.

Editing zone configuration members

1. Log in as admin to one of the switches in a segmented fabric.
2. Enter the **cfgShow** command and print the output.
3. Start another Telnet session and connect to the next fabric as an admin.
4. Enter the **cfgShow** command and print the output.
5. Compare the two fabric zone configurations line by line and look for an incompatible configuration.
6. Connect to one of the fabrics.
7. Run zone configure commands to edit the fabric zone configuration for the segmented switch (see [Zone conflicts](#) on page 72 for specific commands).

If the zoneset members between two switches are not listed in the same order in both configurations, the configurations are considered a mismatch, resulting in the switches being segmented in the fabric.

For example:

[cfg1 = z1; z2] is different from [cfg1 = z2; z1], even though the members of the configuration are the same.

One simple approach to making sure that the zoneset members are in the same order is to keep the members in alphabetical order.

Reordering the zone member list

1. Obtain the output from the **cfgShow** command for both switches.
2. Compare the order in which the zone members are listed. Members must be listed in the same order.
3. Rearrange zone members so that the configuration for both switches is the same. Arrange zone members in alphabetical order, if possible.

Checking for Fibre Channel connectivity problems

Enter the **fcPing** command to:

- Generate an Extended Link Service (ELS) frame ECHO request to the source port specified, and validate the response.
- Generate an ELS ECHO request to the destination port specified, and validate the response.

Regardless of the device's zoning, the **fcPing** command sends the ELS frame to the destination port. A device can take any of the following actions:

- Send an ELS Accept to the ELS request.
- Send an ELS Reject to the ELS request.
- Ignore the ELS request.

Some devices do not support the ELS ECHO request. In these cases, the device either does not respond to the request or sends an ELS reject. When a device does not respond to the ELS request, further debugging is required; however, do not assume that the device is not connected to the Fibre Channel.

The following example is sample output from the **fcPing** command in which one device accepts the request and another device rejects the request:

```
switch:admin> fcping 10:00:00:00:c9:29:0e:c4 21:00:00:20:37:25:ad:05
Source:      10:00:00:00:c9:29:0e:c4
Destination: 21:00:00:20:37:25:ad:05
Zone Check:  Not Zoned
Pinging 10:00:00:00:c9:29:0e:c4 [0x20800] with 12 bytes of data:
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1162 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1013 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1442 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1052 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1012 usec
5 frames sent, 5 frames received, 0 frames rejected, 0 frames timeout
Round-trip min/avg/max = 1012/1136/1442 usec
Pinging 21:00:00:20:37:25:ad:05 [0x211e8] with 12 bytes of data:
Request rejected
Request rejected
Request rejected
Request rejected
Request rejected
5 frames sent, 0 frames received, 5 frames rejected, 0 frames timeout
Round-trip min/avg/max = 0/0/0 usec
```

The following example is sample output from the **fcPing** command in which one device accepts the request and another device does not respond to the request:

```
switch:admin> fcping 0x020800 22:00:00:04:cf:75:63:85
Source:      0x020800
Destination: 22:00:00:04:cf:75:63:85
Zone Check:  Zoned
Pinging 0x020800 with 12 bytes of data:
received reply from 0x020800: 12 bytes time:1159 usec
received reply from 0x020800: 12 bytes time:1006 usec
received reply from 0x020800: 12 bytes time:1008 usec
received reply from 0x020800: 12 bytes time:1038 usec
received reply from 0x020800: 12 bytes time:1010 usec
5 frames sent, 5 frames received, 0 frames rejected, 0 frames timeout
Round-trip min/avg/max = 1006/1044/1159 usec
Pinging 22:00:00:04:cf:75:63:85 [0x217d9] with 12 bytes of data:
Request timed out
Request timed out
Request timed out
Request timed out
Request timed out
5 frames sent, 0 frames received, 0 frames rejected, 5 frames timeout
Round-trip min/avg/max = 0/0/0 usec
```

Checking for zoning problems

1. Enter the **cfgActvShow** command to determine if zoning is enabled.
 - If zoning is enabled, the problem may be caused by zoning enforcement (for example, two devices in different zones cannot detect each other).
 - If zoning is disabled, check the default zone mode by entering the **defZone --show** command. If it is "no access", change it to "all access". To modify the default zone mode from "no access" to "all access", enter the **defZone --allaccess** command and then the **cfgSave** command.

2. Confirm that the specific edge devices that must communicate with each other are in the same zone.
 - If they are not in the same zone and zoning is enabled, proceed to Step 3.
 - If they are in the same zone, perform the following tasks:
 - Enter the **portCamShow** command on the host port to verify that the target is present.
 - Enter the **portCamShow** command on the target.
 - Enter the **nsZoneMember** command with the port ID for the zoned devices on the host and target to determine whether the name server is aware that these devices are zoned together.
3. Resolve zoning conflicts by putting the devices into the same zoning configuration.
4. Verify that no configuration is active by using the **cfgActvShow** command. Enter the **defZone --show** command to display the current state of the zone access mode and the access level. The **defZone** command sets the default zone access mode to No Access.

```
switch:admin> defzone --show
Default Zone Access Mode
committed - No Access
transaction - No Transaction
```

Refer to [Zone conflicts](#) on page 72 for additional information.

Gathering additional information

Collect the data from the **supportSave -n** command. Then collect the data from the **cfgTransShow** command. For the problematic port, collect the data from the **filterPortShow** command.

Diagnostic Features

• Fabric OS diagnostics.....	77
• Diagnostic information.....	77
• Power-on self-test.....	78
• Switch status.....	80
• Using the spinFab and portTest commands.....	83
• Port information.....	86
• Equipment status.....	89
• System message log.....	90
• Port log.....	91
• Syslogd configuration.....	93
• Automatic trace dump transfers.....	94
• Multiple trace dump files support.....	96

Fabric OS diagnostics

The purpose of the diagnostic subsystem is to evaluate the integrity of the system hardware.

Diagnostics are invoked in the one of the following ways:

- Automatically during the power-on self-test (POST).
- Automatically on an individual blade whenever it is installed into a director chassis.
- Manually using Fabric OS CLI commands.

The error messages generated during these test activities are sent to the serial console and system message logs (output formats may differ slightly).

Use the **diagHelp** command to receive a list of all available diagnostic commands.

Diagnostic information

On the switch, you can enter the **supportShow** command to dump important diagnostic and status information to the session screen, where you can review it or capture its data. If you are using a Telnet client, you may have to set up the client to capture the data prior to opening the session.

Most information can be captured using the **supportSave** command and downloaded by FTP off the switch, but when you are collecting information from specialized commands, such as **supportShow**, this information must be captured using a Telnet client.

To save a set of files that customer support technicians can use to further diagnose the switch condition, enter the **supportSave** command. The command prompts for an FTP server, packages the following files, and sends them to the specified server:

- The output of the **supportShow** command.
- Any core files, panic dumps or FFDC files that may have been generated.
- System message (RAS) logs.
- Other special feature logs.

Refer to [Automatic trace dump transfers](#) on page 94.

Power-on self-test

By default, when you power on the system, the boot loader automatically performs power-on self-tests and loads a Fabric OS kernel image. Likewise, if you issue the **slotPowerOn** command, or insert a new blade, a power-on self test is run on that blade.

The POST tests provide a quick indication of hardware readiness when hardware is powered up. These tests do not require user input to function. They typically operate within several minutes, and support minimal validation because of the restriction on test duration. Their purpose is to give a basic health check before a new switch joins a fabric.

These tests are divided into two groups: POST1 and POST2. POST1 validates the hardware interconnect of the device, and POST2 validates the ability of the device to pass data frames between the ports. The specific set of diagnostic and test commands run during POST depends on the switch model.

NOTE

When POST2 is running portloopbacktest mode 8 or mode 11 and if the peer blades are powered off while the test is running, the POST is expected to fail in the core blades as mode 8 and mode 11 tests the peer links and the peer end is no longer active.

NOTE

Gen 5 16-Gbps switches and chassis can take several minutes to complete the POST after it is powered on.

NOTE

The POST2 will run the test for both Fibre Channel and Ethernet frames on Brocade G630 switch and Brocade FC32-64 blade.

On Gen 6 directors, POST2 will run mode 11 (backend link test) from the Core blades only.

You can use the **diagDisablePost** command to disable both POST1 and POST2, and you can re-enable POST1 and POST2 using the **diagEnablePost** command.

The following example shows a typical boot sequence, including POST messages:

```
The system is coming up, please wait...
Read board ID of 0x80 from addr 0x23
Read extended model ID of 0x16 from addr 0x22
Matched board/model ID to platform index 4
PCI Bus scan at bus 0
:      :
:      :
Checking system RAM - press any key to stop test
Checking memory address: 00100000
System RAM test using Default POST RAM Test succeeded.
Press escape within 4 seconds to enter boot interface.
Bootimg "Fabric Operating System" image.
Linux/PPC load:
BootROM command line: quiet
Uncompressing Linux...done.
Now booting the kernel
Attempting to find a root file system on hda2...
modprobe: modprobe: Can't open dependencies file /lib/modules/2.4.19/modules.dep (No such file or directory)
INIT: version 2.78 booting
INIT: Entering runlevel: 3
eth0: Link status change: Link Up. 100 Mbps Full duplex Auto (autonegotiation complete).
INITCP: CPLD Vers: 0x95 Image ID: 0x19
uptime: 2008; sysc_qid: 0
Fabric OS (Paulsa45)
Paulsa45 console login: 2005/03/31-20:12:42, [TRCE-5000], 0,, INFO, ?, trace:, trace_buffer.c, line: 1170
2005/03/31-20:12:42, [LOG-5000], 0,, INFO, SW4100_P45, Previous message repeat 1 time(s), trace_ulib.c,
line: 540
2005/03/31-20:12:43, [HAM-1004], 219,, INFO, SW4100_P45, Processor rebooted - Unknown
SNMP Research SNMP Agent Resident Module Version 15.3.1.4
Copyright 1989, 1990, 1991, 1992, 1993, 1994, 1995, 1996, 1997, 1998, 1999, 2000, 2001 SNMP Research, Inc.
sysctrlid: all services Standby
```

```

FSSK 2: chassis0(0): state not synchronized
FSSK 2: Services starting a COLD recovery
2005/03/31-20:12:48, [FSS-5002], 0,, INFO, SW4100_P45, chassis0(0): state not synchronized, svc.c, line: 318
2005/03/31-20:12:48, [FSS-5002], 0,, INFO, SW4100_P45, Services starting a COLD recovery, mdev.c, line: 638
2005/03/31-20:12:49, [MFIC-1002], 220,, INFO, Paulsa45, Chassis FRU header not programmed for switch NID,
using defaults (applies only to FICON environments).
sysctrlld: all services Active
2005/03/31-20:12:50, [DGD-5001], 0,, INFO, SW4100_P45, Slot 0 has started POST., main.c, line: 1189
POST1: Started running Thu Mar 31 20:12:51 GMT 2005
POST1: Test #1 - Running turboramtest
POST1: Test #2 - Running portregtest
POST1: Script PASSED with exit status of 0 Thu Mar 31 20:12:54 GMT 2005 took (0:0:3)
POST2: Started running Thu Mar 31 20:12:55 GMT 2005
POST2: Test #1 - Running portloopbacktest (SERDES)
POST2: Test #2 - Running minicycle (SERDES)
POST2: Running diagshow
POST2: Script PASSED with exit status of 0 Thu Mar 31 20:13:12 GMT 2005 took (0:0:17)
2005/03/31-20:13:13, [BL-1000], 221,, INFO, Paulsa45, Initializing Ports... Enabling switch...
2005/03/31-20:13:13, [BL-1001], 222,, INFO, Paulsa45, Port Initialization Completed
2005/03/31-20:13:13, [EM-5012], 0,, INFO, SW4100_P45, EM: sent dumpready to ME., em.c, line: 2152
2005/03/31-20:13:13, [DGD-5002], 0,, INFO, SW4100_P45, Slot 0 has passed the POST tests., main.c, line: 936

```

If you choose to bypass POST or after POST completes, various system services are started and the boot process displays additional console status and progress messages.

NOTE

POST must be enabled or disabled in advance (just before power-cycling or rebooting the chassis) Do not enable or disable the POST in the middle of the chassis power ON/reboot sequence, which results in unpredictable behavior. POST must not be disabled or enabled in the middle chassis or slot power ON/reboot sequence.

Disabling POST

A reboot is not required for **diagDisablePost** to take effect.

NOTE

Disabling POST is not recommended and should only be done on the advice of your customer support technician.

1. Connect to the switch and log in with a user account that has admin privileges with the chassis-role permission.
2. Enter the **diagDisablePost** command.

This disables POST1 and POST2.

Enabling POST

A reboot is required for **diagEnablePost** to take effect.

1. Connect to the switch and log in with a user account that has admin privileges with the chassis-role permission.
2. Enter the **diagEnablePost** command to enable POST and reboot the switch for POST tests to run.

```

switch:admin> diagenablepost
Config update Succeeded
Diagnostic POST is now enabled.

```

Switch status

Use the **mapsdb --show** command to display the overall status of the switch, including its power supplies, fans, and temperature. If the status of any one of these components is either marginal or down, the overall status of the switch is also displayed as marginal or down. If all components have a healthy status, the switch displays a healthy status.

To modify the rules used to classify the health of each component, use the **switchStatusPolicySet** command. To view the rules, use the **switchStatusPolicyShow** command.

NOTE

In Fabric OS v7.4.0 and later, the **switchStatusShow** command does work when Flow Vision is enabled. Refer to *Fabric OS Command Reference* and the *Flow Vision Administrator's Guide* for more information.

Viewing the overall status of the switch

1. Connect to the switch and log in as admin.

2. Enter the **mapsdb --show** command.

```
switch:admin> mapsdb --show
1 Dashboard Information:
=====
DB start time: Thu Feb 4 19:17:13 2016
Active policy: dflt_aggressive_policy
Configured Notifications: RASLOG,EMAIL,FENCE
Fenced Ports : 5/60,5/62
Decommissioned Ports : None
Fenced circuits : None
Quarantined Ports : None
Top PIDs <pid(it-flows)>: 0x69b0c0(8) 0x697b00(4)
2 Switch Health Report:
=====
Current Switch Policy Status: CRITICAL
Contributing Factors:
=====
*BAD_PWR (CRITICAL).
*BAD_FAN (MARGINAL).
3.1 Summary Report:
=====
Category |Today |Last 7 days |
-----
Port Health |Out of operating range |No Errors |
BE Port Health |No Errors |No Errors |
GE Port Health |In operating range |No Errors |
Fru Health |Out of operating range |In operating range |
Security Violations |No Errors |No Errors |
Fabric State Changes |Out of operating range |No Errors |
Switch Resource |In operating range |In operating range |
Traffic Performance |In operating range |In operating range |
FCIP Health |No Errors |No Errors |
Fabric Performance Impact|In operating range |In operating range |
3.2 Rules Affecting Health:
=====
Category |Repeat|Rule Name |Execution Time |Object |Triggered |
(Rule Count) |Count | | |Value(Units)|
-----
Fru Health(2)|2 |defALL_PSPS_ |02/04/16 21:32:16 |Power Supply 3 |FAULTY |
| |STATE_FAULTY | | |
| | |Power Supply 4 |FAULTY |
```

For more information on how the overall switch status is determined, refer to the **switchStatusPolicySet** command in the *Fabric OS Command Reference*.

If Flow Vision is enabled the following message appears:

```
ras225:FID128:admin> switchstatusshow
MAPS is enabled, Fabric Watch is disabled. Please use MAPS for monitoring and execute mapsHelp
command for available MAPS commands.
```

Displaying switch information

To display the switch information, perform the following task.

1. Connect to the switch and log in as admin.

2. Enter the **switchShow** command.

[Table 38](#) lists the switch summary information.

TABLE 38 Switch summary information

Variable	Definition
switchType	Switch model and revision numbers
switchName	Switch name
switchState	Switch state: Online, Offline, Testing, or Faulty
switchMode	Switch operation mode: Native, Interop, or Access Gateway
switchRole	Principal, Subordinate, or Disabled
switchDomain	ID: 0-31 or 1-23
switchId	Switch embedded port D_ID
switchWwn	Switch World Wide Name (WWN)
switchBeacon	Switch beaconing state: On or Off
zoning	When Access Gateway mode disabled, the name of the active zone displays in parentheses.
FC Router	FC Router's state: On or Off
FC Router BB Fabric ID	The backbone fabric ID for FC routing

[Table 39](#) lists the following additional properties displayed in the switch summary for Virtual Fabrics-enabled switches.

TABLE 39 VF output values

Variable	Definition
LS Attributes	Displays logical switch attributes, including the fabric ID (FID) associated with the logical switch and the switch role (default switch or base switch).
Allow XISL Use	Allows the switch to use extended interswitch links (XISLs) in the base fabric to carry traffic to this logical switch. Values are ON or OFF.

[Table 40](#) lists the **switchShow** command output information for ports on the specified switch:

TABLE 40 switchShow command output

Variable	Definition
Index	Index follows Area up to 255. Then it continues to the maximum port of the platform. Index identifies the port number relative to the switch. Index column is only displayed on enterprise-class platforms.
Slot	Slot number 1-4 and 7-10.
Port	Port number 0-15, 0-31, or 0-47.
Address	The 24-bit Address Identifier. Address column is only displayed on enterprise-class platforms.
Media	SFP types used.
Speed	The speed of the Port (1G, 2G, 4G, 8G, 10G, N1, N2, N4, N8, AN, UN). The speed can be fixed, negotiated, or auto-negotiated.
State	The port status.
Proto	Protocol support by GbE port.

The details displayed for each switch differ on different switch models. For more information refer to the **switchShow** command in the *Fabric OS Command Reference*.

Displaying the uptime for a switch

1. Connect to the switch and log in as admin.
2. Enter the **uptime** command.

```
ecp:admin> uptime
10:50:19 up 11 days, 6:28, 1 user, load average: 0.49, 0.53, 0.54
```

The **uptime** command displays the length of time the system has been in operation, the total cumulative amount of uptime since the system was first powered on, and the load average over the past one minute. The reason for the last switch reboot is also recorded in the system message log.

Using the spinFab and portTest commands

The **spinFab** command is an online diagnostics command to verify the ISLs between switches at the maximum speed. The routing functionality in the hardware must be set up so that the test frames received by the E_Port are retransmitted on the same E_Port. Several frames are then sent to the port attached to each active E_Port specified. These frames are special frames that never occur during normal traffic, and the default action for such frames is to route them back to the sender. These frames are circulated between switches until the test stops them.



CAUTION

During the spinFab testing, the switch remains in normal operation. However, some performance degradation occurs due to the ISL links being saturated with the test frames. This test should be run with caution on a live fabric.

The following table lists the supported ports for the specified version of Fabric OS when using the **spinFab** command.

TABLE 41 Port type support

Port type	Supported in v6.3.0	Supported in v6.4.0	Supported in v7.2.0 and later
Loopback	Yes	Yes	Yes
D_Ports	No	No	Yes
D_Ports to AG switch ports	No	No	Yes
E_Ports	Yes	Yes	Yes
Trunk Master ports	Yes	Yes	Yes
Ports beyond index 255	Yes	Yes	Yes
Ports with swapped areas	Yes	Yes	Yes
Shared-area ports	Yes	Yes	Yes
Ports in logical switches	Yes	Yes	Yes
Ports in base switches	Yes	Yes	Yes
Trunk slave ports	No	Yes	Yes
Long distance ports	No	Yes	Yes
F_Ports connected to Brocade HBAs	No	Yes	Yes
ICL ports	No	No	Yes
F_Ports connected to Access Gateway	No	No	Yes
EX_Ports	No	No	No
Ports in an Access Gateway switch	No	No	No

NOTE

The **portTest** and **spinFab** commands are not supported on AE and AF ports for Brocade Analytics Monitoring Platform connections.

The **SpinFab** and **portTest** commands on Ethernet ports are not supported for FC32-64 blade and Brocade G630 switch.

The **SpinFab** and **portTest** commands on VE ports are not supported for Brocade 7840 Extension Switch, Brocade 7810 Extension Switch, Brocade SX6 Extension Blade and FX8-24 Extension Blade.

Debugging spinFab errors

Link errors and transmit or receive errors are seen when the **spinFab** test fails.

Link errors

Once the frame is sent out of the port, the **spinFab** command monitors the link errors in the ASIC. If any of the error counters are non-zero, **spinFab** reports an error and the test fails on the port.

```
ERROR: DIAG ERRSTAT spinfab, pass 6,
Pt0/17(7) Ch0/7 CRC_err Error Counter is 109738997 sb 0,

ERROR: DIAG ERRSTAT spinfab, pass 1,
Pt0/3(33) Ch0/33 Enc_in Error Counter is 32 sb 0,
ERROR: DIAG ERRSTAT spinfab, pass 1,
Pt0/3(33) Ch0/33 Enc_out Error Counter is 187725412 sb 0,
ERROR: DIAG ERRSTAT spinfab, pass 1,
Pt0/3(33) Ch0/33 TruncFrm Error Counter is 32 sb 0,
ERROR: DIAG ERRSTAT spinfab, pass 1,
Pt0/3(33) Ch0/33 FrmTooLong Error Counter is 32 sb 0,
ERROR: DIAG ERRSTAT spinfab, pass 1,
Pt0/3(33) Ch0/33 BadOrdSet Error Counter is 32 sb 0,
ERROR: DIAG ERRSTAT spinfab, pass 1,
Pt0/3(33) Ch0/33 BadEOF Error Counter is 1 sb 0,
ERROR: DIAG ERRSTAT spinfab, pass 1,
Pt0/3(33) Ch0/33 DiscC3 Error Counter is 32 sb 0,
```

If you receive any of the link errors, follow the suggested debugging procedures:

- The **spinFab** command does not clear any existing error counters before running the test. You should first clear all error counters and rerun the **spinFab** command.
- Verify that the link comes up by enabling and disabling the local or remote ports.
- Verify that the source of the error is either the local port or the remote port. This can be done by monitoring the port statistics on both ends simultaneously. Refer to [Displaying the port statistics](#) on page 86 for more information on how to display the statistics for a port.
- Verify that the cables and SFPs are inserted properly. Remove and insert them again on both ends.
- Verify that the failing local port is working when connected to another remote port. Similarly, check whether the failing remote port is working when connected to another local port.
- Once the fault is isolated on either the local port or the remote port, replace the cable and SFPs connected to the local port and the remote ports.
- In case of loopback ports, change the loopback plug. Refer to [Marginal links](#) on page 37 for more information on changing the loopback plug.
- Further isolation can be done by running the **portLoopbackTest** command (Offline test) on the failing port to check whether the blade internal ports are having some problems.
 - The **-lb_mode 1** operand verifies that the SFP is working within normal operating parameters. The use of this operand requires that loopback cables are connected.

- The **-lb_mode 2** verifies that the ASIC port is working within normal operating parameters. The use of this operand does not require any loopback cables.

Tx/Rx errors

The following errors are seen when the port fails to transmit or receive the frames.

```
ERROR: DIAG PORTSTOPPED spinfab:spinfab, 0 nMegs,
Pt7/2(2) Ch0/2 No Longer Transmitting, FTX Counter Stuck at 116295726,
ERROR: DIAG TIMEOUT spinfab:spinfab, pass 2,
Pt0/17(7) Ch0/7 Receive Error/Timeout
```

If you receive any of the Tx/Rx errors, follow the suggest debugging procedures:

- Check whether the same port is reporting link errors as discussed in [Link errors](#) on page 84. If yes, follow the same set of debugging procedures as discussed in [Link errors](#) on page 84.
- Check whether the local port or the remote port is beyond port 255. If yes, try connecting to the lower number of ports. This behavior is found in Fabric OS v6.2.0 and earlier versions only.
- Check whether the local port or the remote port port is part of a shared-area region. If yes, try connecting to the non-shared area ports. This behavior is found in Fabric OS v6.2.0 and earlier versions only.
- Check whether the local port or the remote port is having its area swapped. If yes, try connecting to the normal area ports. This behavior is found in Fabric OS v6.2.0 and earlier versions only.

Clearing the error counters

This procedure clears the port hardware statistics, including ALPA-based CRC monitor, End-to-End monitor, and filter-based performance monitor statistics.

1. Connect to the switch and log in as admin.
2. Enter the **portStatsClear** command.

Enabling a port

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the appropriate command based on the current state of the port and whether it is necessary to specify a slot number:
 - To enable a port that is disabled, enter the **portEnable** command.
 - To enable a port that is persistently disabled, enter the **portCfgPersistentEnable** command.

If you change port configurations during a switch failover, the ports may become disabled. To bring the ports online, re-issue the **portEnable** command after the failover is complete.

Disabling a port

1. Connect to the switch and log in using an account with admin permissions.
2. Enter the appropriate command based on the current state of the port and on whether it is necessary to specify a slot number:
 - To disable a port that is enabled, enter the **portDisable** command.
 - To disable a port that is persistently enabled, enter the **portCfgPersistentDisable** command.

Port information

Use the following instructions to view information about ports and to help diagnose if your switch is experiencing port problems.

NOTE

For a detailed discussion of the diagnostic port (D_Port) feature, refer to [Brocade ClearLink Diagnostic Port](#) on page 97.

Viewing the status of a port

1. Connect to the switch and log in as admin.
2. Enter the **portShow** command, specifying the number that corresponds to the port you are troubleshooting. In this example, the status of port 10 is shown:

```
switch:admin> portshow 10
portName:
portHealth: HEALTHY
Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x20b03      PRESENT ACTIVE F_PORT G_PORT U_PORT LOGICAL_ONLINE LOGIN NOELP ACCEPT FLOGI
portType: 18.0
POD Port: Port is licensed
portState: 1      Online
portPhys: 6      In_Sync
portScn: 32      F_Port
port generation number: 14
portId: 020a00
portIfId: 4302000b
portWwn: 20:0a:00:05:1e:41:4a:a5
portWwn of device(s) connected:
    21:00:00:e0:8b:05:e0:b1
Distance: normal
portSpeed: N2Gbps
LE domain: 0
FC Fastwrite: OFF
Interrupts:          0          Link_failure: 0          Frjt:          0
Unknown:             0          Loss_of_sync: 3          Fbsy:          0
Lli:                 18          Loss_of_sig: 6
Proc_rqrd:           161         Protocol_err: 0
Timed_out:           0          Invalid_word: 563851
Rx_flushed:          0          Invalid_crc: 0
Tx_unavail:          0          Delim_err: 0
Free_buffer:         0          Address_err: 0
Overrun:             0          Lr_in:         3
Suspended:           0          Lr_out:        0
Parity_err:          0          Ols_in:        0
2_parity_err:        0          Ols_out:       3
CMI_bus_err:         0
Port part of other ADs: No
```

Refer to the *Fabric OS Command Reference* for additional **portShow** command information, such as the syntax for slot or port numbering, displaying IP interfaces on a GbE port, or displaying FCIP tunnel connection or configuration information.

Displaying the port statistics

1. Connect to the switch and log in as admin.

2. Enter the **portStatsShow** command.

Port statistics include information such as the number of frames received, number of frames sent, number of encoding errors received, and number of class 2 and class 3 frames received.

Refer to the *Fabric OS Command Reference* for additional **portStatsShow** command information, such as the syntax for slot or port numbering.

```
switch:admin> portstatsshow 68
stat_wtx          113535      4-byte words transmitted
stat_wrx          22813      4-byte words received
stat_ftx          9259       Frames transmitted
stat_frx          821        Frames received
stat_c2_frx       0          Class 2 frames received
stat_c3_frx       821        Class 3 frames received
stat_lc_rx        0          Link control frames received
stat_mc_rx        0          Multicast frames received
stat_mc_to        0          Multicast timeouts
stat_mc_tx        0          Multicast frames transmitted
tim_rdy_pri       0          Time R_RDY high priority
tim_txcrd_z       0          Time TX Credit Zero (2.5Us ticks)
time_txcrd_z_vc   0- 3: 0    0          0
time_txcrd_z_vc   4- 7: 0    0          0
time_txcrd_z_vc   8-11: 0    0          0
time_txcrd_z_vc  12-15: 0    0          0
er_enc_in         0          Encoding errors inside of frames
er_crc            0          Frames with CRC errors
er_trunc          0          Frames shorter than minimum
er_toolong        0          Frames longer than maximum
er_bad_eof        0          Frames with bad end-of-frame
er_enc_out        0          Encoding error outside of frames
er_bad_os         0          Invalid ordered set
er_c3_timeout     0          Class 3 frames discarded due to timeout
er_c3_dest_unreach 0          Class 3 frames discarded due to destination unreachable
er_other_discard  0          Other discards
er_type1_miss     0          frames with FTB type 1 miss
er_type2_miss     0          frames with FTB type 2 miss
er_type6_miss     0          frames with FTB type 6 miss
er_zone_miss      0          frames with hard zoning miss
er_lun_zone_miss  0          frames with LUN zoning miss
er_crc_good_eof   0          Crc error with good eof
er_inv_arb        0          Invalid ARB
open              810        loop_open
transfer          0          loop_transfer
opened            409856     FL_Port opened
starve_stop       0          tenancies stopped due to starvation
fl_tenancy        1715       number of times FL has the tenancy
nl_tenancy        331135     number of times NL has the tenancy
zero_tenancy      4          zero tenancy
```

Displaying a summary of port errors for a switch

1. Connect to the switch and log in as admin.

2. Enter the **porterrshow** command.

```

switch:admin> porterrshow
      frames  enc  crc  too  too  bad  enc  disc  link  loss  loss  frjt  fbsy
      tx   rx   in  err  g_eof shrt long eof  out   c3 fail sync sig
=====
0:  665k 7.0k  0  0  0  0  0  0  6  0  0  1  2  0  0
1:  0  0  0  0  0  0  0  0  0  0  0  0  2  0  0
2:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
3:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
4:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
5:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
6:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
7:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
8:  78  60  0  0  0  0  0  0  7  0  0  3  6  0  0
9:  12  4  0  0  0  0  0  0  3  0  0  1  2  0  0
10:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
11:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
12:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
13:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
14:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
15:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
16:  665k 7.4k  0  0  0  0  0  0  6  0  0  1  2  0  0
17:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
18:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
19:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
20:  6.3k 6.6k  0  0  0  0  0  0  7  0  0  1  2  0  0
21:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
22:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
23:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
24:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
25:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
26:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
27:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
28:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
29:  0  0  0  0  0  0  0  0  0  0  0  0  1  0  0
30:  664k 6.7k  0  0  0  0  0  0  6  0  0  1  2  0  0
31:  12  4  0  0  0  0  0  0  3  0  0  1  2  0  0
(output truncated)

```

The **porterrshow** command output provides one output line per port. Refer to [Table 42](#) for a description of the error types.

TABLE 42 Error summary description

Error type	Description
frames tx	Frames transmitted
frames rx	Frames received
enc in	Encoding errors inside frames
crc err	Frames with CRC errors
crc g_eof	CRC errors that occur on frames with good end-of-frame delimiters
too shrt	Frames shorter than minimum
too long	Frames longer than maximum
bad eof	Frames with bad end-of-frame delimiters
enc out	Encoding error outside of frames
disc c3	Class 3 frames discarded
link fail	Link failures (LF1 or LF2 states)
loss sync	Loss of synchronization
loss sig	Loss of signal
frjt	Frames rejected with F_RJT
fbsy	Frames busied with F_BSY

Equipment status

You can display status for fans, power supplies, and temperature.

NOTE

The number of fans, power supplies, and temperature sensors depends on the switch type. For detailed specifications on these components, refer to the switch hardware reference manual. The specific output from the status commands varies depending on the switch type.

Checking the temperature, fan, and power supply

1. Log in to the switch as admin.
2. Enter the **sensorShow** command. Refer to the *Fabric OS Command Reference* for details regarding the sensor numbers.
3. Check the temperature output.

Look for indications of high or low temperatures.

4. Check the fan speed output.

If any of the fan speeds display abnormal RPMs, replace the fan FRU.

5. Check the power supply status.

If any power supplies show a status other than OK, consider replacing the power supply as soon as possible.

Checking the status of the fans

1. Connect to the switch and log in as admin.
2. Enter the **fanShow** command.

```
switch:admin> fanshow
Fan 1 is Absent
Fan 2 is Ok, speed is 6553 RPM
Fan 3 is Ok, speed is 6367 RPM
```

The possible status values are:

- OK—Fan is functioning correctly.
- Absent—Fan is not present.
- Below minimum—Fan is present but rotating too slowly or stopped.
- Above minimum—Fan is rotating too quickly.
- Unknown—Unknown fan unit installed.
- FAULTY—Fan has exceeded hardware tolerance or is not be seated properly.

The output from this command varies depending on switch type and number of fans present. Refer to the appropriate hardware reference manual for details regarding the fan status. You may first consider re-seating the fan (unplug it and plug it back in).

Checking the status of a power supply

1. Connect to the switch and log in as admin.

2. Enter the **psShow** command.

```
switch:admin> psshow
Power Supply #1 is OK
V10645,TQ2Z6452916      ,60-0300031-02, B, QCS ,DCJ3001-02P      , A,TQ2Z64529
Power Supply #2 is faulty
V10704,      TQ2J7040124      ,60-0300031-02, B,CHRKE,SP640-Y01A      ,C ,TQ2J7040
```

The possible status values are:

- OK—Power supply functioning correctly.
- Absent—Power supply not present.
- Unknown—Unknown power supply unit installed.
- Predicting failure—Power supply is present but predicting failure.
- FAULTY—Power supply is present but faulty (no power cable, power switch turned off, fuse blown, or other internal error).

If any of the power supplies show a status other than OK, consider replacing the power supply as soon as possible. For certain switch models, the OEM serial ID data displays after each power supply status line.

Checking temperature status

1. Connect to the switch and log in as admin.
2. Enter the **tempShow** command.

```
switch:admin> tempshow
Sensor  State      Centigrade      Fahrenheit
ID
=====
1      Ok              28              82
2      Ok              16              60
3      Ok              18              64
```

Information is displayed for each temperature sensor in the switch.

The possible temperature status values are:

- OK: Temperature is within acceptable range.
- FAIL: Temperature is outside of acceptable range. Damage might occur.

Refer to the hardware reference manual for your switch to determine the normal temperature range.

System message log

The system message log (RASlog) feature enables messages to be saved across power cycles and reboots.

The Brocade DCX 8510 and X6 family enterprise-class platforms maintain independent and separate RASlogs for each of the two CP blades. Because all RASlog messages are routed to the Active CP, the message CPU ID is added as part of the RASlog message attribute. RASlog message attribute *SLOT* is defined to identify the CPU that generated the message.

For example, in the following message, the identifier *SLOT 6* means the message was generated from the slot 6 blade main CPU:

```
2001/01/07-04:03:00, [SEC-1203], 2,SLOT 6 | FFDC | CHASSIS, INFO, C08_1, Login information: Login
successful via TELNET/SSH/RSH. IP Addr: 192.168.38.2050
```

In the following message, the identifier *SLOT 6/1* means the message was generated from the slot 6 blade co-CPU.

```
2001/01/07-04:03:00, [SEC-1203], 2, SLOT 6/1 , | FFDC | CHASSIS, INFO, C08_1, Login information: Login
successful via TELNET/SSH/RSH. IP Addr: 192.168.38.2050
```

Because RASlog supports Virtual Fabrics and logical switches, the *FID* (Fabric ID) displays on every RASlog message to identify the source of the logical switch that generates the message.

The FID can be a number from 0 to 128, and the identifier *CHASSIS* depends on the instance that generates the message and that it was generated by a chassis instance. The identifier *FID 128* means the message was generated by the default switch instance.

```
2008/08/01-00:19:44, [LOG-1003], 1, SLOT 6 | CHASSIS, INFO, Silkworm12000, The log has been cleared.
2008/09/08-06:52:50, [FW-1424], 187, SLOT 6 | FID 128, WARNING, Switch10, Switch status changed from
HEALTHY to DOWN.
```

For details on error messages, refer to the *Brocade Fabric OS Message Reference*.

Displaying the system message log with no page breaks

1. Connect to the switch and log in as admin.
2. Enter the **errDump** command.

Displaying the system message log one message at a time

1. Connect to the switch and log in as admin.
2. Enter the **errShow** command.

Clearing the system message log

1. Connect to the switch and log in as admin.
2. Enter the **errClear** command.
3. Repeat step 2 on the standby CP for a complete erasure of the message log.

All switch and chassis events are removed from both CPs.

Port log

Fabric OS maintains an internal log of all port activity. The port log stores entries for each port as a circular buffer. For all other switches, the number of lines range from 8192 to 16384. These ranges are for all ports on the switch, not just for one port. When the log is full, the newest log entries overwrite the oldest log entries. The port log is not persistent and is lost over power-cycles and reboots. If the port log is disabled, an error message displays.

NOTE

Port log functionality is separate from the system message log. The port log is typically used to troubleshoot device connections.

Viewing the port log

1. Connect to the switch and log in as admin.

2. Enter the **portlogshow** command:

```

device:admin> portlogshow
time      task      event    port cmd  args
-----
Fri Feb 22 16:48:45 2008
16:48:45.208 SPEE      sn       67    NM  00000009,00000000,00000000
16:48:46.783 PORT      Rx       64    40  02ffffff,00ffffff,02e2ffff,14000000
16:48:46.783 PORT      Tx       64    0   cffffff,00ffffff,02e201bf,00000001
16:48:46.783 FCPH      read     64    40  02ffffff,00ffffff,be000000,00000000,02e201bf
16:48:46.783 FCPH      seq      64    28  22380000,02e201bf,00000c1e,0000001c,00000000
16:48:46.828 SPEE      sn       67    NM  00000009,00000000,00000000
16:48:46.853 PORT      Rx       76    40  02ffffff,00ffffff,02e3ffff,14000000
16:48:46.853 PORT      Tx       76    0   cffffff,00ffffff,02e301c1,00000001
16:48:46.853 FCPH      read     76    40  02ffffff,00ffffff,bf000000,00000000,02e301c1
16:48:46.853 FCPH      seq      76    28  22380000,02e301c1,00000c1e,0000001c,00000000
16:48:47.263 PORT      Rx       79    40  02ffffff,00ffffff,02e4ffff,14000000
16:48:47.263 PORT      Tx       79    0   cffffff,00ffffff,02e401c2,00000001
(output truncated)

```

Use the commands summarized in [Table 43](#) to view and manage port logs. Refer to the *Brocade Fabric OS Message Reference* for additional information about these commands.

TABLE 43 Commands for port log management

Command	Description
portlogclear	Clear port logs for all or particular ports.
portlogdisable	Disable port logs for all or particular ports.
portlogdump	Display port logs for all or particular ports, without page breaks.
portlogenable	Enable port logs for all or particular ports.
portlogshow	Display port logs for all or particular ports, with page breaks.

The **portLogDump** command output (trace) is a powerful tool that is used to troubleshoot fabric issues. The **portLogDump** output provides detailed information about the actions and communications within a fabric. By understanding the processes that are taking place in the fabric, issues can be identified and located.

The **portLogDump** command displays the port log, showing a portion of the Fibre Channel payload and header (FC-PH). The header contains control and addressing information associated with the frame. The payload contains the information being transported by the frame and is determined by the higher-level service or FC_4 upper level protocol. There are many different payload formats based on the protocol.

Because a **portLogDump** output is long, a truncated example is presented:

```

switch:admin> portlogdump
time      task      event    port cmd  args
-----
Fri Feb 22 20:29:12 2008
20:29:12.638 FCPH      write    3     40  00ffffff,00ffffff,00000000,00000000,00000000
20:29:12.638 FCPH      seq      3     28  00300000,00000000,000005f4,00020182,00000000
20:29:12.638 PORT      Tx       3     40  02ffffff,00ffffff,09a5ffff,14000000
20:29:12.638 FCPH      write    9     40  00ffffff,00ffffff,00000000,00000000,00000000
20:29:12.638 FCPH      seq      9     28  00300000,00000000,000005f4,00020182,00000000
20:29:12.639 PORT      Tx       9     40  02ffffff,00ffffff,09a6ffff,14000000
20:29:12.639 PORT      Rx       3     0   cffffff,00ffffff,09a50304,00000001
20:29:12.640 PORT      Rx       9     0   cffffff,00ffffff,09a60305,00000001
20:29:20.804 PORT      Rx       9     40  02ffffff,00ffffff,0306ffff,14000000
20:29:20.805 PORT      Tx       9     0   cffffff,00ffffff,030609a7,00000001
20:29:20.805 FCPH      read     9     40  02ffffff,00ffffff,d1000000,00000000,030609a7
20:29:20.805 FCPH      seq      9     28  22380000,030609a7,00000608,0000001c,00000000
20:29:20.805 PORT      Rx       3     40  02ffffff,00ffffff,02eeffff,14000000
20:29:20.806 PORT      Tx       3     0   cffffff,00ffffff,02ee09a8,00000001
20:29:20.806 FCPH      read     3     40  02ffffff,00ffffff,d2000000,00000000,02ee09a8
20:29:20.806 FCPH      seq      3     28  22380000,02ee09a8,00000608,0000001c,00000000

```

```

20:29:32.638 FCPH      write    3    40  00ffffffd,00ffffffd,00000000,00000000,00000000
20:29:32.638 FCPH      seq      3    28  00300000,00000000,000005f4,00020182,00000000
20:29:32.638 PORT      Tx        3    40  02ffffffd,00ffffffd,09a9ffff,14000000
20:29:32.638 FCPH      write    9    40  00ffffffd,00ffffffd,00000000,00000000, 00000000
20:29:32.638 FCPH      seq      9    28  00300000,00000000,000005f4,00020182,00000000
20:29:32.639 PORT      Tx        9    40  02ffffffd,00ffffffd,09aaffff,14000000
<output truncated>

```

Syslogd configuration

The system logging daemon (syslogd) is an IP-based service for logging system messages made available by default on UNIX and Linux operating systems. It is available as a third-party application for Windows operating systems.

Fabric OS can be configured to use a UNIX-style syslogd process to forward system events and error messages to log files on a remote host system. The host system can be running UNIX, Linux, or any other operating system that supports the standard syslogd functionality.

Fabric OS supports UNIX local7 facilities (the default facility level is 7). Configuring for syslogd involves configuring the host, enabling syslogd on the switch, and, optionally, setting the facility level.

Configuring the host

Fabric OS supports a subset of UNIX-style message severities that default to the UNIX local7 facility. To configure the host, edit the `/etc/syslog.conf` file to map Fabric OS message severities to UNIX severities, as shown in [Table 44](#).

TABLE 44 Fabric OS to UNIX message severities

Fabric OS message severity	UNIX message severity
Critical (1)	Emergency (0)
Error (2)	Error (3)
Warning (3)	Warning (4)
Info (4)	Info (6)

In this example, Fabric OS messages map to local7 facility level 7 in the `/etc/syslog.conf` file:

```

local7.emerg      /var/adm/swcritical
local7.alert      /var/adm/alert7
local7.crit       /var/adm/crit7
local7.err        /var/adm/swerror
local7.warning    /var/adm/swwarning
local7.notice     /var/adm/notice7
local7.info       /var/adm/swinfo
local7.debug      /var/adm/debug7

```

If you prefer to map Fabric OS severities to a different UNIX local7 facility level, refer to [Setting the facility level](#) on page 94.

Configuring the switch

Configuring the switch involves specifying syslogd hosts and, optionally, setting the facility level. You can also remove a host from the list of syslogd hosts.

Specifying syslogd hosts

1. Connect to the switch and log in as admin.
2. Enter the **syslogdIpAdd** command and specify an IP address.
3. Verify that the IP address was entered correctly using the **syslogdIpShow** command.

The **syslogdIpAdd** command accepts IPv4 and IPv6 addresses. You can specify up to six host IP addresses for storing syslog messages, as shown in this example:

```
switch:admin> syslogdipadd 1080::8:800:200C:417A
switch:admin> syslogdipadd 1081::8:800:200C:417A
switch:admin> syslogdipadd 1082::8:800:200C:417A
switch:admin> syslogdipadd 10.1.2.4
switch:admin> syslogdipadd 10.1.2.5
switch:admin> syslogdipadd 10.1.2.6
switch:admin> syslogdipshow
syslog.IP.address.1080::8:800:200C:417A
syslog.IP.address.1081::8:800:200C:417A
syslog.IP.address.1082::8:800:200C:417A
syslog.IP.address.4 10.1.2.4
syslog.IP.address.5 10.1.2.5
syslog.IP.address.6 10.1.2.6
```

Setting the facility level

1. Connect to the switch and log in as admin.
2. Enter the **syslogdfacility -l n** command:

The *n* variable is a number from 0 through 7, indicating a UNIX local7 facility. The default is 7.

You need to set the facility level only if you specified a facility other than local7 in the host `/etc/syslog.conf` file.

Removing a syslogd host from the list

1. Connect to the switch and log in as admin.
2. Enter the **syslogdIpRemove** command followed by the IP address of the host that you want to remove.

```
switch:admin> syslogdipremove 10.1.2.1
```

3. Enter the **syslogdIpShow** command to verify if the IP address was deleted.

Automatic trace dump transfers

You can set up a switch so that diagnostic information is transferred automatically to a remote server. If a problem occurs, you can then provide your customer support representative with the most detailed information possible. To ensure the best service, you should set up for automatic transfer as part of standard switch configuration, before a problem occurs.

We always check `cfgload` attribute also while uploading the trace dump and use it as a protocol when it is configured as SCP. Otherwise we use FTP by default. Also we don't allow the user to configure `supportftp` protocol as FTP when `cfgload` attribute is configured as secured.

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.
 - In the case of Gen 5 devices, the trace dump files overwrite each other by default; sending them to a server preserves information that would otherwise be lost.
 - In the case of Gen 6 devices, the trace dump files are stored on the switch by its type up to a limit and transferred to an FTP server when the limit exceeds. The files are deleted after they are transferred to an FTP server.
- Setting up a periodic checking of the remote server so that you are alerted if the server becomes unavailable and you can correct the problem.

After the setup is complete, you can run the **supportSave -c** command to save RASlog, TRACE, supportShow, core file, FFDC data and other diagnostic support information to the server without specifying server details.

The following procedures describe the tasks for setting up automatic transfer.

Specifying a remote server

1. Verify that the FTP service is running on the remote server.
2. Connect to the switch and log in as admin.
3. Enter the **supportFtp -s** command and respond to the prompts.

Enabling the automatic transfer of trace dumps

1. Connect to the switch and log in as admin.
2. Enter the **supportFtp -e** command.

Setting up periodic checking of the remote server

1. Connect to the switch and log in as admin.
2. Enter the **supportFtp -t** command.

Example of setting the interval in hours

```
switch:admin> supportftp -t 4
supportftp: ftp check period changed
```

The minimum interval is 1 hour. Specify 0 to disable the checking feature.

Saving comprehensive diagnostic files to the server

1. Connect to the switch and log in as admin.
2. Enter the **supportSave -c** command and respond to the prompts.

```
switch:admin> supportsave -c
This command will collect RASLOG, TRACE, supportShow, core file, FFDC data
and other support information and then transfer them to a FTP/SCP server
or a USB device. This operation can take several minutes.
NOTE: supportSave will transfer existing trace dump file first, then
automatically generate and transfer latest one. There will be two trace dump
files transferred after this command.
OK to proceed? (yes, y, no, n): [no] y
```

Multiple trace dump files support

The following devices with warm memory support multiple trace dump files as described in this section.

- Brocade 7840 Extension Switch
- Brocade G620
- Brocade X6-4
- Brocade X6-8

Each of the following commands and scenarios trigger a separate trace dump from the warm memory to a file.

- **traceDump -n**
- **supportSave**
- When a software VERIFY error is detected.
- When panic dump is triggered.

Auto FTP support

You can automatically upload the trace dump triggered by FFDC and the **traceDump -n** command to a remote FTP server using the auto FTP feature.

- In case of **traceDump -n**, the tar file is automatically uploaded to the remote location. You can use the **supportDecode** command to decode.
- In case of FFDC, two tar files are generated. The *core_file.tar* collects RASlog and *trace.tar* file collects FFDC trace logs.

Trace dump support

You can run the **traceDump** command to display the trace dump details. The following is an example from a Gen 6 device. For Gen 5 devices, refer to *Brocade Fabric OS Command Reference*.

```
switch:admin> tracedump
Dump status for switch:
Type                                Timestamp
-----
CLI                                2015/08/31 17:58
Panicdump                          2015/08/31 17:58
FFDC (EM-1100)                     2015/08/21 01:33
```

You can use the **traceDump -r** command to remove the trace dump files generated by CLI, panic dump, and FFDC messages. Fixed-port switches do not support **-s** option.

NOTE

In the case of Gen 6 platforms, the trace dump files are automatically deleted from the switch when the Auto FTP feature is enabled and you run the **supportSave** command.

Brocade ClearLink Diagnostic Port

• Enhanced support for 32Gbps QSFPs.....	97
• Supported platforms for D_Ports.....	99
• Licensing requirements for D_Ports.....	101
• Understanding D_Ports.....	101
• Topology 1: ISLs.....	106
• Topology 2: ICLs.....	107
• Topology 3: Access Gateways.....	107
• Topology 4: HBA to switch.....	109
• Using a D_Port in static-static mode between switches.....	110
• Using D_Ports in dynamic mode.....	112
• Using D_Port mode between switches and HBAs.....	114
• Using a D_Port in on-demand mode	118
• Using the fabriclog command.....	119
• Calculating buffers for long-distance cables.....	120
• Support for audit logs.....	120
• Using D_Port show commands.....	120

Brocade ClearLink Diagnostic Port (D_Port) mode allows you to convert a Fibre Channel port into a diagnostic port for testing traffic and running electrical loopback and optical loopback tests. The test results can be very useful in diagnosing a variety of port and link problems.

NOTE

This feature uses Brocade ClearLink™ proprietary technology.

Enhanced support for 32Gbps QSFPs

Fabric OS 8.1.0a and later, supports electrical and optical loopback tests on 32G-QSFPs that support electrical wrap and optical wrap.

Not all QSFPs currently available support electrical wrap and optical wrap. Generally, electrical wrap must be set on both ends of the link, and optical wrap must be set on the remote end. ISL-to-ISL connections are supported but ICL -to- ICL connections are not. (This support is available on breakout cables.) Both switches must have QSFPs that support electrical wrap and optical wrap. The following are supported connections:

- 32G QSFP ISL < -- > 32G QSFP ISL
- 32G QSFP ISL < -- > 32G SFP ISL
- 32G QSFP ICL < -- > 32G QSFP ICL

Limitations and Considerations

The following QSFPs do not support this feature:

- JAF1: 32G Finisar version 1
- JAF2: 32G Finisar version 2
- JAA1: 32G Avago version 1
- AD1: 32G Avago with early hardware upgraded to JAA2 firmware
- JAA2: 32G Avago version 2

- The earlier versions of the 32G QSFP, having serials number starting with "ZTA" and next 5-digit number less than "11547" (for example, ZTA11516000000K, ZTA11517000001F) supports D_ports and link traffic but does not support electrical wrap and optical wrap. The Later version ZTA11547 and above supports all the three (electrical wrap, optical wrap, and Link Traffic).

In addition, if one end of the link has older 16G- or 32G-QSFPs or QSFPs that do not support electrical wrap or optical wrap, and the other end supports this feature, electrical loopback test will run on the end which supports the feature. Optical loopback test is not supported.

Till Fabric OS 8.1.0, both ends run Electrical loopback test simultaneously. From Fabric OS 8.2.0, one end (initiator switch) will run the electrical loopback test first. After the initiator completes the electrical loopback test, the responder will run the electrical loopback test. The rest of the D_Port test sequencing remains the same. If one end runs older version, the older switch will run the Electrical loopback test first.

Till Fabric OS 8.1.0, when one switch alone is rebooted, the configured D_Ports will run the Electrical loopback test alone and mark the overall test as FAIL. From Fabric OS 8.2.0, when one switch alone is rebooted, the D-Port test will not run on the configured D_Ports and the test status will be marked as STOPPED. When both ends are booted simultaneously, the D-Port test will run automatically on the configured D_Ports as in Previous releases.

Prior to Fabric OS 8.2.0, when the D-Port test is started on multiple QSFP channels, the test will run simultaneously on the multiple channels: for example, electrical loopback test on channel 1, then electrical loopback test on channel 2, followed by optical loopback test on channel 1, etc. With Fabric OS 8.2.0, the tests will be sequenced within a QSFP: the entire test will run on one channel and then the test will be started on the next channel in the QSFP. For example: electrical loopback test on channel 1, optical loopback test on channel 1, link traffic test on channel 1, followed by electrical loopback test on channel 2, etc. The sequencing is only within the quad.

In Fabric OS 8.2.0, support is provided for D-port test if the channels within the QSFP are configured in different LS. Channels within the QSFP still needs to be configured as static D-Port to run the D-port test.

For D-port test on SmartOptics SFP's, User must follow the cable length requirements as per the SFP specification of the modules.

ATTENTION

The following restrictions apply to both 32G/16G breakout QSFP and 32G fixed-speed QSFPs. These restrictions are not applicable if these QSFPs are in ICL ports or they do not support electrical wrap/optical wrap.

- Fabric OS enforces the same static D_Port configuration for all channels in the QSFP ports, for both 32G/16G breakout QSFPs and 32G fixed-speed QSFPs.
- Fabric OS blocks dynamic and skips on-demand D_Port tests for 32G/16G breakout QSFPs and 32G fixed-speed QSFPs.

A warning message appears when the static D_Port configuration for QSFPs is changed.

With these restrictions, an Emulex HBA cannot run on-demand D_Port tests. For an Emulex HBA, you must configure a static D_Port on the switch side and initiate D_Port tests from the Emulex HBA side.

Example output

Prior to Fabric OS 8.1.0, electrical and optical D_Port tests were displayed as "SKIPPED" as they were unsupported, as in the following example output.

For SmartOptics SFP on ISL ports:

```
switch> portdporttest --show 55
D-Port Information:
=====
Port:                55
Remote WWNN:         10:00:00:27:f8:f0:26:38
Remote port index:   55
Mode:                Manual
```

```

No. of test frames:      1 Million
Test frame size:        1024 Bytes
FEC (enabled/option/active):  Yes/No/Yes
CR (enabled/option/active):  Yes/No/No
Start time:             Mon Apr 18 09:57:43 2016
End time:               Mon Apr 18 09:57:50 2016
Status:                 PASSED
=====
Test          Start time    Result          EST(HH:MM:SS)    Comments
=====
Electrical loopback  -----    SKIPPED         -----    No SFP or chip support
Optical loopback     09:57:44    SKIPPED         -----    -----
Link traffic test    09:57:46    PASSED          -----    -----
=====
Roundtrip link latency:      1116 nano-seconds
Approximate cable distance:  unknown
Buffers required:           1 (for 2112 byte frames at 32Gbps speed)
Egress power:               Tx: Not Avail, Rx: -1.6 dBm.
Ingress power:              Rx: -0.5 dBm, Tx: -23.6dBm, Diff: 0.0 dBm (No Loss) 62:admin>

```

With current support, results are displayed as either "PASSED" or "FAILED" as in the following example output.

```

device:admin> portdporttest --show 55
D-Port Information:
=====
Port:                55
Remote WWNN:         10:00:00:27:f8:f0:26:30
Remote port index:   55
Mode:                Automatic
No. of test frames:  1 Million
Test frame size:     1024 Bytes
FEC (enabled/option/active):  Yes/No/Yes
CR (enabled/option/active):  Yes/No/No
Start time:          Mon Apr 18 09:43:42 2016
End time:            Mon Apr 18 09:43:51 2016
Status:              PASSED
=====
Test          Start time    Result          EST(HH:MM:SS)    Comments
=====
Electrical loopback  09:43:42    PASSED         -----    -----
Optical loopback     09:43:44    PASSED         -----    -----
Link traffic test    09:43:47    PASSED         -----    -----
=====
Roundtrip link latency:      1113 nano-seconds
Approximate cable distance:  unknown
Buffers required:           1 (for 2112 byte frames at 32Gbps speed)
Egress power:               Tx: -18.7dBm, Rx: -0.5 dBm, Diff: 0.0 dBm (No Loss)
Ingress power:              Rx: -1.6 dBm, Tx: Not Avail.
62:admin>

```

Supported platforms for D_Ports

The ports use Brocade-branded or certain Brocade-qualified SFP or QSFP transceivers. D_Port functionality is supported on the following Gen 5 and Gen 6 Fibre Channel (FC) platforms.

TABLE 45 Supported transceivers for D_Ports with Fabric OS

Transceiver	SFP	QSFP
Gen 5	10/16-Gbps FC 8/16-Gbps Long Wave Length (LWL) FC 8/16-Gbps Extended Long Wave Length (ELWL) FC - SmartOptics 16-Gbps 40-km SFP+ (see Note below)	16 Gbps
Gen 6	10/16/32-Gbps FC	16/32 Gbps

TABLE 45 Supported transceivers for D_Ports with Fabric OS (continued)

Transceiver	SFP	QSFP
	16/32-Gbps LWL FC 8/16-Gbps ELWL FC	

NOTE

The part numbers for the supported SmartOptics SFP+ transceivers with support for electrical wrap and optical wrap are P/N 16G-ER-D xxx-BR2.

In Fabric OS 8.2.0, support is provided for electrical and optical loopback tests only on 16Gbps BR2 Smartoptics SFP that supports electrical wrap and optical wrap.

2KM 32Gbps QSFP's only supports link traffic test.

The following table lists the Brocade devices and Fabric OS releases that support D_Ports.

TABLE 46 Supported platforms for D_Ports with Fabric OS

Product	Fabric OS release and later
Brocade DCX 8510-4 Backbone	7.0.0
Brocade DCX 8510-8 Backbone	7.0.0
Brocade 6510 Switch	7.0.0
Brocade 6505 Switch	7.0.1
Brocade 6520 Switch	7.1.0
Brocade 7840 Extension Switch	7.3.0
Brocade G620 Switch	8.0.0
Brocade X6-4 Director	8.0.1
Brocade X6-8 Director	8.0.1
Brocade G610 Switch	8.1.0
Brocade G630 Switch	8.2.0
Brocade 7810 Extension Switch	8.2.1

D_Port functionality is supported on the following HBAs:

- Brocade 16-Gbps HBA (Brocade Fabric Adapter 1860) ports operating in HBA mode with a 16-Gbps SFP+ on Brocade 16-Gbps switches running Fabric OS version 7.1 or later.
- Non-Brocade 16-Gbps HBAs: Emulex LPe16002B-M6, QLogic QLE-2672, QLogic 1860-2P
- Non-Brocade 32-Gbps HBAs: Emulex LPe32000B, QLogic QLE2764, QLogic QLE2742

Brocade HBA v3.1 provides limited support for D_Ports. Brocade HBA v3.2 provides extensive support for D_Ports, including dynamic D_Port mode.

- Non-Brocade 16-Gbps HBAs (must have HBA D_Port support)

For applicable topologies, refer to the "Supported topologies" section.

For D_Port functionality with 8-Gbps SFP transceivers, the switch must be running Fabric OS 7.2.1 or later.

To run D_Port tests on a link with 8-Gbps SFP transceivers, both ends of the link must have the same type of SFP. That is, both ends of the link must have LWL SFPs or both ends must have ELWL SFPs.

Licensing requirements for D_Ports

The D_Port feature does not require a license if you are running tests between a pair of Brocade devices, whether the devices are switches, Access Gateways, or HBAs.

If you want to run D_Port tests between a switch and a non-Brocade HBA, the Fabric Vision license for Gen 6 and Gen 5 or the combination of the Fabric Watch license and the Advanced Performance Monitoring license for Gen 5 only is required. Also, the HBA vendor must have implemented Brocade HBA D_Port support.

In Fabric OS 7.3.2 and later, you can also run D_Port tests if you have a Fabric Watch and Performance Monitor combo license.

Understanding D_Ports


A port in D_Port mode does not carry any user traffic, and is designed to run only specific diagnostics tests for identifying link-level faults or failures.

The following figure illustrates an example D_Port connection between a pair of switches through SFP transceivers (port assignments will vary).

NOTE

D_Port functionality is supported only on 16-Gbps (Gen 5) and 32-Gbps (Gen 6) platforms; it is not supported on 8-Gbps (Gen 4) platforms. However, the user can configure the link speed on a Gen 5 or Gen 6 platform to come up at 8 Gbps, for example, as D_Port functionality remains available on links at lower speeds on those platforms.

FIGURE 3 Example of a basic D_Port connection between switches



To bring up a port in D_Port mode, follow these basic steps:

1. Disable the ISL capability on the ports. (**portdisable port**)
2. Enable D_Port functionality on both ends of the link. (**portcfgdport --enable port**)
3. Enable the ISL capability on the ports. (**portenable port**)

NOTE

Detailed configuration examples are presented later in this chapter.

Once the ports are configured and enabled as D_Ports, the following basic test suite is executed in the following order, depending on the SFPs installed:

1. Electrical loopback (with 16/32-Gbps SFP+ and 32-Gbps QSFP) (See the following Note.)
2. Optical loopback (with 16/32-Gbps SFP+ and 32-Gbps QSFP) (See the following Note.)
3. Link traffic

4. Link latency, distance, and power measurement (with 8-Gbps SFPs, 10-Gbps SFPs, 16/32-Gbps SFP+ and QSFP+, G620 QSFP+)

NOTE

32G QSFPs having a serial number beginning with ZTA11547 or newer to supports only link traffic test in 8.0.1.

NOTE

For loopback testing, both direct (port-to-port) optical cabling and optical loopback (by means of an optical loopback connector) are supported. External loopback tests require a loopback connector. When non-Brocade HBAs are used, refer to the manufacturer for supported hardware and additional testing details.

NOTE

Electrical and optical loopback tests are not supported on 8-Gbps LWL/ELWL SFPs, 10-Gbps SFPs, or QSFP/QSFP+ FC16-64 ports.

The user configures the desired ports on either end, or both ends, of the connection. (The default D_Port mode of a port is dynamic mode. If the user configures one end as D_Port static mode, then the other end behaves as a D_Port. Alternatively, the user can configure both ends as D_Port static mode.) Once both sides are configured, a basic test suite is initiated automatically when the link comes online. After the automatic test is complete, the user can view results through the CLI or a GUI and rectify issues (if any) that are reported. The user can also start (and restart) the test manually to verify the link.

D_Port configuration modes and testing

A D_Port can be configured in one of three modes:

- **Static** — This explicitly configures the port as a D_Port. In this mode, the port remains a D_Port until you explicitly remove the D_Port configuration.
- **Dynamic** — The port is automatically set to a D_Port based on an external request from a remote port on the other end of the connection. In this mode, the port remains a D_Port until all the diagnostic tests are completed and the remote port reverts to normal mode. For the port to become a dynamic D_Port, the remote port on the other end of the connection must be either a static D_Port or an on-demand D_Port. Dynamic D_Port mode is supported on connections between a switch and a Host Bus Adapter (HBA), or an Access Gateway and an HBA/device, ISLs, and ICLs.

By default, a switch has the capability to support dynamic D_Port mode. You can disable this capability by using the **configure** command, as shown in the following example.

```
switch:admin> configure
Configure...

Fabric parameters (yes, y, no, n): [no]
Virtual Channel parameters (yes, y, no, n): [no]
F-Port login parameters (yes, y, no, n): [no]
D-Port parameters (yes, y, no, n): [no] y

Dynamic D-Port (on, off): [on] off
```

NOTE

Changes are reflected in the RASLog.

- **On-demand** — This mode is off by default and must be enabled. The port then becomes a D_Port as a result of an internal request or event within the local switch such as the **slotpoweroff**, **slotpoweron**, **slot insert**, **portcfgPersistentDisable** and **portcfgPersistentEnable** commands. In this mode, the port remains a D_Port until all the diagnostic tests are completed successfully. If any of the tests fail, the port continues to remain a D_Port. For a switch port to work as an on-demand D_Port, the other end of the connection must support dynamic D_Port capability. With Fabric OS 7.3.0 on switches and chassis, on-

demand D_Port mode can be configured by means of a switch-wide command, and is supported by default on switches and chassis (ISLs and ICLs). The following internal events within a switch or chassis can trigger a port to become a D_Port:

- **slotpoweroff** and **slotpoweron**
- Slot or blade insert
- **portcfgPersistentDisable** and **portcfgPersistentEnable** commands

Once on-demand D_Port configuration is enabled on the switch and any of the on-demand triggers have been issued (such as **slotpoweroff**, **slotpoweron**, **slot insert**), then the port comes up in on-demand D_Port mode.

If dynamic D_Port is supported, the on-demand D_Port forces the remote port to dynamic D_Port mode, triggers the diagnostic tests automatically, and then changes back to normal port mode after successful completion of the tests.

By default, a switch does not support on-demand D_Port mode. You can turn this capability on by using the **configure** command, as shown in the following example.

```
switch:admin> configure
Configure...

Fabric parameters (yes, y, no, n): [no]
Virtual Channel parameters (yes, y, no, n): [no]
F-Port login parameters (yes, y, no, n): [no]
D-Port parameters (yes, y, no, n): [no] y

Dynamic D-Port (on, off): [on]
On Demand D-Port (on, off): [off] on
```

NOTE

Even if the on-demand D_Port option is enabled, if any static D_Port configuration is enabled on a port, then that static configuration takes precedence.

The following table summarizes D_Port test initiation modes and test start behavior.

TABLE 47 D_Port configuration mode and nature of test

D_Port mode/nature of test		Description
Mode	Static	You must configure the port explicitly. Port remains a D_Port until you remove the configuration.
	Dynamic	No user configuration is required. D_Port mode is initiated by external request/event from the remote port. The remote port can either be a static or on-demand D_Port.
	On-demand	No user configuration is required. D_Port mode is initiated by internal request within the local switch. The remote port should be a dynamic D_Port.
Nature of test	Automatic	Test automatically starts when the port comes online.
	Manual	User starts test from the switch or Access Gateway side (using the portdporttest command with either the --start or --restart keyword), or from the HBA side (refer to "BCU D_Port commands" in this guide) for Brocade HBAs.

When the tests complete, the port behavior depends on the mode:

- For static D_Ports, you must remove the D_Port configuration at either or both ends of the link to bring up the port as a regular E_Port or F_Port.
- For a switch port in dynamic D_Port mode, the port automatically reverts back to an E_Port or an F_Port if the port at the other end reverts to a regular port.
- For a switch in on-demand D_Port mode, the port automatically reverts back to an E_Port after the tests are completed successfully.

- Beginning with Fabric OS 8.0.1, the following test options is available only for optical loopback tests: `-framesize`, `-nframes`, `-pattern`, `-payload`, `-time`. Both ends must be running the same release (Fabric OS 8.0.1 or later) or the test reverts to legacy behavior.
- By default, Forward Error Correction (FEC) is enabled on 32-Gbps and always the D-port test will run with FEC enabled .

General limitations and considerations for D_Port tests

Consider the following issues when running D_Port tests:

- The link to be tested must be marginally functional and able to carry a minimal number of frames before it can become a D_Port link.
- D_Port testing is useful for diagnosing marginal faults only. A complete failure of any component cannot be detected.
- D_Port configuration is not supported on mezzanine cards.
- D_Ports do not support a loop topology.
- D_Port testing is not supported on adapter ports configured in CNA mode.
- Toggling the port on either side of the link during testing may result in a test failure.
- On ICL ports and FC16-64 blade ports that are configured as D_Ports, only link traffic tests can be run. Electrical loopback and optical loopback tests are not supported. A connection between a QSFP port of an FC16-64 and a core blade is not a valid connection.
- When a large number of D_Ports are configured, the test is run on one port per blade at a time, and other ports wait until the test is completed. No test begins until the fabric is stable.
- Before running On-demand D_Port tests on the links having Dense Wavelength Division Multiplexing (DWDM) or Course Wavelength Division Multiplexing (CWDM), you must provision the port by using the **portcfgdport** command with **-dwdm** option. Optical loopback tests are skipped for D_Ports with DWDM.
- The D_Port DWDM provisioning is applicable to static D_Ports, On-demand D_Ports, and Dynamic D_Ports. During the configuration of the Static D_Port using **portcfgdport --enable**, DWDM is automatically configured if the port was previously provisioned for DWDM.
- When you run a D_Port test on the links between a FC16-64 port blade and a fixed-port switch or blade, run the test on one link at time for short distance links. If you have 100-km links, you should start the test on other links only after the 100-km link test is completed.
- You can run the D_Port test on a maximum of eight links at a time. If you run on more than eight links simultaneously, false alarms may be raised.
- In case of switch-to-Host Bus Adapter (HBA) or Access Gateway-to-HBA connections with Brocade HBA v3.2.3 or later, Brocade recommends that D_Port tests be limited to a maximum of eight D_Ports at once. Otherwise, there is a possibility of false alarms.
- If a remote SFP is not capable of running an optical loopback test or a QSFP is connected by means of a breakout cable to an SFP, the test is skipped with the reason "No remote SFP or chip support" .
- During a dynamic D_Port configuration, when electrical or optical loopback testing is in progress and the switch is rebooted or there is a failover, it is possible for the ports to end up in an unsynchronized state. To clear this condition, enter the **portDportTest --clear all** command on a D_Port, an E_Port, or both, as appropriate.
- Beginning with Fabric OS 8.0.1, it is possible to specify the number of transmitted frames and the duration of the test for optical loopback test.
- D_port test will fail, if each switch in the test has the same domain ID. When you run the D-port test on ISLs and spinfab, each switch should have unique domain IDs.

NOTE

Long-duration optical loopback tests can be run on only one port at a time. Long-duration electrical loopback tests are not supported.

NOTE

D_Port is not supported on ethernet port.

- Refer to the following topics in this guide for additional specific information:
 - "8 Gbps LWL and ELWL SFP transceiver limitations for D_Ports"
 - "High Availability limitations and considerations for D_Ports"
 - "Access Gateway limitations and considerations for D_Ports"
 - "Host Bus Adapter limitations and considerations for D_Ports"

Access Gateway limitations and considerations for D_Ports

In addition to the items listed in "General limitations and considerations for D_Ports," you should keep in mind the following limitations and considerations when using a D_Ports with Access Gateways:

- D_Ports on an Access Gateway are supported only when there is no mapping between F_Ports and N_Ports; this includes static mapping and preferred port mapping. In addition, device (WWN) mapping is also not retained when a D_Port is used. If an Access Gateway port to be tested is mapped, the port mapping (including static and preferred port mapping) must be removed before the D_Port can be used. Refer to "Saving port mappings on an Access Gateway" in this chapter.
- Access Gateway supports D_Port dynamic mode only with HBAs. If the port on the connected HBA is configured as a D_Port, the Access Gateway port automatically changes to D_Port mode. Once the Access Gateway port becomes a D_Port, the tests are run automatically, and then when the HBA becomes a normal port, the Access Gateway port also switches to normal F_Port mode. However, Access Gateway-to-Access Gateway and Access Gateway-to-switch port connections require static configuration to become D_Ports.
- If a pre-Fabric OS 7.3.0 static D_Port is connected to a Fabric OS 7.3.0 and later port that is not a D_Port, and dynamic D_Port mode is not enabled, credit loss may happen on the non-D_Port port.
- If a switch makes an on-demand D_Port request to a port on an Access Gateway running a version of Fabric OS earlier than 7.3.0 that is not a D_Port, the Access Gateway will disable the port. This restriction applies only to this combination of switch and Access Gateway.

High Availability limitations and considerations for D_Ports

Consider the following High Availability (HA) limitations and considerations when using D_Ports:

- There is no HA support for D_Port test options and results. Any information from a previous test is lost following a failover or reboot.
- During an HA failover reboot on one side of the link, the link is reinitialized and may restart the test. However, the test cannot proceed if the remote port is not ready to proceed further (the remote port may already be done with the D_Port test and in the final state). In such a case, the test will eventually fail with a "Remote port not ready" message. Restarting the test from either side will recover the port.

8-Gbps LWL and ELWL SFP transceiver limitations for D_Ports

Consider the following 8-Gbps LWL and ELWL SFP transceiver limitations when using D_Ports:

- On these transceivers, only link traffic tests can be run. Electrical and optical loopback tests are not available on 8-Gbps SFP transceivers.

- On these transceivers, if the cable length is less than or equal to 100 meters, the length is displayed as “unknown”. If the cable length is greater than 100 meters, the length is displayed accurately.

Support for audit logs

In previous releases, log messages were not generated when a port was either configured or unconfigured as a D_Port.

The **switchShow** command was the only option available to verify whether a port was configured as a D_Port or not. With Fabric OS 8.1.0, the following log messages are provided.

NOTE

The following messages are supported only for static D_Port configurations.

This example log message is generated when a port is configured as a D_Port.

```
624 AUDIT, 2016/04/21-10:36:06 (UTC), [FABR-1075], INFO, RAS, admin/admin/172.26.3.151/telnet/CLI, ad_0/Dport_DCX/FID 128,, Port is configured as D_port.
```


This example log message is generated when a port is unconfigured as a D_Port.

```
625 AUDIT, 2016/04/21-10:36:06 (UTC), [FABR-1075], INFO, RAS, admin/admin/172.26.3.151/telnet/CLI, ad_0/Dport_DCX/FID 128,, Port is not configured as D_port.
```

Topology 1: ISLs

The following figure illustrates inter-switch links (ISLs) that connect multiple switches through a pair of chassis. The letter E represents E_Ports to be configured as D_Ports.

FIGURE 4 ISLs connecting multiple switches and chassis




Note the following:

- Only static-static and static-dynamic D_Port modes are supported on the ISLs.

Topology 2: ICLs

The following figure illustrates inter-chassis links (ICLs) between slots 5 and 8 in corresponding chassis. The letter E represents E_Ports to be configured as D_Ports.

FIGURE 5 ICLs connecting chassis blades




Static-static, static-dynamic, and on-demand-dynamic D_Port modes are also supported on the ICLs.

Topology 3: Access Gateways

The following figure illustrates a switch configured as a single Access Gateway connected to a fabric switch. The letters N and F represent, respectively, an N_Port and an F_Port to be configured as D_Ports.

The Access Gateway must be a Brocade 6505 or Brocade 6510 or Brocade G610 or Brocade G620.


FIGURE 6 Single Access Gateway to switch



The above topology is supported only with static-static D_Port modes.

The following figure illustrates multiple Access Gateways connected to a switch in a cascaded topology. The letters N and F represent, respectively, an N_Port and an F_Port to be configured as D_Ports.

FIGURE 7 Multiple Access Gateways cascaded to switch




The above topology is supported only with static-static D_Port modes.

NOTE

N_Port-to-F_Port and device (WWN) mappings must be removed from an Access Gateway port before configuring the Access Gateway port as a D_Port. Refer to [Saving port mappings on an Access Gateway](#) on page 109.

The following figure illustrates connectivity between an HBA and an Access Gateway. The letter F represents an F_Port to be configured as a D_Port.

FIGURE 8 Access Gateway to HBA

Static-static and static (HBA) - dynamic (AG) D_Port modes are supported.

Saving port mappings on an Access Gateway


Before configuring ports as D_Ports on a switch configured as an Access Gateway, you must remove N_Port-to-F_Port and device (WWN) mappings. Fabric OS commands are available to save N_Port mappings. Once you save them, you can display the saved N_Port mappings to reconfigure them after D_Port is disabled. A command is also available to delete saved N_Port mappings.

For more details, refer to Chapter 2, "Configuring Ports in Access Gateway Mode," in the Brocade Access Gateway Administration Guide.

Topology 4: HBA to switch

The following figure illustrates connectivity between an HBA and a switch. The letter "F" represents an F_Port to be configured as a D_Port. This topology supports static-static and dynamic (switch) - static (HBA) D_Port mode. In dynamic mode, the switch port does not need to be configured explicitly as a D_Port. It comes up in D_Port mode when it receives a request from the remote port.

FIGURE 9 HBA to switch



For configuration details, refer to "Using D_Port with HBAs" in this chapter.

Using a D_Port in static-static mode between switches

You can configure D_Ports in static-static modes between switches (ISLs), chassis (ICLs), Access Gateways, and switch-Access Gateway links.

The sections "Enabling D_Port" and "Disabling D_Port" apply to topologies 1, 2, 3, and 4:

Enabling a D_Port in static mode

Use this procedure to configure a basic, automatic D_Port diagnostics session in static mode between two switches. The summary steps are as follows:

1. Disable ISL functionality on the ports.
2. Configure D_Port static mode on both ends of the link.
3. Enable ISL functionality on the ports.

ATTENTION

The automatic test might fail if you do not follow the sequence of steps exactly.

NOTE

"Port 1" and "Port 2" simply represent corresponding peer ports at opposite ends of the link to be tested. Switch A and Switch B can be the same platform or different platforms supporting an end-to-end D_Port connection.

The detailed steps are as follows:

1. Disable Port 1 on Switch A.

```
switchA:admin> portdisable 1
```

2. Configure Port 1 on Switch A as a D_Port in static mode.

```
switchA:admin> portcfgdport --enable 1
```

3. Repeat Step 1 and Step 2 for the corresponding port (in this example, Port 2) on Switch B.

```
switchB:admin> portdisable 2
switchB:admin> portcfgdport --enable 2
```

4. Re-enable Port 1 on Switch A.

```
switchA:admin> portenable 1
```

5. Re-enable Port 2 on Switch B.

```
switchB:admin> portenable 2
```

The basic test suite starts as soon as both ports are enabled and ready to perform the test.

6. While the test is running, you can enter a variety of **show** commands to confirm the D_Port configuration and view test results on an interface. See "Using D_Port show commands" at the end of this chapter.

To view D_Port status and test results for an interface:

```
switch:admin> portdporttest --show 10/0/39
```

D-Port Information:

=====

```
Port: 7
Remote WWNN: 10:00:00:05:33:81:43:00
Remote port index: 71
Mode: Automatic
No. of test frames: 1 Million
Test frame size: 1024 Bytes
FEC (enabled/option/active): Yes/No/No
CR (enabled/option/active): Yes/No/No
Start time: Sun Dec 7 22:49:05 2014
End time: Sun Dec 7 22:49:30 2014
Status: PASSED
```

```
=====
Test Start time Result EST(HH:MM:SS) Comments
=====
```

```
Electrical loopback 22:49:07 PASSED -----
Optical loopback 22:49:21 PASSED -----
Link traffic test 22:49:26 PASSED -----
=====
```

```
Roundtrip link latency: 277 nano-seconds
Estimated cable distance: 3 meters
Buffers required: 1 (for 2112 byte frames at 16Gbps speed)
Egress power: Tx: -2.6 dBm, Rx: -3.3 dBm, Diff: 0.7 dBm (Loss is within tolerable limit)
Ingress power: Rx: -2.5 dBm, Tx: -2.7 dBm, Diff: 0.0 dBm (No Loss)
```

- To confirm D_Port test results for all ports, use the following command.

```
switch:admin> portdporttest --show all
```

Port	State	SFP Capabilities	Test Result
24	ONLINE	E,O	PASSED
26	ONLINE	E,O	PASSED
33	OFFLINE	---	FAILED

- Optional* : To stop the test on both ends, use the **portDportTest --stop** command on both the ends. After stopping the test, verify manually that the test has stopped on both ends. You can restart the test on both ends using the **portDportTest --start/restart** command on any one of the ends. You do not need to run the same command on both ends of a single connection

Disabling a D_Port in static mode

Use this procedure to disable a D_Port diagnostics session in static mode, as configured in "Enabling D_Port in static mode."

NOTE

"Port 1" and "Port 2" simply represent corresponding peer ports at opposite ends of the link to be tested.

- Disable Port 1 on Switch A.

```
switchA:admin> portdisable 1
```

- Disable the D_Port functionality in static mode on Port 1 on Switch A.

```
switchA:admin> portcfgdport --disable 1
```

- Repeat Steps 1 and Step 2 for Port 2 on Switch B.

```
switchB:admin> portdisable 2
switchB:admin> portcfgdport --disable 2
```

- Reenable Port 1 on Switch A.

```
switchA:admin> portenable 1
```

- Reenable Port 2 on Switch B.

```
switchB:admin> portenable 2
```

Using D_Ports in dynamic mode

Enabling dynamic D_Port switch-wide configuration forces the ports on that switch or chassis to respond to D_Port requests from the other end of the connection. It basically responds to a remote port request to change its mode to D_Port mode, and run diagnostic tests automatically. For more information on enabling dynamic D_Port mode for all ports in a switch or chassis, refer to "D_Port configuration mode and testing" in this chapter.

Preprovisioning D_Ports

In a normal scenario, you must disable a port before configuring it as a static D_Port. This is a disruptive operation and any command or procedure error could cause the wrong port to go down. To avoid such errors, you can preprovision a port that needs to be configured as a D_Port. Any preprovisioned port does not require to be disabled before being configured as a D_Port.

1. To add one or more ports to the D_Port provision list, use the **portcfgdport --provision -add** *[slot/] port_list* command. If the **-dwdm** option is specified, the ports in the provision list are provisioned for D_Ports over links with DWDM or CWDM.

With the preprovisioning of D_Ports, links with DWDM or CWDM can support dynamic D_Port testing and on-demand D_Port testing.

The following example shows how to provision a port with the DWDM option.

```
Switch:admin> portcfgdport --provision -add -dwdm 7/16
```

The following example shows how to provision a port without the DWDM option.

```
Switch:admin> portcfgdport --provision -add 7/20
```

2. (Optional). To remove one or more ports from the provision list, use the **portcfgdport --provision -delete** *[slot/] port_list* command. Specifying the **-dwdm** is optional. When a port is removed from the provision list, the D_Port DWDM provisioning is automatically cleared.

The following example shows how to delete a port with the DWDM option.

```
Switch:admin> portcfgdport --provision -delete -dwdm 4/12
```

The following example shows how to delete a port without the DWDM option.

```
Switch:admin> portcfgdport --provision -delete 4/16
```

3. To list the ports in the provision list, use the **portcfgdport --provision -show** command, as in the following example.

```
Switch:admin> portcfgdport --provision -show
Slot Port D-Port provision DWDM
=====
  3    4      ON          OFF
  3    5      ON          OFF

Switch:admin> portcfgdport --provision -show 3/2-5
Slot Port D-Port provision DWDM
=====
  3    2      OFF         OFF
  3    3      OFF         OFF
  3    4      ON          OFF
  3    5      ON          OFF
```

The following example shows a port that is not preprovisioned.

```
Switch:admin> portcfgdport --provision -show 3/2
Slot Port D-Port provision DWDM
=====
  3    2      OFF         OFF
```

4. To enable D_Port mode on a preprovisioned port, use the **portcfgdport --enable** command. You do not need to disable the port before running this command. However, this command will display an error message if the particular port is not preprovisioned using the previous steps.

Specifying **-dwdm** is optional. If the port is preprovisioned as DWDM, the port is automatically configured as DWDM.

The following example enables the port as D_Port.

```
Switch:admin> portcfgdport --enable 4/10
```

The following example enables the port as DWDM.

```
Switch:admin> portcfgdport --enable -dwdm 5
```

5. To disable D_Port mode on a preprovisioned port, use the **portcfgdport --disable** command.

The following example disables the port as D_Port.

```
Switch:admin> portcfgdport --disable 4/10
```

The following example disables the port as DWDM.

```
Switch:admin> portcfgdport --disable -dwdm 5
```

Using D_Port mode between switches and HBAs

When HBAs are used, D_Port mode initiates electrical loopback, optical loopback, and link-traffic diagnostic tests on the link between the HBA and the connected switch port. Results can be viewed from the switch by means of Fabric OS commands and from the Brocade HBA by way of the Brocade Command Line Utility (BCU) and Brocade Host Connectivity Manager (HCM) during or after the test. Once in D_Port mode, the adapter port does not participate in fabric operations, log in to a remote device, or run data traffic.

NOTE

The syntax in this section applies to Brocade HBAs only. Where HBAs are supplied by other vendors, view the appropriate GUI or CLI interface to see the results on the HBA side, and use the commands as provided by the vendor.

In dynamic D_Port mode, you can disable the physical port by using the **bcu port --disable** command. However, that command will not exit dynamic D_Port mode. When you enable the port again, the switch will again force the adapter port into D_Port mode if the switch port is still enabled as a D_Port.

The following sections apply to "Topology 4: HBA to switch":

- "Automatic mode configuration"
- "Dynamic mode configuration"
- "BCU D_Port commands"

Enabling a D_Port in static mode between a switch and an HBA

This procedure enables a D_Port diagnostic session from the connected switch to an HBA. After the default test suite is run automatically, you can run specific tests manually to obtain additional detail.

1. Disable the switch port by using the **portDisable** command.
2. Configure the switch port that you want to enable as a D_Port by using the **portCfgDport --enable** command.
3. Disable the adapter port by using the adapter **bcu port --disable** command.
4. Enable the switch port by using the **portEnable** command.

5. Enable the adapter port as a D_Port by using the adapter **bcu diag --dportenable** command and configure test parameters.
For more details on adapter configuration, refer to the *Brocade Adapters Administrator's Guide*.
6. Enable the switch port.

Using non-Brocade HBAs for D_Port testing

This section illustrates how to enable Emulex or QLogic HBA cards for D_Port testing.

Cards must be installed in a server according to the manufacturer's instructions and connected to a Brocade switch.

Support is provided for non-Brocade HBAs beginning with Fabric OS 8.0.1 on Gen 6 platforms.

Enabling D_Port testing with an Emulex HBA

This task illustrates how to enable D_Port testing with an Emulex HBA.

D_Port testing is supported on Emulex 16G and 32G cards. For the latest firmware and product-specific details, go to www.emulex.com.

1. Configure the Brocade switch into dynamic D_Port mode as follows.
 - a) Disable the switch by entering the **switchDisable** command.

NOTE

Dynamic D_Port is enabled by default.

```
switch:admin> switchdisable
```

- b) Enter the **configure** command, and then select "D-Port Parameters" and ensure that dynamic D_port testing is enabled and on-demand port testing is disabled.

Configure...

```
Fabric parameters (yes, y, no, n): [no]
Virtual Channel parameters (yes, y, no, n): [no]
F-Port login parameters (yes, y, no, n): [no]
D-Port Parameters (yes, y, no, n): [no] y
```

```
Dynamic D-Port (on, off): on
On Demand D-Port (on, off): off
```

2. Re-enable the switch by using the **switchEnable** command.


```
switch:admin> switchenable
```
3. From the host server, start D_Port testing by selecting the Emulex "OneCommand Manager."
4. For each port under test, click the **Diagnostics** tab, then click **D_Port Tests**.

Enabling D_Port testing with a QLogic HBA

This task illustrates how to enable D_Port testing with a QLogic HBA.

D_Port testing is supported on QLogic 16G and 32G cards. For the latest firmware/driver and product-specific details, go to www.qlogic.com. If the latest firmware/driver supporting D_Port testing is installed in the server, nothing further needs to be done on the server side.

NOTE

The QLogic cards support only electrical and optical loopback tests. They do not support link traffic tests.

1. On the Brocade switch, convert the port from an F_Port to a D_Port, disabling the specified F_Port by entering the **portDisable** *port_ID* command, as in the following example.

```
switch:admin> portdisable 1/1
```

2. Enter the **portCfgDport --enable** *port_ID* command.

```
switch:admin> portcfgdport --enable 1/1
```

3. Enable the port by using the **portEnable** *port_ID* command.

```
switch:admin> portenable 1/1
```

The D_Port test begins automatically.

BCU D_Port commands

The following BCU commands can be used for D_Port configuration and control:

NOTE

These commands are for Brocade HBAs only. Refer to documentation from other vendors as appropriate.

- **bcu diag --dportenable** -- Enables D_Port on a specific port, sets the test pattern, and sets the frame count for testing.
- **bcu diag --dportdisable** -- Disables D_Port on a specific port and sets the port back to an N_Port or NL_Port.
- **bcu diag --dportshow** -- Displays test results for a test in progress on a specific port.
- **bcu diag --dportstart** -- Restarts a test on a specific port when the test has completed.
- **bcu port --list** -- Displays the D_Port enabled or disabled state on the adapter and connected switch.

Host Bus Adapter limitations and considerations for D_Ports

In addition to the items listed in [General limitations and considerations for D_Port tests](#) on page 104, you should keep in mind the following limitations and considerations when using a D_Port with a Host Bus Adapter (HBA):

- D_Port is supported only on Brocade 16-Gbps HBA ports operating in HBA mode with a 16-Gbps SFP+ on Brocade 16-Gbps switches running Fabric OS 7.1 or later. In addition, the Brocade adapter must be using driver version 3.2.0 or higher.
- D_Port is supported on non-Brocade 16-Gbps HBAs on Gen 5 platforms if
 - You have a Fabric Vision license or the combination of Fabric Watch license and Advanced Performance Monitoring license present on the switch and
 - The HBA vendor has implemented the Brocade HBA D_Port support.
- D_Port is supported on non-Brocade 32-Gbps HBAs on Gen 6 platforms if you have a Fabric Vision license present on the switch.
- D_Ports are not supported in a configuration of an HBA to another HBA (in target mode).
- D_Ports on the HBA do not support forward error correction (FEC) and credit recovery (CR). If these features are enabled on the switch side, the HBA ignores them.
- Because of SFP electrical wrap (EWRAP) bleed-through, during the beginning of switch electrical loopback testing the HBA will receive some broken frames. This causes the port statistic error counter to increase. Examples are “CRC err”, “bad EOF”, and “invalid order set”. Similar results occur for the optical loopback test. You should ignore these port statistics on the HBA.

- The following commands from the switch are not supported by the HBA, and the HBA will drop them:
 - portdporttest --restart**
 - portdporttest --setarg**
- The Qlogic HBA does not support manual D_Port tests that are started from the switch side by means of the **portdporttest** command with suboptions (--start, --setargs, --restart, and so on). It supports only automatic tests when the port is configured as a D_Port by means of the **portcfgdport --enable** command from the switch side.
- When Emulex HBAs are used, D_Port tests can be triggered only from the Emulex HBA side only. The supported D_Port configuration is "Switch (Dynamic D_Port) /Emulex HBA (On-Demand D-port)". Static D_Port configuration is not supported in this case.
- The maximum number of D_Ports on which the tests can run simultaneously depends on the HBA firmware version.

TABLE 48 Limitation on number of D_Ports for simultaneous tests

HBA firmware version	Maximum number of D_Ports on which tests can be run simultaneously
HBA v3.2.0	4
HBA v3.2.3	8

- You must configure D_Port using **portcfgdport** command from the switch side and then start the test from the HBA side for 32 Gb/s QSFP.
- As soon as the HBA is configured as the static D_Port, the switch changes to dynamic D_Port mode. After the test is completed and the switch is rebooted, the switch port changes back to G_Port mode. To resolve this issue, remove the static D_Port configuration on the HBA.
- Powering off and on or plugging in and out slots containing ports in D_Port mode results in those ports losing the dynamic D_Port state when the slot or port is back up. If this happens, you must reconfigure the static D_Port mode on the HBA.
- D_Port is supported on Qlogic® HBAs only at link speed of 32-Gbps and 16-Gbps.

The following considerations apply to 32-Gbps SFPs and QSFPs.

With the introduction of the new SFPs and QSFPs, FEC and FEC with Transmitter Training Signal (TTS) mode is enabled by default for 32 Gbps. FEC cannot be disabled in this case. The following table summarizes the state of the FEC fields and the resulting behavior.

TABLE 49 FEC field state and behavior for 32 Gbps tests

Field	State/Behavior
Enabled	State is "Yes" if port is FEC capable and FEC is configured by means of the portcfgfec command. For 32-Gbps SFP/QSFPs at link speed of 32 Gbps, FEC is enabled by default and cannot be disabled.
Option	State is "Yes" if portdporttest command is used with the -fec option. For a link speed of 32 Gbps, FEC is auto enabled, irrespective of this option.
Active	State is always "Yes" for 32 Gbps alone. This depends on the capability of local and remote ports and does not depend on the -fec option alone.

The output of the **portdporttest --show** command displays "FEC = Yes" for both Enabled and Active for 32 Gbps only. The output is not changed for port speeds of 10 or 16 Gbps.

Confirming SFP and link status with an HBA

The steps in the following example illustrate how the **bcu diag --dportenable** command will fail with an SFP installed but without a connection to the switch.

1. Confirm the initial port status.

```
# bcu port --list
```

Port#	FN	Type	PWWN/MAC	FC Addr/ Eth dev	Media	State	Spd
1/0	-	fc	10:00:8c:7c:ff:1c:e9:00	160000	sw	Linkup	16G*
	0	fc	10:00:8c:7c:ff:1c:e9:00	160000	sw	Linkup	16G*
1/1	-	fc	10:00:8c:7c:ff:1c:e9:01	--	sw	Linkdown	---
	1	fc	10:00:8c:7c:ff:1c:e9:01	--	sw	Linkdown	---

2. Disable the port.

```
# bcu port --disable 1/0
port disabled
```

3. Remove the connection to the switch and attempt to enable the D_Port.

```
# bcu diag --dportenable 1/0
ERROR: Timer expired - Retry if persists contact support
```

4. Install an SFP and attempt to enable the D_Port.

```
# bcu diag --dportenable 1/0
ERROR: Switch port is not D_Port capable or D_Port is disabled
```

5. Connect to the HBA without the SFP and disable the native port.

```
# bcu port --disable 1/0
port disabled
```

6. Attempt to enable the D_Port.

```
# bcu diag --dportenable 1/0
ERROR: SFP is not present.
D-port will be enabled but it will be operational only after inserting a valid SFP.
```

Using a D_Port in on-demand mode

Enabling on-demand D_Port switch-wide configuration forces the ports on that switch or chassis to respond to an internal requests within the switch as a result of certain events. The switch basically responds to internal request to change a port mode to D_Port mode, and run diagnostic tests automatically. For more information on enabling on-demand D_Port mode for all ports in a switch or chassis, refer to "D_Port configuration modes and nature of test," above.

When an on-demand D_Port-capable switch or chassis comes online, the switch checks if the other end of the connection supports dynamic D_Port mode. If dynamic D_Port is supported on the other end, the switch changes the remote port to D_Port mode, and then triggers diagnostic tests automatically. The D_Ports change to normal port mode after successful completion of the tests.

Using the fabriclog command

This command can be enabled or disabled on either a fabric log or a D_Port log.

During D_Port testing that involves a large number of ports, the logs for specific ports can become “wrapped around.” Log file wrapping overwrites existing data with incoming data when the maximum log file size is exceeded. This creates a problem for logging on ports of interest, particularly if a test fails and logging continues unnecessarily.

To remedy this, a new failstop option has been added to the **fabriclog** command (**fabriclog -f**). When this option is enabled, The D_Port fabric is stopped and disabled when a test fails, making the log available for debugging purposes. To resume logging, the user must reenale logging explicitly.

The complete syntax is as follows. Adding the optional dport keyword applies the command only to D_Port logs.

To display fabric logs: **fabriclog -s | --show dport**

To clear fabric logs: **fabriclog -d | --clear dport**

To disable fabric logging: **fabriclog -c | --disable dport**

NOTE

If the argument **dport** is specified, only D_Port fabric logs are disabled. If it is not specified, only fabric logs are disabled.

To enable fabric logging: **fabriclog -e | --enable dport**

NOTE

If the argument **dport** is specified, only D_Port fabric logs are disabled. If it is not specified, only fabric logs are disabled.

To stop and disable D_Port fabric logging upon failure of the first D_Port test. **fabriclog -f | --failstop dport**

NOTE

The **failstop** option is applicable only when the **dport** argument is applied. If the **failstop** option is set, the D_Port fabric log is stopped and disabled when the first D_Port test fails. This option is cleared automatically when the fabric log for D_Port is reenaled.

The following example shows the effect of this operand.

```
switch:admin> fabriclog -failstop dport
fabriclog failstopset successfully
sw0:FID128:admin> fabriclog -s dport
Time Stamp Input and *Action S, P Sn,PnPort Xid
=====
Switch 0; Thu Oct 6 03:27:13 2016 GMT (GMT+0:00)
03:27:13.301915 DP:D_PORT_FINAL(TEST_START)->D_PORT_INIT A2,P3 A2,P30 NA
03:27:13.302013 MSG_FAB_DIAG_CMD, stage (cold done) A2,P3 A2,P30 NA
03:27:13.302034 DiagCmdSend:Cnt=1,Cmd=136 A2,P3 A2,P30 NA
:
Number of entries: 119
Max number of entries: 2048
fabriclog is disabled
failstopis set
```

To display the help for the command: **fabriclog -h | --help**

Note the following considerations and limitations for this option:

- This option is not synchronized to a standby control processor and must be specified again following a high-availability (HA) failover.
- This option is not persistent and must be specified again on the device following a manual or HA reboot.

- This option stops the D_Port fabric log when the first D_Port test fails. Consequently, the debugging of multiple test failures may not be possible simultaneously, as the fabric log may not be available for subsequent failures. To remedy this, the user must resolve the first test failure before addressing subsequent failures.

Calculating buffers for long-distance cables

Consider the following when making D_Port measurements over long-distance cables.

When D_Port testing is done over long-distance cables, the values returned by the **portbuffershow** command are calculated according to the buffers required, the frame size, and the cable length provided by the user. For D_Ports the "Buffers required" value displayed by the **portdporttest --show** command is calculated according to link latency, link speed, and a fixed frame size of 2112 (the default). If the user-provided values are the same, then the buffer values displayed by both the **portbuffershow** and **portdporttest --show** commands will be the same. If the user-provided values are different, then the displayed values will be different.

With the **portbuffershow** command, the buffers are calculated according to user-provided values as well as those provided by the **portcfglongdistance** command. The buffers required, frame size, and cable length are provided by the user. The default frame size used in this case is 2048.

With the **portdporttest** command, the "Buffers required" value is calculated according to link latency, link speed, and frame size (2112, the default). The link latency is calculated according to the actual cable distance, and not the user-specified value.

Support for audit logs

In previous releases, log messages were not generated when a port was either configured or unconfigured as a D_Port.

The **switchShow** command was the only option available to verify whether a port was configured as a D_Port or not. With Fabric OS 8.1.0, the following log messages are provided.

NOTE

The following messages are supported only for static D_Port configurations.

This example log message is generated when a port is configured as a D_Port.

```
624 AUDIT, 2016/04/21-10:36:06 (UTC), [FABR-1075], INFO, RAS, admin/admin/172.26.3.151/telnet/CLI,
ad_0/Dport_DCX/FID 128,, Port is configured as D_port.
```

This example log message is generated when a port is unconfigured as a D_Port.

```
625 AUDIT, 2016/04/21-10:36:06 (UTC), [FABR-1075], INFO, RAS, admin/admin/172.26.3.151/telnet/CLI,
ad_0/Dport_DCX/FID 128,, Port is not configured as D_port.
```

Using D_Port show commands

This section presents a variety of options for viewing the results of D_Port testing.

In addition to using the **portdporttest** command to start or stop D_Port tests, you can also use it to show a variety of detailed test results.

You can display the complete results from either the responder or the initiator switch. If the initiator switch is running Fabric OS v7.1.x or earlier, the responder displays only the local D_Port results, and you must query the initiator to see the complete results.

The following example shows basic D_Port results for a specified port.

```
device:admin> portdporttest --show 26
D-Port Information:
```



```

=====
Port:                26
Remote WWNN:         10:00:00:05:33:13:2f:b5
Remote port index:   42
Mode:                Automatic
Start time:          Wed Feb  2 01:41:43 2011
End time:            Wed Feb  2 01:43:23 2011
Status:              PASSED
=====
Test                  Start time    Result          EST(secs)    Comments
=====
Electrical loopback    01:42:08      PASSED          --            -----
Optical loopback       01:42:16      PASSED          --            -----
Link traffic test      01:43:15      PASSED          --            -----
=====
Roundtrip link latency: 1108 nano-seconds
Estimated cable distance: 20 meters
Buffers required:      1 (for 1024 byte frames at 16Gbps speed)
Egress power:          Tx:-3.3 dBm, Rx:-3.7 dBm, Loss:0.4 dB (within tolerable limits)
Ingress power:         Rx:-3.5 dBm, Tx:-3.2 dBm, Loss:0.3 dB (within tolerable limits)

```

The following example shows the **portdporttest --show** output for port 26 where the electrical and optical tests pass but the link traffic test fails.

```

device:admin> portdporttest --show 26
D-Port Information:
=====
Port:                26
Remote WWNN:         10:00:00:05:33:13:2f:b5
Remote port index:   42
Mode:                Automatic
Start time:          Wed Feb  2 01:41:43 2011
End time:            Wed Feb  2 01:43:23 2011
Status:              PASSED
=====
Test                  Start time    Result          EST(secs)    Comments
=====
Electrical loopback    01:42:08      PASSED          --            -----
Optical loopback       01:42:16      PASSED          --            -----
Link traffic test      01:43:15      FAILED          --            -----
=====
Roundtrip link latency: 1108 nano-seconds
Estimated cable distance: 20 meters
Buffers required:      1 (for 1024 byte frames at 16Gbps speed)
Egress power:          Tx:-3.3 dBm, Rx: Not Avail
Ingress power:         Rx:-3.5 dBm, Tx: Not Avail

```

Enter **portdporttest --show all** to display the capabilities and test results of all the D_Ports in a switch.

```

switch:admin> portdporttest --show all
Port  State      SFP Capabilities  Test Result
=====
24    ONLINE     E,O               PASSED
26    ONLINE     E,O               FAILED
33    ONLINE     E,O               PASSED

```

Enter **switchshow** to see detailed switch and D_Port information.

```

device:admin> switchshow
switchName:          switch_10
switchType:          109.1
switchState:         Online
switchMode:          Native
switchRole:          Principal
switchDomain:         1
switchId:             fffc01
switchWwn:           10:00:00:05:33:13:2f:b4
zoning:              OFF
switchBeacon:        OFF
FC Router:           OFF

```

```

Allow XISL Use: ON
LS Attributes: [FID: 10, Base Switch: No, Default Switch: No, Address Mode 0]
Index Port Address Media Speed State Proto
=====
 24 24 010000 id N16 Online FC D-Port Loopback->Port 24
 26 26 010200 id N16 Online FC D-Port segmented, (D-Port mode mismatch)
 33 33 010300 id N8 Online FC D-Port 10:00:00:05:33:13:2f:b5

```

The following example illustrates D_Port test output at 32-Gbps.

```

device:admin> portdporttest --show 2
D-Port Information:
=====
Port: 2
Remote WWNN: 10:00:00:27:f8:f0:26:40
Remote port index: 0
Mode: Manual
No. of test frames: 1 Million
Test frame size: 1024 Bytes
FEC (enabled/option/active): Yes/No/Yes
CR (enabled/option/active): No/No/No
Start time: Tue Jun 23 07:27:43 2015
End time: Tue Jun 23 07:28:15 2015
Status: PASSED
=====
Test Start time Result EST (HH:MM:SS) Comments
=====
Electrical loopback 07:27:45 PASSED -----
Optical loopback 07:28:04 PASSED -----
Link traffic test 07:28:10 PASSED -----
=====
Roundtrip link latency: 272 nano-seconds
Estimated cable distance: 2 meters
Buffers required: 1 (for 2112 byte frames at 32gbps speed)
Egress power: Tx: 0.3 dBm, Rx: -3.6 dBm, Diff: 3.9 dBm (Loss is within tolerable limit)
Ingress power: Rx: -4.1 dBm, Tx: 0.3 dBm, Diff: 4.4 dBm (Loss is within tolerable limit)

```

The following example illustrates D_Port test output for longer test duration.

```

device:admin> portdporttest --show 31
D-Port Information:
=====
Port: 31
Remote WWNN: 10:00:50:eb:1a:0d:be:08
Remote port index: 31 (External loopback)
Mode: Manual
Test Duration (HH:MM): 00:01
Test frame size: 1024 Bytes
FEC (enabled/option/active): Yes/No/No
CR (enabled/option/active): Yes/No/No
Start time: Fri Apr 15 05:14:20 2016
End time: -----
Status: IN PROGRESS
=====
Test Start time Result EST (HH:MM:SS) Comments
=====
Electrical loopback 05:14:22 PASSED -----
Optical loopback 05:14:37 IN PROGRESS 00:01:00 -----
Link traffic test ----- NOT STARTED -----
=====
Roundtrip link latency: 148 nano-seconds
Estimated cable distance: 0 meters
Egress power: Tx: -2.9 dBm, Rx: Not Avail.
Ingress power: Rx: -2.8 dBm, Tx: Not Avail.

```

Enter **portcfgshow** to see which ports are D_Port-enabled.

```

device:admin> portcfgshow
Ports of Slot 0 24 26 27
-----+-----+-----
Octet Speed Combo 1 1 1

```

```

Speed                AN  AN  AN
AL_PA Offset 13      .. .. ..
Trunk Port           ON  ON  ON
Long Distance        .. .. ..
.....
.....
Port Auto Disable    .. .. ..
CSCTL mode           .. .. ..
D-Port mode          ON  ON  ON
D-Port over DWDM     .. .. ..
Compression          .. .. ..
Encryption           .. .. ..
FEC                  ON  ON  ON
Fault Delay          0   0   0
    where AE:QoSAutoEnable, AN:AutoNegotiate, ..:OFF, -:NotApplicable, ?:INVALID

```

Switch Type and Blade ID

The switchType is a field displayed in the **switchshow** command output. When you are gathering information to give to your switch support provider, you may be asked the switch model. If you do not know the model, see "SwitchType to B-Series Model Conversions" table to convert the switchType to a B-Series model number.

```
device:admin> switchshow
switchName:    4thfloor_A
switchType:    162.0 <=== convert this number using the following table.
switchState:   Online
switchMode:    Native
switchRole:    Principal
switchDomain:  1
switchId:      fffc00
switchWwn:     10:00:00:00:00:00:00:00
zoning:        OFF
switchBeacon:  OFF
FC Router:     OFF
FC Router BB Fabric ID: 128
```

In the example above, the number 162 is the switchType and .0 is the revision of the motherboard of the switch. The revision number is not necessary when converting the number. Convert the value using the following table.

TABLE 50 switchType to B-Series model conversions

switchType	B-Series switch model	Base switch speed
109	6510	16 Gb 48-port switch
117	6547	16 Gb 48-port Blade Server SAN I/O Module
118	6505	16 Gb 24-port switch
120	DCX 8510-8	16 Gb 512-port core fabric backbone
121	DCX 8510-4	16 Gb 256-port core fabric backbone
129	6548	16 Gb 28-port Blade Server SAN I/O Module
130	M6505	16 Gb 24-port Blade Server SAN I/O Module
133	6520	16 Gb 96-port switch
148	7840	16 Gb 24-FC ports, 16 10GbE ports, 2 40GbE ports extension switch
149	6546	16 Gb 24-port Blade Server SAN I/O Module
150	6545	16 Gb 26-port Blade Server SAN I/O Module
152	6549	16 Gb 28-port Blade Server SAN I/O Module
156	6543	16 Gb 24-port Blade Server SAN I/O Module
157	6558	16 Gb 48-port Blade Server SAN I/O Module
158	6559	16 Gb 48-port Blade Server SAN I/O Module
162	G620	32 Gb 48-port and four 4x32Gb QSFP-port switch
165	X6-4	32 Gb 192-port core fabric backbone
166	X6-8	32 Gb 384-port core fabric backbone
167	6542	16 Gb 32-port Blade Server SAN I/O Module
170	G610	32 Gb 24-port switch
178	7810	32 Gb 12-FC ports, 6 10GbE ports, 2 1GE copper ports extension switch

NOTE

For more information about the B-series switch models, refer to the relevant hardware installation guide.

See the following table to find the description of the blade model displayed in the output from the **slotshow** command.

```
device:admin> slotshow
Slot  Blade Type      ID      Model Name      Status
-----
  1      CP BLADE        175      CPX6            ENABLED
  2      CP BLADE        175      CPX6            ENABLED
  3      SW BLADE        178      FC32-48         ENABLED
  4      SW BLADE        178      FC32-48         ENABLED
  5      SW BLADE        178      FC32-48         ENABLED
  6      AP BLADE        186      SX6             ENABLED
  7      CORE BLADE      177      CR32-8          ENABLED
  8      CORE BLADE      177      CR32-8          ENABLED
  9      SW BLADE        178      FC32-48         ENABLED
 10      SW BLADE        178      FC32-48         ENABLED
 11      SW BLADE        178      FC32-48         ENABLED
 12      AP BLADE        186      SX6             ENABLED
```

TABLE 51 B-series blade model descriptions

Blade ID	B-series blade model	Description
50	CP8	Gen 5 DCX 8510 Director control processor blade
75	FX8-24	24-FC port with 10 1GbE and two 10GbE ports Fibre Channel routing and FCIP blade
96	FC16-48	16 Gb 48-FC ports blade
97	FC16-32	16 Gb 32-FC ports blade
98	CR16-8	4x16 Gb 16-FC ports core routing blade
99	CR16-4	4x16 Gb 8-FC ports core routing blade
153	FC16-64	16 Gb 64-FC ports blade
175	CPX6	Gen 6 X6 Director control processor blade
176	CR32-4	4x32 Gb 8-FC ports core routing blade
177	CR32-8	4x32 Gb 16-FC ports core routing blade
178	FC32-48	32 Gb 48-FC ports blade
186	SX6	16-FC port FCIP extension blade

Hexadecimal Conversion

- Hexadecimal overview.....126
- Example conversion of the hexadecimal triplet 0x616000.....126
- Decimal-to-hexadecimal conversion table.....127

Hexadecimal overview

Hexadecimal, also known as hex, is a numeral system with a base of 16, usually written by means of symbols 0-9 and A-F (or a-f). Its primary purpose is to represent the binary code that computers interpret in a format easier for humans to remember. It acts as a form of shorthand, in which one hexadecimal digit takes the place of four binary bits. For example, the decimal numeral 79, with the binary representation of 01001111, is 4F (or 4f) in hexadecimal, where 4 = 0100 and F = 1111.

Hexadecimal numbers can have either an *Ox* prefix or an *h* suffix. The address 0xFFFFFA is the same address as FFFFFAh. This type of address with 6 digits representing 3 bytes, is called a hex triplet. Fibre Channel uses hexadecimal notation in hex triplets to specify well-known addresses and port IDs.

Example conversion of the hexadecimal triplet 0x616000

Notice the PID (610600 - in bold) in the **nsShow** output is in hexadecimal.

```
switch:admin> nsshow
{
  Type Pid      COS      PortName      NodeName      TTL (sec)
  N      610600;  2,3;10:00:00:00:c9:29:b3:84;20:00:00:00:c9:29:b3:84; na
  FC4s: FCP
  NodeSymb: [36] "Emulex LP9002 FV3.90A7 DV5-5.10A10 "
  Fabric Port Name: 20:08:00:05:1e:01:23:e0
  Permanent Port Name: 10:00:00:00:c9:29:b3:84
  Port Index: 6
  Share Area: No
  Device Shared in Other AD: No
  Redirect: No
  LSAN: Yes
  The Local Name Server has 1 entry }
```

1. Separate the 6 digits into triplets by inserting a space after every 2 digits: 61 06 00.
2. Convert each hexadecimal value to a decimal representation:

61 = Domain ID = 97

06 = Area (port number) = 06

00 = Port (ALPA) = 0 (not used in this instance, but is used in loop, shared areas in PID assignments on blades, NPV, and Access Gateway devices)

Result: hexadecimal triplet 610600 = decimal triplet 97,06,00

Refer to when [Decimal-to-hexadecimal conversion table](#) on page 127 converting hexadecimal values to a decimal representation.

Decimal-to-hexadecimal conversion table

TABLE 52 Decimal-to-hexadecimal conversion table

Decimal	01	02	03	04	05	06	07	08	09	10
Hex	01	02	03	04	05	06	07	08	09	0a
Decimal	11	12	13	14	15	16	17	18	19	20
Hex	0b	0c	0d	0e	0f	10	11	12	13	14
Decimal	21	22	23	24	25	26	27	28	29	30
Hex	15	16	17	18	19	1a	1b	1c	1d	1e
Decimal	31	32	33	34	35	36	37	38	39	40
Hex	1f	20	21	22	23	24	25	26	27	28
Decimal	41	42	43	44	45	46	47	48	49	50
Hex	29	2a	2b	2c	2d	2e	2f	30	31	32
Decimal	51	52	53	54	55	56	57	58	59	60
Hex	33	34	35	36	37	38	39	3a	3b	3c
Decimal	61	62	63	64	65	66	67	68	69	70
Hex	3d	3e	3f	40	41	42	43	44	45	46
Decimal	71	72	73	74	75	76	77	78	79	80
Hex	47	48	49	4a	4b	4c	4d	4e	4f	50
Decimal	81	82	83	84	85	86	87	88	89	90
Hex	51	52	53	54	55	56	57	58	59	5a
Decimal	91	92	93	94	95	96	97	98	99	100
Hex	5b	5c	5d	5e	5f	60	61	62	63	64
Decimal	101	102	103	104	105	106	107	108	109	110
Hex	65	66	67	68	69	6a	6b	6c	6d	6e
Decimal	111	112	113	114	115	116	117	118	119	120
Hex	6f	70	71	72	73	74	75	76	77	78
Decimal	121	122	123	124	125	126	127	128	129	130
Hex	79	7a	7b	7c	7d	7e	7f	80	81	82
Decimal	131	132	133	134	135	136	137	138	139	140
Hex	83	84	85	86	87	88	89	8a	8b	8c
Decimal	141	142	143	144	145	146	147	148	149	150
Hex	8d	8e	8f	90	91	92	93	94	95	96
Decimal	151	152	153	154	155	156	157	158	159	160
Hex	97	98	99	9a	9b	9c	9d	9e	9f	a0
Decimal	161	162	163	164	165	166	167	168	169	170
Hex	a1	a2	a3	a4	a5	a6	a7	a8	a9	aa
Decimal	171	172	173	174	175	176	177	178	179	180
Hex	ab	ac	ad	ae	af	b0	b1	b2	b3	b4
Decimal	181	182	183	184	185	186	187	188	189	190
Hex	b5	b6	b7	b8	b9	ba	bb	bc	bd	be
Decimal	191	192	193	194	195	196	197	198	199	200
Hex	bf	c0	c1	c2	c3	c4	c5	c6	c7	c8

TABLE 52 Decimal-to-hexadecimal conversion table (continued)

Decimal	201	202	203	204	205	206	207	208	209	210
Hex	c9	ca	cb	cc	cd	ce	cf	d0	d1	d2
Decimal	211	212	213	214	215	216	217	218	219	220
Hex	d3	d4	d5	d6	d7	d8	d9	da	db	dc
Decimal	221	222	223	224	225	226	227	228	229	230
Hex	dd	de	df	e0	e1	e2	e3	e4	e5	e6
Decimal	231	232	233	234	235	236	237	238	239	240
Hex	e7	e8	e9	ea	eb	ec	ed	ef	ee	f0
Decimal	241	242	243	244	245	246	247	248	249	250
Hex	f1	f2	f3	f4	f5	f6	f7	f8	f9	fa
Decimal	251	252	253	254	255					
Hex	fb	fc	fd	fe	ff					

Revision History

FOS-821-TD-UG101; September 28, 2018

- Revised the publication number.
- Revised the "Rolling Reboot Detection" topic.
- Added "Revision History" to the document.
- Updated the document template.

FOS-821-TD-UG100; August 28, 2018

- Added Portloopbacktest mode 8 support for the Brocade 7810 Extension Switch.
- Updated the supported platforms for D_Ports with Fabric OS.
- Updated the "Preinstallation messages" section.
- Updated the switch type and the blade ID for the Brocade 7810 Extension Switch.
- Added a note on using the spinFab and portTest commands.