

BROCADE



EDUCATION SOLUTIONS

BCFA in a Nutshell 8 Gbps Study Guide for Exam 143-410

Exam Preparation Materials

Corporate Headquarters - San Jose, CA USA

T: (408) 333-8000
info@brocade.com

European Headquarters - Geneva, Switzerland

T: +41 22 799 56 40
emea-info@brocade.com

Asia Pacific Headquarters - Singapore

T: +65-6538-4700
apac-info@brocade.com

© 2010 Brocade Communications Systems, Inc. All Rights Reserved.

Brocade, the Brocade B-weave logo, Fabric OS, File Lifecycle Manager, MyView, Secure Fabric OS, SilkWorm, and StorageX are registered trademarks and the Brocade B-wing symbol and Tapestry are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. FICON is a registered trademark of IBM Corporation in the U.S. and other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

Revision: January, 2010

BCFA in a Nutshell 8 Gbps First Edition



Objective: The BCFA Nutshell guide is designed to help you prepare for the BCFA Certification, exam number 143-410.

Audience: The BCFA Nutshell self-study guide is intended for those who have successfully completed the CFA 280 Introduction to Fibre Channel Administration and Theory course, and who wish to undertake self-study or review activities before taking the actual BCFA exam. The BCFA guide is not intended as a substitute for classroom training or hands-on time with Brocade products.

How to make the most of the BCFA guide: The BCFA guide summarizes the key topics on the BCFA exam for you in an easy to use format. It is organized closely around the exam objectives. We suggest this guide be used in conjunction with our free online knowledge assessment test. To benefit from the BCFA guide, we strongly recommend you have successfully completed the CFA 280 Introduction to Fibre Channel Administration and Theory course.

We hope you find this useful in your journey towards BCFA Certification, and we welcome your feedback by sending an email to jcannata@brocade.com.

Helen Lautenschlager
Director of Education Solutions

Joe Cannata
Certification Manager

A handwritten signature in black ink, appearing to read "Helen Lautenschlager".

A handwritten signature in blue ink, appearing to read "Joe Cannata".

Table of Contents

1 - Fibre Channel Concepts	9
Fibre Channel Networking Model	9
Port Types	10
Classes of Service	11
Fibre Channel Network Addressing	12
Shared Area Addressing	13
Well-Known Addresses	14
Device Communication	15
Fabric Initialization Process	17
Buffer-to-Buffer Credits	18
2 - Product Hardware Features	19
8 Gbps Stand-alone Brocade Switches	19
Brocade 300	20
Brocade 5100	21
Brocade 5300	22
Brocade 7500/7500E	23
Brocade 7800	24
Brocade 8000	25
Brocade Encryption Switch (BES)	26
Chassis-based Switches	27
Brocade DCX Backbone	27
Brocade DCX-4S Backbone	28
Brocade 48000 Director	29
Available Blades and Compatibility	30
ICLs	32
ICL Overview	32
ICL Licensing	32
Supported ICL Configurations	33
Brocade FC HBAs	34
Brocade CNAs	35
3 - Installation and Configuration	36
Installation Concerns	36
Securing Management Access	36
Set the Management IP Address	37
Set the Command Line Session Timeout	37
Set the Login Banner	38
Activate License Features	38
Set the Switch Name	38
Set the Chassis Name	39
Set the syslog Server	39
Set Password Rules	39
Role Based Access Control	40
RADIUS and LDAP	41
Interoperability Modes	41
McDATA Fabric Mode	42
McDATA Open Fabric Mode	42
4 - FCP Routing	43
Fabric Terminology	43
Principal Switch Path	43
Principal Switch Commands	44
Routing Policies	44
Exchange-Based Routing	45
Exchange-Based Routing and DLS	46
4 and 8 Gbps Trunking Overview	47
The Deskew Counter	47
5 - Zoning	48
Maximum Zone Database Size	48

Zoning Best Practices	49
Default Zoning	49
Types of Zoning Enforcement	50
6 - Management	51
Management Interfaces and Tools	51
DCFM Overview and Features	51
Brocade SAN Health	53
Host Connectivity Manager (HCM)	54
Fabric Watch	54
SNMP Authentication	55
Configuration Files	56
Configuration Upload and Download	57
Firmware Download	57
Stand-alone Switches	57
Chassis-based Switches	57
USB Storage	58
Slow Drain Device Detection	58
7 - Troubleshooting	59
Problem Approach	59
Gather Information	60
Diagraming the Fabric	60
Common SAN Problems	61
Web Tools Switch / FRU Status	62
Blade Status LEDs	63
Switch / FRU Status Commands	64
HCM SupportSave	64
Fabric OS Support Data	65
Taking the Test	66

List of Tables

Classes of Service	11
Core Blade Compatibility Matrix	30
Port Blade Compatibility Matrix	30
Extension and Application Blade Compatibility Matrix	31
aaaconfig Options	41
Zoning Enforcement	50
Gathering Troubleshooting Information	60
FRU Status Commands	64

List of Figures

Fibre Channel Networking Model	9
Node and Port Types	10
24-bit Addressing	12
Shared Area Addressing	13
Device Communication Example	16
Fabric Initialization	17
8 Gbps Switches	19
Brocade 300	20
Brocade 5100	21
Brocade 5300	22
Brocade 7500/7500E	23
Brocade 7800	24
Brocade 8000	25
Brocade Encryption Switch	26
Brocade DCX Backbone	27
Brocade DCX-4S Backbone	28
Brocade 48000 Director	29
Supported ICL Configurations	33
Brocade FC HBAs	34
Brocade CNAs	35
Principal Switch Path	43
Exchange-based Routing	45
Exchange-Based Routing and DLS	46
Trunking	47
Maximum zoning database size	48
Brocade SAN Health	53
Configuration Files	56
Slow Drain Device Detection	58
Web Tools Switch / FRU Status	62
Blade Status LEDs	63
Sample NDA	66

1 - Fibre Channel Concepts

Fibre Channel Networking Model

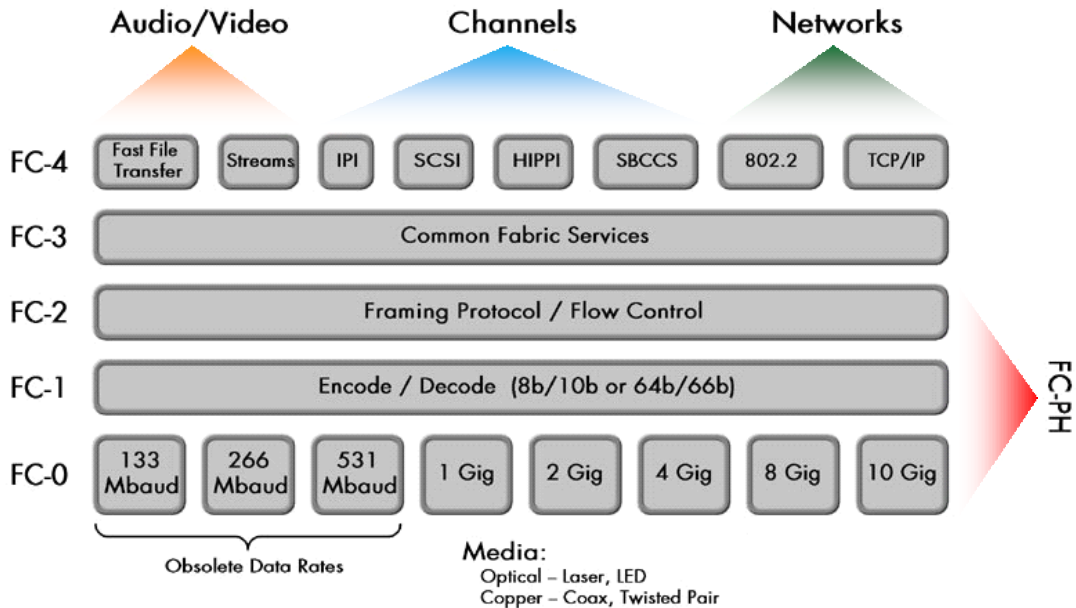


Figure 1: Fibre Channel Networking Model

- The FC-0 and FC-1 layers specify physical and data link functions needed to physically send data from one port to another
- FC-0 specifications include information about feeds and speeds
- FC-1 layer contains specifications for 1, 2, 4 and 8 Gbps 8b/10b encoding, ordered set and link control communication functions. 10 Gbps communication uses 64b/66b encoding
- FC-2 specifies content and structure of information along with how to control and manage information delivery. This layer contains basic rules needed for sending data across the network. This includes: (1) how to divide the data into frames, (2) how much data should be sent at one time before sending more (flow control), and (3) where the frame should go. It also includes Classes of Services, which define different implementations that can be selected depending on the application.
- FC-3 defines advanced features such as striping (to transmit one data unit across multiple links) and multicast (to transmit a single transmission to multiple destinations) and hunt group (mapping multiple ports to a single node). While the FC-2 level concerns itself with the definition of functions with a single port, the FC-3 level deals with functions that span multiple ports.
- FC-4 provides mapping of Fibre Channel capabilities to pre-existing protocols, such as IP, SCSI, or ATM, etc.

Port Types

Device Ports (Nx_Ports)

- N_Port – Node Port, a Fabric device directly attached
- NL_Port – Node Loop Port, a device attached to a loop

Switch Ports

- U_Port – Universal Port, a port waiting to become another port type
- FL_Port – Fabric Loop Port, a port to which a loop attaches
- G_Port – Generic Port, a port waiting to be an F_Port or E_Port
- F_Port – Fabric Port, a port to which an N_Port attaches
- E_Port – Expansion Port, a port used for inter-switch links (ISLs)

Configured Ports

- EX_Port – A type of E_Port used to connect to an FC Router fabric1
- VE_Port – Virtual E_Port (used in FCIP fabrics)
- VEX_Port – VEX_Ports are no different from EX_Ports, except underlying transport is IP rather than FC

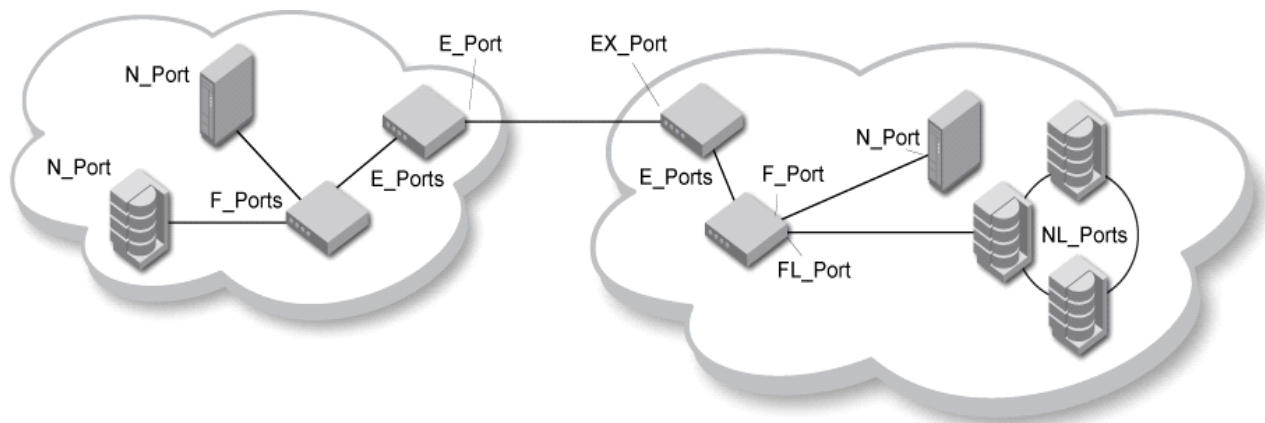


Figure 2: Node and Port Types

Classes of Service

Class	Description	ACK	Brocade Supported
1	Dedicated connection, allocates full bandwidth between ports	Yes	No
2	Connectionless, switch-to-switch communication between ports, transfers frames	Yes	Yes
3	Connectionless, switch-to-switch communication between ports, transfers frames	No	Yes
4	Dedicated connection, allocates requested amount of bandwidth between ports, uses virtual circuits	Yes	No
6	Dedicated connection for multicast service	Yes	No
F	Switch to switch communication	Yes	Yes

Table 1: Classes of Service

Different types of information have different delivery requirements. In order to accommodate the different delivery requirements, Fibre Channel associates a set of delivery characteristics into what is called a Class of Service. The characteristics relate to the type of connection between the ports, confirmation of delivery, flow control mechanisms and how errors are handled.

- Class-1 is a connection-oriented circuit that dedicates 100% of the bandwidth between the sending and receiving ports. It also provides for a confirmation of delivery (ACK).
- Class-2 is a connectionless class with an acknowledgement (confirmation of delivery). No bandwidth is allocated or guaranteed. IP uses this class. Uses both Buffer-to-Buffer (BB) credits and End-to-End (EE) credits for flow control.
- Class-3 is a connectionless class without an acknowledgement (confirmation of delivery). No bandwidth is allocated or guaranteed. FCP uses this class. Uses Buffer-to-Buffer (BB) credits for flow control, does not use End-to-End (EE) credits.
- Class-4 is a connection-oriented class that uses virtual circuits and confirmation of delivery. Unlike Class-1 that reserves the entire bandwidth, Class-4 can allocate a requested amount of bandwidth.
- Class-6 is a variation of Class-1 that provides a one-to-many multicast service with a confirmation of delivery
- Class-F is a connectionless class with acknowledgements (confirmation of delivery). between two switches

Note: Brocade supports Class-2, Class-3, and Class-F only.

Fibre Channel Network Addressing

When a node attaches to the fabric, it must receive a unique 24-bit address. The network address is a three-byte address based upon the Domain ID, the Area ID and, if a loop device, its AL_PA. This address is the source address and is used for routing data thru the fabric from one device to another.

Fabric-attached devices use an address format of “DD AA 00”. This is the address of any Fabric-attached device that has logged into the fabric as point-to-point.

Public Loop attached devices use an address format of “DD AA PP”. The “DD AA” bytes of the address come from the fabric login process and the “PP” byte is assigned during FC_AL initialization.

NPIV attached devices use an address format of “DD AA PP”. The “DD AA” bytes of the address come from the fabric login process and the “PP” byte is assigned during login process.

Each switch is responsible for assigning unique addresses

Addresses are 24 bits:

- Domain ID (8 bits) 0x01 - 0xEF
- Area ID (8 bits) 0x00 - 0xFF
- Node Address (8 bits) 00 / AL_PA / NPIV / Shared Area

Address types:

- Fabric DD AA XX2
- Public loop / NPIV DD AA PP
- Shared Area DD AA 80




Figure 3: 24-bit Addressing

Shared Area Addressing

- A shared area is an Area ID that exists more than once in a single domain
 - These shared areas are differentiated by their Node Addresses
 - This allows for more than 256 ports in a single domain
- The FC8-48 has some ports that use shared areas
 - Ports 16-47 of the FC8-48 blade use shared areas
- Shared Area PIDs use a Node Address of either 0x00 or 0x80
 - Example of two shared areas on a FC8-48 blade in slot:
 - 018000 – Port 16
 - 018080 – Port 40
- Shared Area IDs will use the Node Address to allow 384 ports to be addressed in a single domain.
- The FC8-48 blade does not use shared areas when installed into a DCX-4S since the total port count in the domain would not exceed 256.

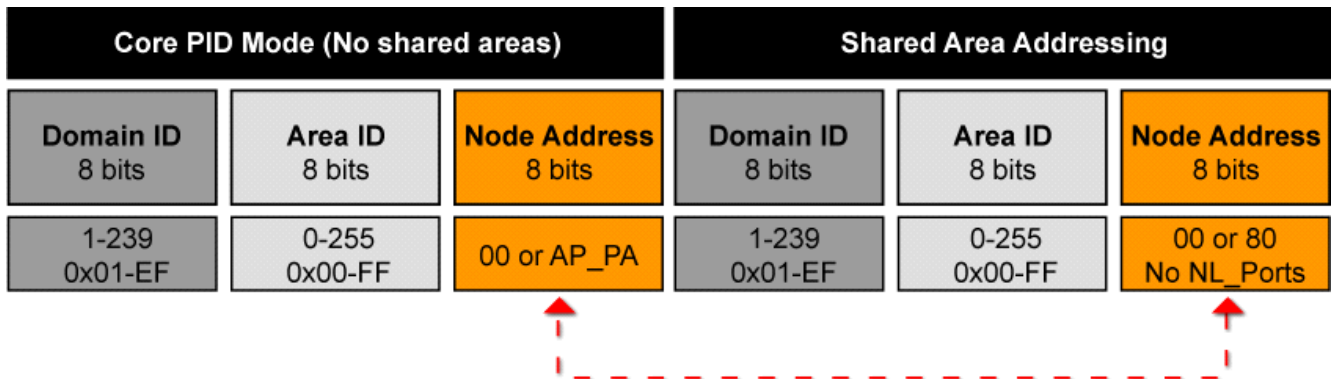


Figure 4: Shared Area Addressing

Well-Known Addresses

Every switch has reserved 24-bit addresses known as 'Well Known Addresses'. The services residing at these addresses provide a service to either nodes or management applications in the fabric.

FFFFF6 Clock Synchronization Server: Clock Synchronization over Fibre Channel is attained through a Clock Synchronization Server that contains a reference clock. The Server synchronizes client's clocks to the reference clock on a periodic basis, using either Primitive Signals or ELS frames.

FFFFF7 Security Server: The security-key distribution service offers a mechanism for the secure distribution of secret encryption keys.

FFFFF8 Alias Server: The Alias Server manages the registration and deregistration of Alias IDs for both hunt groups and multicast groups. The Alias Server is not involved in the routing of frames for any group.

FFFFFA Management Server: The Management server provides a single point for managing the fabric.

FFFFFB Time Server: The time server sends to the member switches in the fabric the time on either the principal switch or the Primary FCS switch.

FFFFFC Directory (Name Server): The directory server/name server is where fabric/public nodes register themselves and query to discover other devices in the fabric.

FFFFFD Fabric Controller: The fabric controller provides state change notifications to registered nodes when a change in the fabric topology occurs.

FFFFFE F_Port (Fabric Server Login): Before a fabric node can communicate with services on the switch or other nodes in the fabric, an address is assigned by the fabric login server. Fabric addresses assigned to nodes are 3 bytes long and are a combination of the domain ID plus the port area number of the port the node is attached to.

FFFFFF Broadcast Server: When a frame is transmitted to this address, the frame is broadcast to all operational N and NL ports.

Device Communication

Below is an example of the frame communication between a host device and the switch (fabric).

Note: The HBA, among other things, is responsible for framing packets, physical addressing and link level error checking.

FLOGI: Fabric Login command: Used to establish a 24-bit address for the device logging in. Also establishes Buffer-to-Buffer credits, class of service supported.

PLOGI: Port Login command: Device must login into the Directory (Name) Server to Register its information as well as query for devices this device is zoned with.

SCR: State Change Registration: Device needs to register for State Change Notification so if there is a change in the fabric, such as a zoning change or a change in the state of a device that this device has access to, the device will receive an RSCN.

Registration: A device will exchange registration information with the Directory (Name) Server.

Query: Devices can query the Directory (Name) Server for information about the device it has access to.

PLOGI: Port Login command: Initiator must login into the target.

PRLI: Process Login command: This establishes the operating (SCSI is the most common) environment between the two N_Ports.

Inquiry: This example uses an INQ command, it could be something different such as a report LUNs command for example. Which command is used is determined by the driver on the initiator. The command's purpose is to get a list of LUNs the initiator has access to.

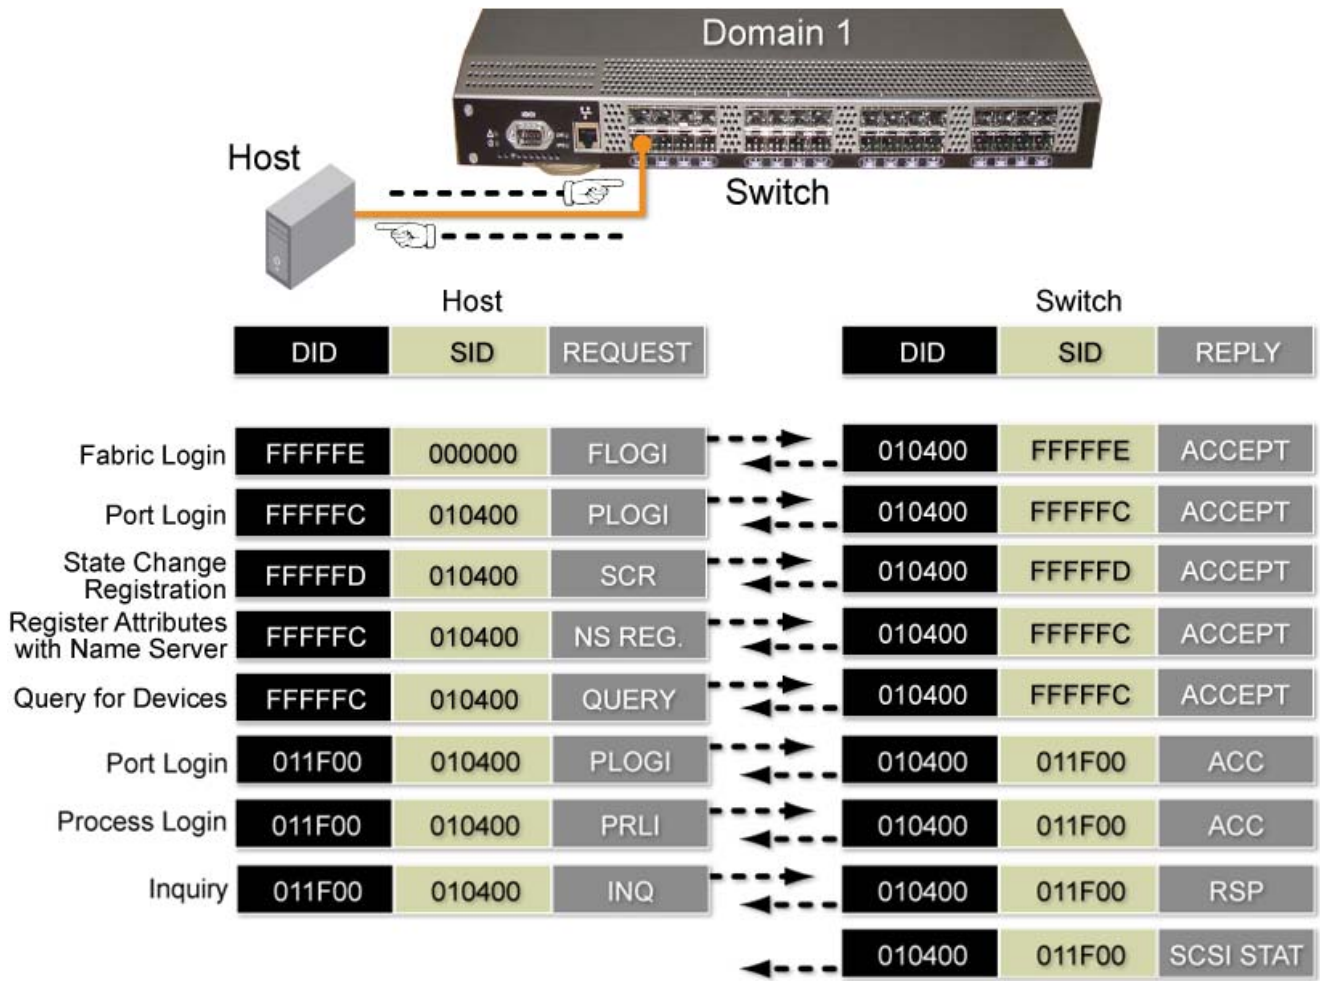


Figure 5: Device Communication Example

Fabric Initialization Process

When a switch joins a fabric several Class F frames are used to exchange various parameters:

- ELP (Exchange Link Parameters)
 - Contains sender information
 - R_A_TOV / E_D_TOV
 - PWWN / Switch Name
 - Flow control used
- ESC (Exchange Switch Capabilities)
 - Vendor specific info
 - Virtual Fabric support
- EFP (Exchange Fabric Parameters)
 - Principal switch selection
 - Principal switch priority
 - Switch name
 - Domain ID list
 - Vendor specific info
 - Virtual Fabric support
 - Zoning database

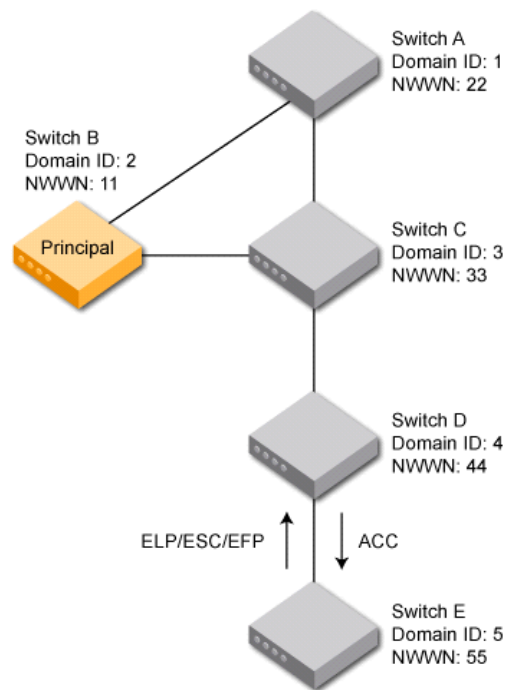


Figure 6: Fabric Initialization

Buffer-to-Buffer Credits

- BB Credits are used as a flow control mechanism to ensure that the transmitter does not overwhelm the receiver with too many frames
- BB Credits are exchanged during login:
 - FLOGI – Accept for Nx_Ports
 - ELP – Accept for E_Ports
- Both sides (example device and switch) do not have to have the same amount of credits
- Distance, link speed, and frame size affect required credits to fill the link

2 - Product Hardware Features

8 Gbps Stand-alone Brocade Switches




Figure 7: 8 Gbps Switches

Brocade 300



Figure 8: Brocade 300

- 1U form factor
- Single FRU
- USB port
- 24 FC Ports
- Ports on Demand (8-port increments)
- 8 Gbps FC port speed
- 1 GoldenEye2 ASIC
- 676 user BB credits per ASIC
- 1:1 subscription ratio
- 3 x 8-port trunk groups
- Integrated Routing (IR) is NOT supported
- Access Gateway capable
- Comes with 8 licensed ports; subsequent ports are licensed in groups of 8 to support 16 and 24 ports.
- Supports all speeds up to and including 8 Gbps.
- Access Gateway capable, but must have all ports activated with POD licenses.

Brocade 5100



Figure 9: Brocade 5100

- 1U form factor
- FRUs: Two 125 W Power Supply/Fan Assemblies
- USB port
- 40 FC Ports
- Ports on Demand (8-port increments)
- 8 Gbps FC port speed
- 1 Condor2 ASIC
- 2012 user BB credits per ASIC
- 1:1 subscription ratio
- 5 x 8-port trunk groups
- Support for Integrated Routing (IR) – Available per port
- Access Gateway capable
- Comes with 24 licensed ports; subsequent ports are licensed in groups of 8 to support 32 and 40 ports
- Supports all speeds up to and including 8 Gbps
- Integrated Routing support requires an Integrated Routing license

Brocade 5300



Figure 10: Brocade 5300

- 2U form factor
- FRUs: Two 125 W Power Supplies and three fans
- USB port
- 80 FC Ports
- Ports on Demand (16-port increments)
- 8 Gbps FC port speed
- 9 GoldenEye2 ASICs
- 292 user BB credits per ASIC
- 1:1 subscription ratio
- 10 x 8-port trunk groups
- Support for Integrated Routing (IR) – Available per port
- No Access Gateway support
- Comes with 48 licensed ports; subsequent ports are licensed in groups of 16 to support 64 and 80 ports
- Supports all speeds up to and including 8 Gbps

Brocade 7500/7500E




Figure 11: Brocade 7500/7500E

- The Brocade 7500 is a 1U switch for FC-FC routing and FCIP
 - Sixteen 4 Gbps FC F, FL, E or EX_Ports
 - Two 1 GbE ports supporting FCIP connections with multiple tunnels, data compression, traffic shaping, and routing
 - Port speeds of 4, 2, or 1 Gbps
 - ISL Trunking, Extended Fabrics (FC ports only)
 - Two 8-port groups (ports 0-7 and ports 8-15)
 - 377 available user buffer credits per ASIC
 - Available as a bladed product, the FR4-18i
 - FCIP functionality not compatible with the Brocade 7800 switch or FX8-24 blade
- The Brocade 7500E is an expanded version of the 7500
 - Same chassis as 7500
 - Economical solution for connecting remote sites using FCIP
 - Four ports:
 - Two 4 Gbps Fibre Channel (E, F, FL, EX) ports
 - Two 1 GbE ports (VE, VEX) ports (up to 50 Mbps each)
 - Software license upgrade available

Brocade 7800




Figure 12: Brocade 7800

- The Brocade 7800 has two configurations
- Brocade 7800 4/2 – four 8 Gbps FC Ports and two 1 GbE Ports
- Brocade 7800 16/6 – sixteen 8 Gbps FC Ports and six 1 GbE Ports
- Auto-sensing link speeds up to 8 Gbps
- One Condor2 ASIC
- Available as a bladed product, the FX8-24

Brocade 8000




Figure 13: Brocade 8000

- Essentially two switches in one: a FC switch and a CEE switch
- 8 Gbps FC ports link to SAN devices and 10 GbE (CEE) ports connect to Ethernet LAN devices
- Capable of forwarding frames between CEE/FCoE and FC
- Available as a bladed product, the FCOE10-24

Brocade Encryption Switch (BES)




Figure 14: Brocade Encryption Switch

- 32 FC Ports
- 8, 4, 2, and 1 Gbps FC port speed
- Two 1 GbE ports for out-of-band cluster HA interconnect
- Two 300W power supplies
- USB port
- Three fans
- 2U form factor
- Two Condor2 ASICs
- Available as a bladed product, the FS8-18
- Smart Card reader used for:
 - System Card – Can be required for starting encryption services
 - Authentication Cards – A quorum of cards can be required to make changes to encryption configurations
- Requires Fabric OS v6.3.0 or higher

Chassis-based Switches

Brocade DCX Backbone



Figure 15: Brocade DCX Backbone

- 12 keyed slots (blades cannot be installed in the wrong slot)
- 8 port blades (1-4, 9-12)
- 2 CP8 Control Processor blades (6, 7)
- 2 CR8 Core Routing blades (5, 8)
- Up to 384 user ports
- Dual WWN boards for redundancy
- Four 2000 watt power supplies provide power and redundancy
- Three 200 mm blowers for chassis cooling
- Air intakes are located on the blower FRU (Field Replaceable Unit) and underneath the power supplies
- Air flow in the DCX is from the non-port side to the port-side of the chassis

Brocade DCX-4S Backbone



Figure 16: Brocade DCX-4S Backbone

- Provides a platform for midsize and large enterprise
- Smaller chassis and lower cost than the DCX
- 8 keyed slots (blades cannot be installed in the wrong slot)
- 4 port blades (1-2, 7-8)
- 2 CP8 blades (4, 5)
- 2 CR4S-8 blades (3, 6)
- Up to 192 user ports
- Dual WWN boards for redundancy (located behind the logo plate)
- Two 2000 watt power supplies for N+1
- Two 200 mm blowers for chassis cooling
- Air intakes on blower FRUs and right side of the chassis

Brocade 48000 Director



Figure 17: Brocade 48000 Director

- 10 blade slots
- 8 port blades (1-4, 7-10)
- 2 CP4 blades (5, 6)
- Processing and core switching is handled on the same blades unlike the DCX and DCX-4S
- Up to 384 user accessible Fibre Channel ports
- Four 1000 watt power supplies for N+1 redundancy located behind right side of front bezel
- Single WWN card contains two redundant partitions
- 3 blowers for chassis cooling
- Airflow in the 48000 chassis is from non-port side to port-side

Available Blades and Compatibility

- Blades available for the chassis are divided into several categories: CP, Core, Port, Application, and Extension blades
- Control Processor (CP) blades provide processing functionality
- Core Routing blades provide switching between the port blades within the same chassis as well as linking to other DCX and DCX-4S chassis
- Port blades provide Fibre Channel ports for connecting devices
- Application blades are designed for running fabric applications within the switch
- Extension and bridging blades provide services for extending a fabric such as FCoE and FCIP

	DCX	DCX-4S	48000
Core Processing (CP8)			
Core Routing (CR8)			
Core Routing (CR4S-8)			
Core Processing/Routing (CP4)			

Table 2: Core Blade Compatibility Matrix

	DCX	DCX-4S	48000
8 Gbps 16-port blade (FC8-16)			
8 Gbps 32-port blade (FC8-32)			
8 Gbps 48-port blade (FC8-48)			
4 Gbps 16-port blade (FC4-16)			
4 Gbps 32-port blade (FC4-32)			
4 Gbps 48-port blade (FC4-48)			
10 Gbps 6-port blade (FC10-6)			

Table 3: Port Blade Compatibility Matrix

	DCX	DCX-4S	48000
FCIP extension blade (FX8-24)	✓	✓	
FCoE extension blade (FCOE10-24)	✓	✓	
Application blade (FA4-18)	✓	✓	✓
Router/FCIP blade (FR4-18i)	✓	✓	✓
Storage encryption blade (FS8-18)	✓	✓	
iSCSI bridge blade (FC4-16IP)			✓

Table 4: Extension and Application Blade Compatibility Matrix

ICLs

ICL Overview

- ICLs provide dedicated connections between DCX and DCX-4S chassis
- ICL connectors are located on the CR8 and CR4S-8 core blades
- Allows up to three chassis to be connected
- Requires Fabric OS v6.3, earlier versions of firmware only support two chassis
- For FICON purposes, the ICL connection is not considered a hop
- Not supported on the 48000
- ICLs are 8 Gbps ISL connections between two or three DCX and DCX-4S chassis
- Speed locked at 8 Gbps
- Copper-based proprietary connector
- No SFPs
- Each port provides up to 16 x 8 (128) Gbps uni-directional bandwidth on the DCX and up to 8 x 8 (64) Gbps uni-directional bandwidth on the DCX-4S
- ICL cables are 2 meters in length
- ICL license required for both DCX and DCX-4S
- ICL 8-link license only activates half the bandwidth of the DCX, and is valid for either a DCX or DCX-4S
- ICL 16-link license activates all the bandwidth, valid for DCX only
- Allows ISL connectivity without consuming user ports

ICL Licensing

- ICL 16-Link License:
 - Activates all 16 links per ICL port on the DCX
 - Available on the DCX only
- ICL 8-Link License:
 - Activates all eight links on ICL ports on a DCX-4S chassis or half of the ICL bandwidth on the DCX
 - Allows users to purchase half the bandwidth of DCX ICL ports initially and upgrade with an additional 8-Link license to utilize the full ICL bandwidth at a later time
 - Useful for environments that wish to create ICL connections between a DCX and a DCX-4S
 - Available on the DCX-4S and DCX

Supported ICL Configurations

- Fabric OS v6.0 – v6.2
 - ICLs can run between two chassis only
- Fabric OS v6.3
 - Up to three chassis in full-mesh





Figure 18: Supported ICL Configurations


Brocade FC HBAs



BR-415
4 Gbps x 1 Port



BR-425
4 Gbps x 2 Ports



BR-815
8 Gbps x 1 Port



BR-825
8 Gbps x 2 Ports

Figure 19: Brocade FC HBAs

- Platform Support
 - Intel IA32, IEM64T, and IA64
 - AMD64, and x86
 - SUN SPARC, and x86
- Operating Systems
 - Windows 2003 and 2008
 - Redhat and Suse Linux - Version RHEL 4/5, SLES 9/10
 - Solaris - Version 9/10 (x86 and SPARC)
 - VMware - Version 3.5
- Support for Multi-Path I/O
 - Windows MPIO driver from Microsoft
 - Linux Device Mapper driver included in the kernel
- Both Full Height and Low Profile brackets supported and shipped with each HBA

Brocade CNAs



Figure 20: Brocade CNAs

- BR-1010: One 10 GbE port
- BR-1020: Two 10 GbE ports
- Both full height and low profile brackets supported and shipped with each CNA
- FCoE traffic support requires direct connect to FCoE Forwarder (FCF)

3 - Installation and Configuration

Installation Concerns

- Power
 - Cable dual power switches to dual power circuits
- Air
 - Most B-Series switches: air flow is from non-port side to port side
 - M-EOS switches: air flow is from port side to non-port side
- Cables
 - Allow for manageable cable slack to minimize stress
 - Do not mix single (longwave) with multimode (shortwave) in patch panel
 - Secure with Velcro straps
 - Be wary of distances - total can add up quickly with patch panels
 - Create a port map
- Monitor switch environment
 - Displays power status: `psshow`
 - Displays fan status: `fanshow`
 - Displays temp readings: `tempshow`
 - Displays all sensor readings: `sensorshow`

Securing Management Access

- Use the IPFilter policy to disable telnet or http (Web Tools)
- IPFilter policies are not automatically distributed to other switches in the fabric
 - To ensure that all switches are using the same policy use the `distribute` command
- If telnet is disabled from within a telnet session, all active telnet sessions will be terminated
- To avoid losing your session, disable telnet through an alternate interface
 - Serial port session
 - SSHv2 session
 - Web Tools

Set the Management IP Address

- Use the `ipaddrset` command to configure the switch IP address
- Default IP Address for Fabric OS switches 10.77.77.77
- Default netmask Fabric OS switches: 255.255.255.0
- Obtain addressing information for your network
 - IP Address & netmask
 - Default gateway
- Backbones/Directors require more than one IP address on the same subnet
 - One IP Address required for switch management
 - This IP address will be assigned dynamically to the active CP
 - One IP Address required per Control Processor
 - Default IP Addresses for Directors: 10.77.77.77 (switch management),
.75 (cp0), .74 (cp1)

Set the Command Line Session Timeout

- Automatically terminate a telnet or SSH session after a period of inactivity
 - Timeout value is specified in minutes
 - Setting a timeout value of 0 disables automatic session timeout
 - Valid settings include 0, or a value between 1 and 99,999 minutes
 - To display the current setting, type `timeout` with no arguments
- Default timeout on switches is 10 minutes. When changing the timeout value you can use the `login` command to restart the login session and use the new timeout value.

```
B51:admin> timeout
```

```
Current IDLE Timeout is 0 minutes
```

```
B51:admin> timeout 15
```

```
IDLE Timeout Changed to 15 minutes
```

```
The modified IDLE Timeout will be in effect after NEXT login
```

```
B51:admin> login
```

```
B51 login: admin
```

```
Password:
```

```
B51:admin> timeout
```

```
Current IDLE Timeout is 15 minutes
```

Set the Login Banner

- A login banner is displayed after a successful CLI or Web Tools login
- Viewed from command line or Web Tools
- Set using the `bannerset` command
- Remove using `bannerset ""`

Activate License Features

- Used to enable Fabric OS features
- Based on the switch License ID
 - `licenseidshow`
- License string is up to 32 case-sensitive characters
- A single license key may activate one feature or a bundle of features
- License commands
 - `licenseshow`
 - `licenseadd`
 - `licenseremove`
 - `licenseslotcfg`
 - Only used for Directors to install specific slot based licenses

Set the Switch Name

- Switch names should be unique for easier administration
- Naming suggestions
 - Site or building where switch is located
 - Floor or room where switch is located
 - Indicate topology (core switch vs. edge switch)
 - Rack ID
 - Switch Type
 - Fabric ID
 - Domain ID
- Example: SJC2_C4_E_B5100
 - San Jose, building 2
 - Room C4
 - Rack E
- Switch name is assigned using the `switchname` command

Set the Chassis Name

- The chassis name for the switch should be set along with the switch name
- With Fabric OS v6.3 and later the chassis name is used for naming `supportsave` files
- The chassis name can be set using the `chassisname` command
 - `chassisname SJC12_C4_E_B5100`
- Used to distinguish the physical chassis from logical switches in Virtual Fabric mode

Set the syslog Server

- The system logging daemon (`syslogd`) on hosts can receive system events and error messages from Brocade switches
- If all switches and control processors escalate messages to `syslogd`, the administrator may view a fabric-wide log of events
- Configuration is simple
 - `syslogdipadd`
 - `syslogdipremove`
 - `syslogdipshow`
- `syslog` records are tagged as belonging to a facility
 - Supports the `local1` through `local7` facilities
 - The default facility is `local7`
 - Change the facility, on a per-switch basis, using the `syslogdfacility` command
- Additional host configuration may be necessary, see server documentation

Set Password Rules

- Avoid stale passwords by setting a password expiration policy¹
 - Minimum age: `minpasswordage`
 - Maximum age: `maxpasswordage`
 - Expiration warning (days): `warning`
- Set the account lockout policy
 - Password failures allowed: `lockoutthreshold`
 - Set lockout duration (minutes): `lockoutduration`

Role Based Access Control

- Fabric OS implements two classes of accounts:
- Default Accounts (root, factory, admin, user)
 - Each default account has a hard-coded set of permissions
 - The permissions define roles with privileges corresponding to the account name
 - The privileges and account names cannot be changed
 - The accounts can be disabled if necessary
- User-Defined Accounts
 - 256 user-defined accounts available per switch
 - 32 simultaneous login sessions per switch (includes the default accounts)
- Use the `userconfig` command to administer accounts
 - `userconfig --show`
 - `userconfig --change`
 - `userconfig --add`
 - `userconfig --delete`
- User-defined accounts assist in tracking who did what, when
 - Enable enhanced change tracking with `trackchangeset 1`

RADIUS and LDAP

- When configured to use RADIUS or LDAP user accounts are managed from the authentication servers
- The local switch database is bypassed
- Account changes made to the local database are not propagated back to the authentication servers
- If authentication is configured to use the local switch database as a backup the local accounts will need to be created and maintained separately
- Configuration is done using the `aaaconfig --authspec` command

Arguments	Use
"local"	Local database only
"radius"	RADIUS only
"radius;local"	RADIUS first. If RADIUS authentication fails, use local database.
"radius;local" -backup	RADIUS first. Only use local database if RADIUS is unavailable.
"ldap"	LDAP only
"ldap;local"	LDAP first. If LDAP authentication fails, use local database.
"ldap;local" -backup	LDAP first. Only use local database if LDAP is unavailable.

Table 5: `aaaconfig` Options

Interoperability Modes

- Fabric OS v6.0 and later supports three fabric interop settings:
 - Interopmode 0 (Brocade Native Mode)
 - Interopmode 1 (no longer available)
 - Interopmode 2 (McDATA Fabric Mode)
 - Interopmode 3 (McDATA Open Fabric Mode)

McDATA Fabric Mode

- Fabric OS McDATA Fabric Mode (interopmode 2):
- Used for adding Fabric OS products to M-EOS fabrics in McDATA Fabric Mode
- Only supports connections to M-Series products – no other vendors
- Supports zoning from both B-series and M-Series switches
- When interopmode 2 is first enabled all existing zone configurations (defined and effective) are erased
- D,I and Port WWN zoning is supported
- Non-disruptive firmware upgrade
- Uses domain IDs 1-31 by default
- SCC policies are supported however DCC Policies cannot be used
- Requires M-EOS 9.07.02+
- Required interop mode for FICON

McDATA Open Fabric Mode

- Fabric OS McDATA Open Fabric Mode (interopmode 3):
- Used to add Fabric OS products to M-EOS fabrics in Open Fabric Mode
- Only supports connections to M-Series products – no other vendors
- When interopmode 3 is first enabled all existing zone configurations (defined and effective) are erased
- M-EOS 9.07.02 minimum required
- Uses domain IDs 97-127 by default
- Zoning can only be done via M-EOS switches
- Only Port WWN zoning is supported
- Fabric OS features that are not supported include but are not limited to: Traffic Isolation zones, Ingress Rate Limiting, Quality of Service, Advanced Performance Monitoring, and Top Talkers

4 - FCP Routing

Fabric Terminology

- **Principal Switch**
 - Selected when the fabric initializes, before routing is established
 - Manages the assignment of unique Domain IDs
 - Provides time synchronization to all other switches in the fabric
- **Principal ISL**
 - ISL used to communicate between the Principal Switch and other switches in the fabric

Principal Switch Path

FSPF uses several frames to perform its functions. Since it may run before fabric routing is set up, FSPF does not use the routing tables to propagate the frames, but floods the frames throughout the fabric hop-by-hop. At the beginning, frames are flooded on all the Inter-Switch Links (ISLs); as the protocol progresses, it builds a spanning tree rooted on the Principal Switch. Frames are then sent only on the ISLs that belong to the spanning tree. These ISLs are called “Principal ISLs”.

Where there are multiple ISLs between switches, the first ISL to respond to connection requests becomes the Principal ISL. Only one ISL from each switch will be used as the Principal ISL.

‘Upstream’ means going out that E_Port is going toward the Principal Switch. ‘Downstream’ means going out that E_Port is going away from the Principal Switch. These designations are seen in the switchshow output.

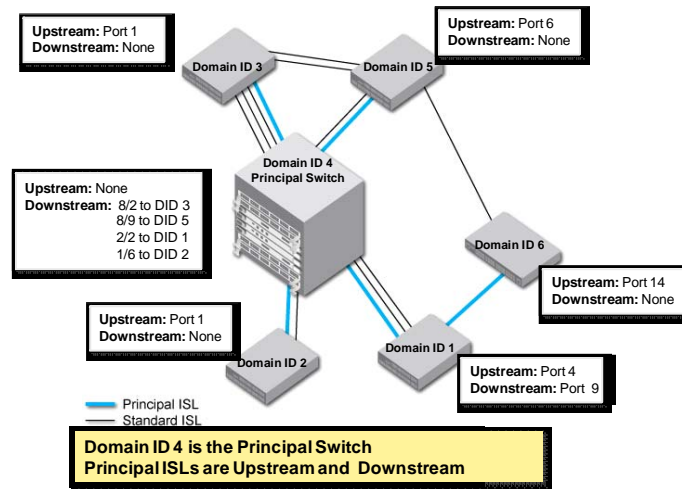


Figure 21: Principal Switch Path

Principal Switch Commands

- Set the preferred Principal Switch priority using:
`fabricprincipal -enable -p 0xff`
 - p – sets the principal selection priority for the switch
 - f – forces a fabric rebuild immediately after enabling on a switch
- Selection process:
 - If none of the switches has a priority setting, switch with the lowest WWN becomes Principal
 - Only switches with a priority setting participate in the selection
 - Switch with the highest priority (0xff is highest) becomes Principal
 - If more than one switch has highest priority, the switch with the lowest WWN becomes Principal

Routing Policies

- The routing policy is unidirectional and responsible for selecting a route based on one of two user-configurable routing policies:
 - Port-based routing
 - Exchange-based routing
- Each switch has its own routing policy
 - Different policies can exist in the same fabric
- 2 Gbps ASIC routing is handled by the Fabric Shortest Path First (FSPF) protocol and uses only Port-based routing
- 4/8 Gbps ASICs use the FSPF protocol and either Port-based routing or Exchange-based routing
 - Exchange-based routing is Brocade's factory default setting

Exchange-Based Routing

Exchange-based routing is also known as Dynamic Path Selection (DPS).

The choice of routing path is based on the Source ID (SID), Destination ID (DID), and Fibre Channel originator exchange ID (OXID), optimizing path utilization for the best performance. Thus, every exchange can take a different route through the fabric. Exchange-based routing requires the use of the Dynamic Load Sharing (DLS) feature. When this policy is in effect, you cannot make any changes to DLS.

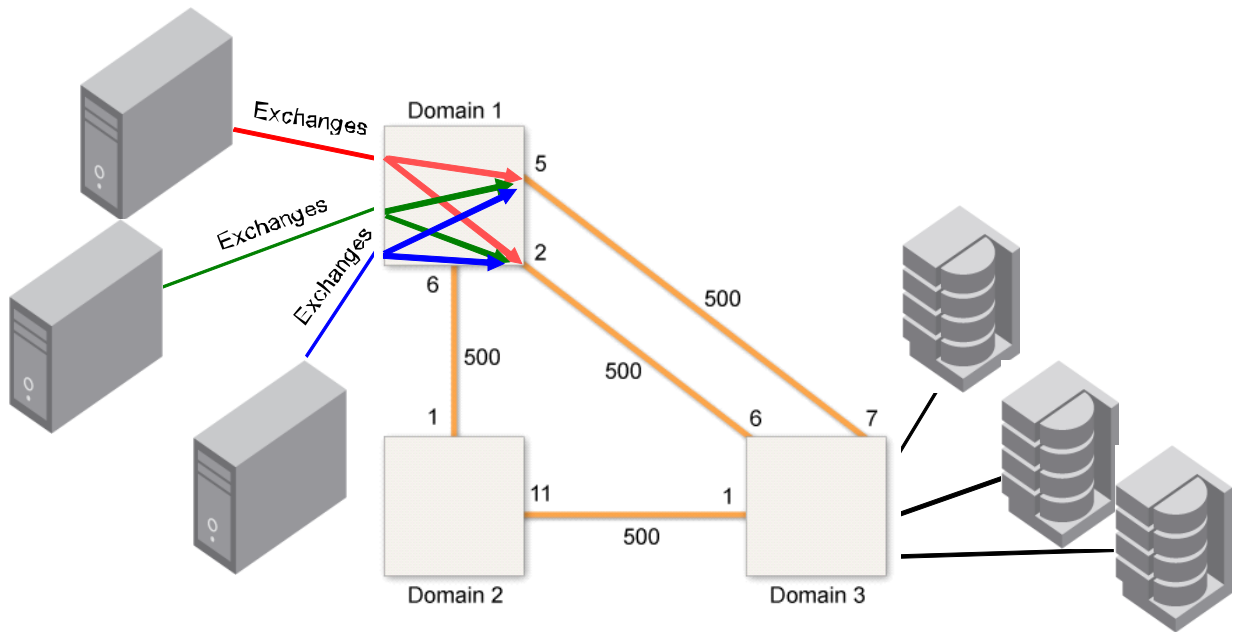


Figure 22: Exchange-based Routing

Exchange-Based Routing and DLS

The Link State Database calculates the cost of each link and determines the lowest cost path within each switch. The input port from the source is assigned to all output ports toward the destination (Dynamic Route Selection).

- Exchanges are allocated via round-robin assignment
- Chosen routes are used regardless of whether or not other devices in the fabric go offline or fabric changes occur
- Changes in the fabric, when Dynamic Load Sharing is enabled (DLSset), causes FSPF to recalculate the Dynamic distribution of exchanges to the remaining output ports to continue to distribute devices across equal cost routes.

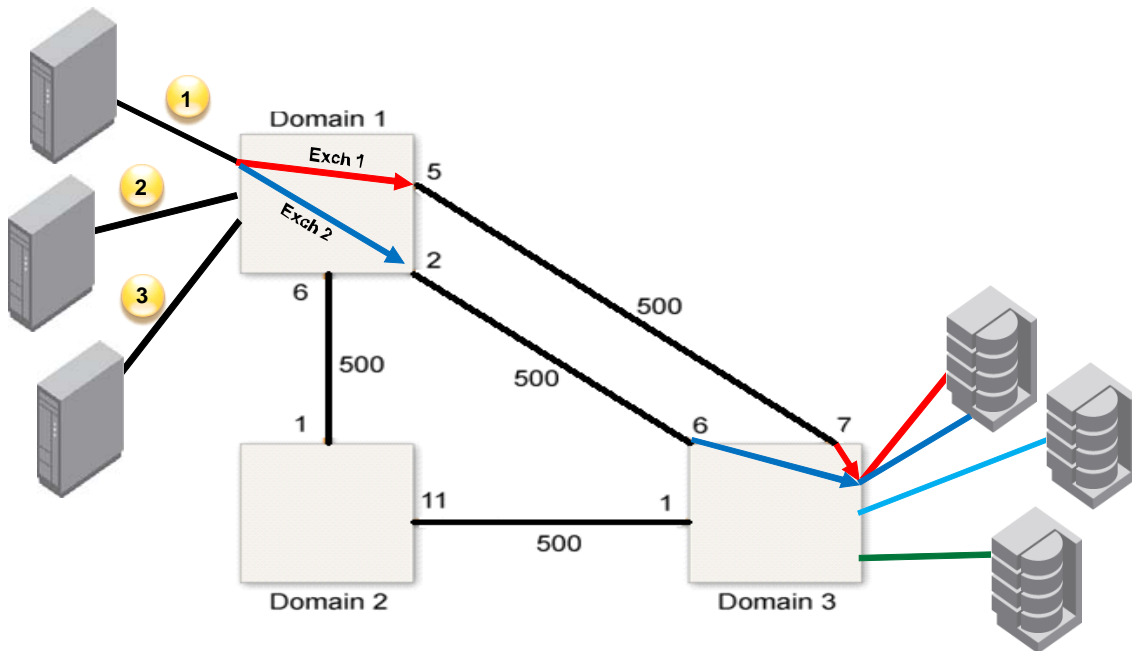


Figure 23: Exchange-Based Routing and DLS

4 and 8 Gbps Trunking Overview

- Automatically aggregates up to 8 ISLs when the switches are connected
 - Condor ASICs provide up to 32 Gbps of aggregate bandwidth
 - Condor2 ASICs provide up to 64 Gbps of aggregate bandwidth
- All ports in a trunk group must operate at a common speed
 - Condor ASICs support multiple 2/4 Gbps trunks between switches
 - Condor2 ASICs support multiple 2/4/8 Gbps trunks between switches
- Trunking port groups include: ports 0-7, 8-15, and so on
- An ISL Trunking license is required for all switches participating in trunking
 - Trunking is available when the license is installed and the ports are reinitialized
- Trunking is enabled by default
 - If it has been disabled, it must be re-enabled on the trunk ports using the `portcfgtrunkport` CLI command
- When trunking criteria is met, the trunk forms automatically




Figure 24: Trunking

The Deskew Counter

- Deskew values are related to distance and link quality
 - Deskew units represent the time difference for traffic to travel over each ISL as compared to the shortest ISL in the group
 - The system automatically sets the minimum deskew value of the ISL with the least latency (shortest round-trip time) to 15 deskew units
 - The deskew for the remaining ISLs is calculated in relation to the ISL with the least latency
- The deskew value is a representation of an ISLs transmission capabilities
 - Differences in deskew can be caused by signal degradation which affects the transmission time of frames through the link
 - Can also be caused by excessive differences in cable length
- Deskew values are displayed in the `trunkshow` command output

5 - Zoning

Maximum Zone Database Size

This is the maximum size of the local database (this switch) and may not be the fabric-wide max size. The switch with the lowest max zone database size, typically the switch with the lowest version of Fabric OS, will determine the maximum zoning database size in a fabric.

Use the `cfgsize` command to display the size details of the zone database. The size details include the Zone DB maximum size, the committed size, and the transaction size. All sizes are in bytes.

Zone DB max size is the upper limit for the defined configuration, determined by the amount of flash memory available for storing the defined configuration.

Committed size is the size of the defined configuration currently stored in flash memory.

Transaction size is the size of the uncommitted defined configuration. This value will be nonzero if the defined configuration is being modified, otherwise it is zero.

```
sw300:admin> cfgsize
Zone DB max size - 1045274 bytes
Available Zone DB size - 1044056 bytes
    committed - 206
    transaction - 0
```

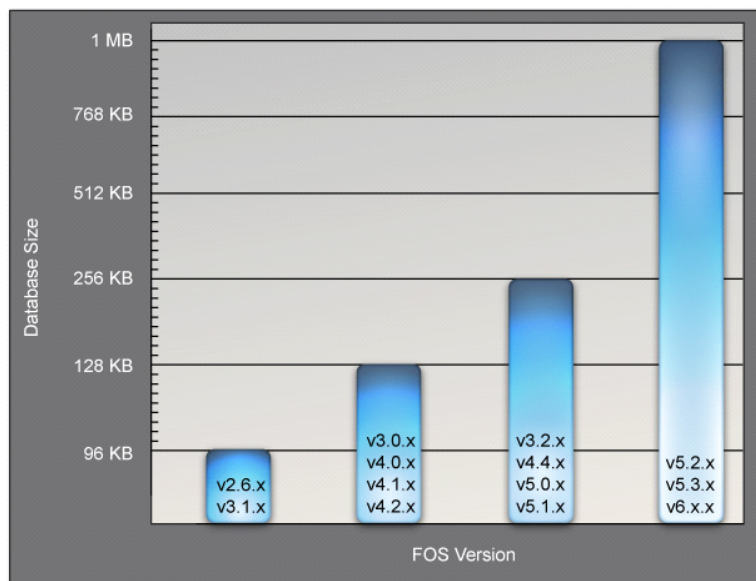


Figure 25: Maximum zoning database size

Zoning Best Practices

- Always implement zoning, even if LUN Masking is used
- Use PWWN identification for all zoning configuration unless special circumstances require domain,index identification (for example, FICON)
- Make zone object names only as long as they need to be to be meaningful
- Define all zones so that hardware enforcement is used
- Use single initiator zoning
- Use separate zones for tape and disk traffic if an HBA is carrying both types of traffic
- Implement `defzone --noaccess`
- Use the free Brocade SAN Health software to validate zoning configurations
- Analyze zones to verify correct devices are communicating
- Backup with `configupload`

Default Zoning

- In early versions of Fabric OS, when zoning was not implemented or a `cfgdisable` command was issued, all devices in the fabric could access each other
- Starting with Fabric OS v5.1.0+, a default zone is available, which:
 - Controls what device access is allowed within a fabric when zoning is not enabled
 - Can enable all device access with `defzone --allaccess` (default)
 - Can disable all device access with `defzone --noaccess`
- The `defzone` setting is:
 - In effect when a user-specified zone configuration is not enabled
 - Not in effect when a user-specified zone configuration is enabled

Types of Zoning Enforcement

- Session Enforcement
 - Name Server restricts PLOGIs
- Hardware Enforcement
 - Source device is denied access to destination device if they are not defined in the same zone
 - Available through ASIC hardware logic checking at the destination port
 - More secure than session enforcement
- Enforcement is based on how members in a zone are defined

Devices that are Session enforced cause any PLOGIs to the device to be rejected.

Devices that are Hardware enforced cause any frames that do not comply with the effective zone configuration to be rejected. This blocking is performed at the transmit side of the port where the source device is located. This is the highest level of protection for a device.

The decision for what enforcement a device receives is based on how the members in a given zone are defined.

Zone Membership	Example	2, 4 and 8 Gbps ASIC Enforcement
All Domain, Index	Z1="dom1, index1; dom1, index2"	Hardware
All WWNs	Z2="wwn1; wwn2; wwn3"	Hardware
Mixed	Z3="dom1, index3; wwn4"	Session

Table 6: Zoning Enforcement

6 - Management

Management Interfaces and Tools

- Command Line Interface
 - Serial Communication (HyperTerm or tip)
 - Telnet (port 23)
 - SSHv2 (port 22)
- SMI-S (Storage Management Initiative Specification)
- DCFM (Data Center Fabric Manager)
- Web Tools
 - HTTP
 - HTTPS requires a Digital Certificate to be installed on the switch
- SNMPv1 and SNMPv3

DCFM Overview and Features

Brocade Data Center Fabric Manager is a key component of the Brocade Data Center Fabric (DCF) architecture. It is designed to unify management of data center fabrics—from the storage ports to the Host Bus Adapters (HBAs) attached to physical or virtualized servers. It can be used to configure and manage the Brocade DCX Backbone along with Brocade Directors, routers, and switches. DCFM supports Brocade encryption capabilities for data-at-rest and HBA products.

Brocade DCFM manages pure Fabric OS fabrics, mixed Fabric OS/M-EOS fabrics, and pure M-EOS fabrics (Enterprise or Professional Plus Edition required) with product-specific element managers and enhanced group management functions.

- Centralizes management of Brocade fabrics within and across data centers, including those with different protocols
- Supports NPIV and server virtualization
- Maximizes productivity by automating tasks and providing easy-to-use operations
- Secures fabrics by managing user access controls and ensuring consistent security settings
- Delivers real-time and historical performance monitoring
- Adds the ability to do parallel firmware upgrades

DCFM is available in three editions that differ in supported features, hardware platforms, and scalability limits

- DCFM Enterprise
 - Top level management application for enterprise customers and large environments
 - Supports up to 9,000 switch ports and 20,000 devices
- DCFM Professional Plus
 - Upgrade to DCFM Professional
 - Manage up to four fabrics with either Fabric OS or M-EOS switches
 - Supports up to 2,560 switch ports
 - Includes all features of DCFM Enterprise except:
 - FICON support
 - DCX management (the DCX-4S can still be managed)
 - FX8-24 FCIP blade support
- DCFM Professional
 - No cost, included with switch purchase
 - Single fabric management for pure Fabric OS fabrics
 - Supports up to 10 switches, 640 switch ports, and 1,000 hosts or storage devices

Brocade SAN Health

- Automates documenting of a SAN and comes in different formats
- SAN Health is a free utility that aids in creating:
 - Comprehensive documentation
 - Historical performance graphs
 - Detailed topology diagrams
 - Best practice recommendations

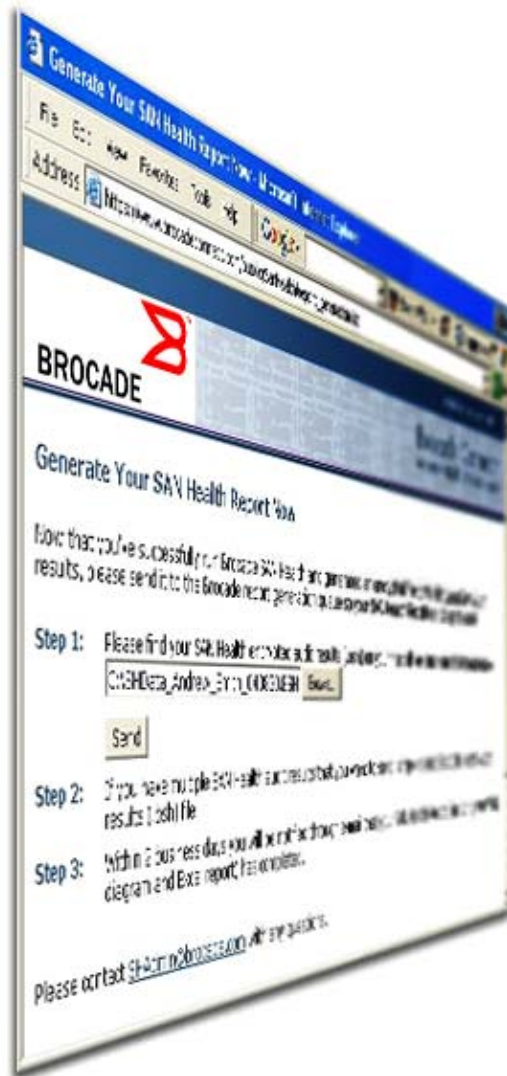


Figure 26: Brocade SAN Health

Host Connectivity Manager (HCM)

- HBA management tool loaded on the host that contains the HBAs
- Consists of two parts:
 - The GUI, Host Connectivity Manager software
 - The CLI, Brocade Command Line Utility (BCU)
- Used for HBA configuration, troubleshooting, and firmware upgrades

Note: Firmware upgrades for HBAs can also be performed using DCFM

Fabric Watch

- Optionally licensed per switch, it monitors the performance and status of the switch, including:
 - **Fabric events** – Fabric reconfigs, zone changes, and new logins
 - **Switch status** – Environmental (fans, power supplies, and temperature), SFP (Tx/Rx power, current, & voltage), Security, resource and FRU
 - **Port status** – Monitors F/FL/E_Port signal quality parameters
 - **Performance options** – monitor end-to-end performance
- Fabric Watch maintains a set of counters for each of the monitored conditions
 - Tracks the number of occurrences of each condition
 - Each counter is compared with an upper boundary and lower boundary
- Fabric Watch parameters can be updated several ways:
 - Using the `fwconfigure` command from the switch CLI
 - Using the `portthconfig` command from the switch CLI
 - By making the changes in the switch configuration file and downloading to the switch using the `configdownload` command from the switch CLI

SNMP Authentication

SNMPv1 Authentication between the NMS and the switch is performed using unsecure clear text “community strings” only.

SNMPv3 Authentication between the NMS and the switch is performed using an MD5 or SHA cryptographic hash of the “Username” stored on both the switch and the NMS station. In addition privacy is also supported in SNMPv3 by encrypting packets using DES. SNMPv3 supports three security levels:

- No authentication and no privacy – Username is not hashed and data is not encrypted.
- Authentication and no privacy – requires an auth password be configured.
- Authentication and privacy – requires an auth password and privacy password.
- No access – Disallow get or set requests altogether

The security subsystem can contain more than one security model. In the case of Brocade switches, for example, it contains: SNMPv1 with community strings; SNMP v3 user-names, authentication, and privacy; and Access Control Lists (ACLs). Brocade enabled SNMP ACLs in all Fabric OS versions.

Authentication occurs using either MD5 or SHA algorithm checks. MD5, message-digest algorithm, is an extension of the MD4 algorithm. The MD5 algorithm takes an input a message of arbitrary length and produces an output 128-bit "fingerprint" or "message-digest". It is designed for circumstances where a large file must be "compressed" in a secure manner before being encrypted with a private (secret) key under a public-key cryptosystem such as RSA. **RSA** is a public-key encryption technology created by Ron Rivest, Adi Shamir, and Leonard Adleman. RSA was built from the fact that, given that it is pretty easy to multiply two very large prime

Configuration Files

Configuration files, in general, are an editable, executable listing of configuration settings made on a switch during initial setup, and subsequent configuration.

- In Fabric OS v6.2.0+, the configuration file is divided into three sections:
 - Header
 - Chassis
 - One or more switch sections (logical switches)
- When upgrading firmware to Fabric OS v6.2.0 and later, pre-v6.2.0 configuration files are no longer valid
- It is recommended to perform the `configupload -all` command to upload both chassis and switch information after the firmware has been upgraded

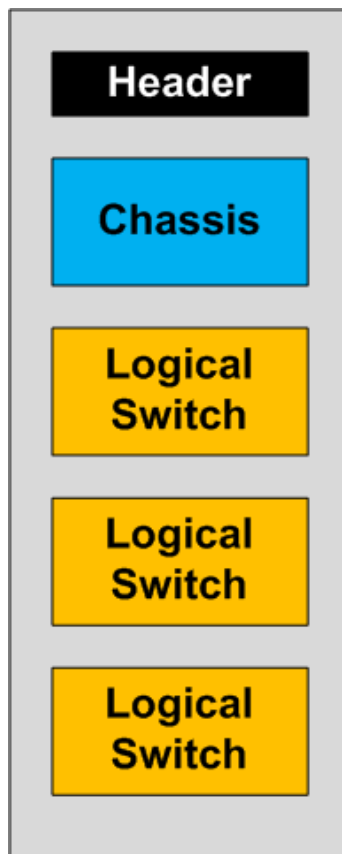


Figure 27: Configuration Files

Configuration Upload and Download

- Fabric OS v6.2.0 had no way to back up or restore the Virtual Fabric configuration
- Fabric OS v6.3.0 has resolved this issue with an option added to the existing `configupload/` `configdownload` CLI commands
 - The `-vf` option used in conjunction with these CLI commands will back up and restore the logical switch configuration
 - Example:

```
SW1:admin> configupload -vf -p ftp 10.20.1.10,user,  
/home/user/config/vf_config.txt,password
```

Firmware Download

Stand-alone Switches

1. The `firmwaredownload` command is entered
2. Firmware is downloaded to Secondary Partition
3. Primary and Secondary boot pointers are swapped
4. CP boots from firmware in new Primary Partition
5. Firmware in Primary Partition is committed to Secondary
6. Download complete

Chassis-based Switches

1. Run `firmwaredownload` command on the active CP
2. The standby CP blade downloads firmware
3. The standby CP blade reboots and comes up with the new Fabric OS
4. The active CP blade forces a failover and reboots to become the standby CP blade
5. The new standby CP blade (the active CP blade before the failover) receives the firmware from the now active CP
6. The new standby CP blade reboots and comes up with the new Fabric OS
7. The `firmwarecommit` command runs automatically on both CP blades

USB Storage

- Available on all Brocade 8 Gbit/sec products
- For a Backbone, only the Active CP is able to mount the USB device, however you can have USB drives installed on both CPs
- Only a Brocade branded USB device is supported
- Unsupported USB devices will fail with device not found message
- Brocade USB devices are pre-formatted with the required directory structure

```
r6-st02-dcx1:admin> usbstorage -l

firmwarekey\          0B      2007 Sep 25 13:54
config\               0B      2007 Sep 25 13:54
support\              0B      2007 Sep 25 13:54
firmware\             0B      2007 Sep 25 13:54

Available space on usbstorage 100%
```

Slow Drain Device Detection

- A slow drain device returns credits slower than the sender wants to send frames to it
- Slow drain can exist at any link utilization level
- Achieved throughput into the slow drain port is lower compared to intended throughput
- Slow drain spreads into the fabric and can slow down unrelated flows in the fabric
- Bottleneck detection finds bottlenecks at the egress side of the port
- Utilize the `bottleneckmon` command to enable/disable on fabric ports




Figure 28: Slow Drain Device Detection

7 - Troubleshooting

Problem Approach

- Gather and analyze data so that the problem can be divided and conquered and determine where the breakdown occurred
- Analyze problem and relevant information to determine the nature of the problem
 - What are the problem symptoms?
 - Is there a problem accessing a switch or switches in a fabric?
 - Is it a firmware download problem?
 - Is it a switch or end-device connectivity problem?
 - Does the support provider have release note documentation that could help?
 - Has anything changed?
 - Is the problem reproducible?
- What type of switch data was or can be gathered to help solve the problem?
 - Has at least one `supportsave` been collected?
 - If pre-Fabric OS v6.2 was it run on both CPs for directors¹?
 - Is the current switch configuration backed up?
- Are any best practice switch supportability processes being used?
 - Topology diagrams
 - `supportsave` / DCFM Data Collection
 - Baseline switch configuration saved
 - Product documentation
 - Support provider recommendations
 - Change management procedures
 - SAN Health

Gather Information

Below is a short list of possible issues and what tools to use to gather information:

Problem Area	Investigate	Tools
End-device connectivity	<ul style="list-style-type: none"> • Link, login • Fabric (zoning, security) • End-to-end device connectivity 	<ul style="list-style-type: none"> • Switch LEDs • Switch Commands • End device (storage/host) parameters
Switch connectivity	<ul style="list-style-type: none"> • Marginal Links • Incorrect zone, Access Control List (ACL), or Virtual Fabric (VF) configurations • Incorrect switch parameters 	<ul style="list-style-type: none"> • Switch LEDs • Product Knowledge • Switch commands • Web and GUI-based monitoring and management software tools
Firmware download	<ul style="list-style-type: none"> • Network connectivity • FTP parameters • Switch parameters 	<ul style="list-style-type: none"> • Use host tools plus analyze host FTP parameters • Check FTP parameters on switch
Switch access	<ul style="list-style-type: none"> • Network connectivity • Switch authentication or configuration parameters 	<ul style="list-style-type: none"> • Host tools including CLI • Switch: check network parameters and user authentication

Table 7: Gathering Troubleshooting Information

Diagramming the Fabric

- A fabric diagram is often needed to troubleshoot many problems
- A fabric diagram can be drawn using a Fabric OS `supportsave`
 - A `supportsave` from all core switches in the fabric is required to build a basic diagram or one from each switch in the fabric to map all ISL port numbers
- Use the following Fabric OS commands from `SSHOW_FABRIC` to map the switches in a fabric:
 - `fabricshow`
 - `islshow`
 - `trunkshow`

Common SAN Problems

- **Segmented Fabric**
 - Wrong product license, domain ID conflicts, zoning conflicts, incompatible switch parameters
 - `switchshow`, `configshow`, `fabricshow`, `fabstatssh`, `portshow`, `portcfgshow`, check zone related commands, and license configuration. You can also use the DCFM Zone Merge Tool to merge the zone configurations.
- **Configuration**
 - Port, device or switch is not correctly configured
 - `portcfgshow`, `configshow`, `portlogdump`, `portshow`, `fabricshow`, `trunkshow`, `portcfglongdistance`, `licenseshow`, and `portshow`
- **Missing Device**
 - `Nx_Port` is not registered with the name server, zoning is enabled and device is not in the zone, LUN masking is enabled and device is not properly defined, application accessing device is not configured correctly
 - Check physical connectivity using `switchshow`, `portshow`, and `fcping`. Check fabric connectivity with `nsallshow`, `nsshow`, `nscamshow`, `zoning(zoneshow, etc.)` and port configuration commands (`portcfgshow`, `portshow`). Optionally use a diagnostic tests such as `porttest`; this will test link components and port. LUN masking may also prevent a host from seeing the storage.
- **Timeout/sluggishness**
 - ISL congestion, insufficient BB-credits, marginal link, slow drain device
 - `urouteshow`, `topologyshow`, `porterrshow`, `portshow`, `portstatssh`, `portcfgshow`, `portbuffershow`, `aptpolicy`, and `bottleneckmon` (Fabric OS v6.3 or later). You can also use DCFM or Fabric Watch to identify congestion problems.
- **Licensing**
 - Customers do not have the license to do what they are attempting
 - `licenseshow`, `licenseadd`
- **Marginal Links**
 - Problems related to performance or problems that occur when connecting switches or end-devices can be related to marginal links or SFPs
 - Check physical connectivity, check `porterrshow`, `errdump` (look for Fabric Watch link events).
 - A marginal switch port is defined as a switch port that has one or both of the following conditions:
 - It is receiving a marginal incoming signal
 - It has a switch receiver that is not functioning properly
- **Zoning**
 - End-devices not able to access each other can be related to zoning
 - `configshow`, `configure`, `fabricshow`, `zonehelp`, `cfgshow`, `zoneadd`, DCFM Fabric Merge, `switchshow`, `errshow`
 - The `switchshow`, `errshow`, and `fabstatssh` commands can all be issued from a single fabric to determine why the segmentation occurred.

Web Tools Switch / FRU Status

The status of the environmentals can be found at the top of the window. The status of the switch ports can be determined by the colored outline on each port in the switch picture. The status of the switch can be found by checking the Switch Information tab in the Switch Events, Information section of the window. Check the *Hardware Reference Manual* for information on how to read the status of the LEDs.

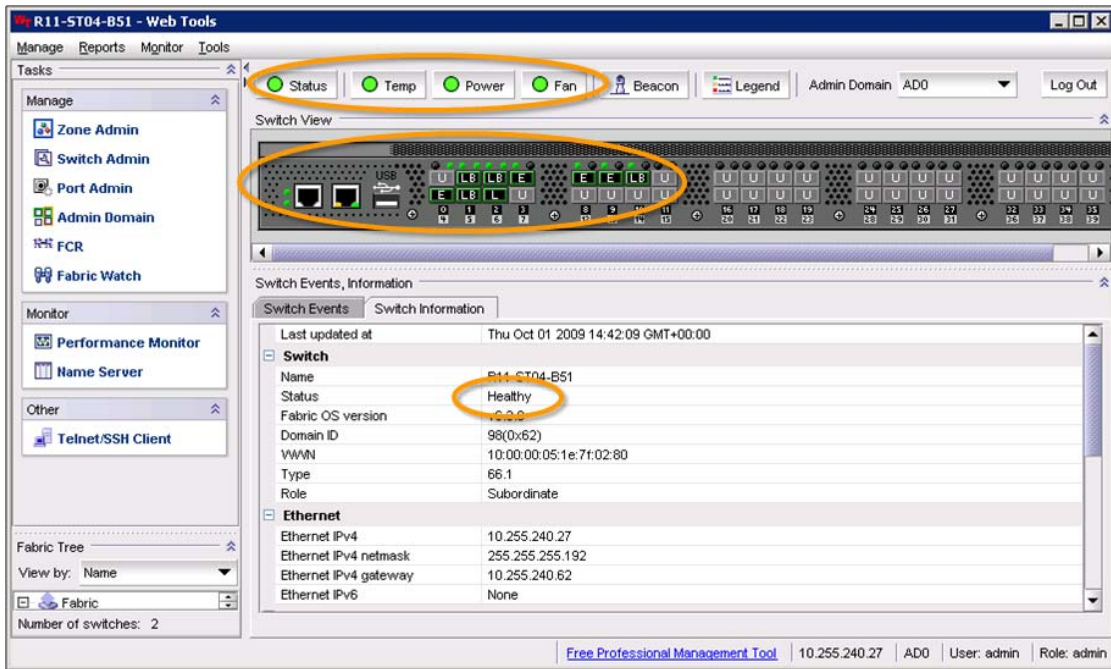


Figure 29: Web Tools Switch / FRU Status

Blade Status LEDs

- DCX WWN Card: LED indications:
 - 1 - 4 Port blades
 - 5 CR8 (core) 0
 - 6 CP8 0
 - 7 CP8 1
 - 8 CR8 (core) 1
 - 9 - 12 Port blades



Blade LEDs:
Amber: Needs attention
Green: Power

DCX WWN Card LEDs:
Amber: Needs attention
Green: Power




Figure 30: Blade Status LEDs

Switch / FRU Status Commands

Use the following CLI commands as a starting point to check the status of a switch, or its FRUs:

Command	Displays Status Of
<code>switchshow</code>	Switch and ports
<code>hashow</code>	CPs and failover
<code>slotshow</code>	Blades
<code>psshow</code>	Power supplies
<code>tempshow</code>	Temperature sensors
<code>fanshow</code>	Fans
<code>sensorshow</code>	Environmental sensors

Table 8: FRU Status Commands

HCM SupportSave

- Captures and packages the following data:
 - Firmware and driver traces
 - Supportshow information
 - Adapter configuration data
 - o Windows registry data
 - o Linux/VMware/Solaris `bfa.conf` file
 - Master and Application logs
- Stored in the local file system
- Supported from both HCM and BCU

Fabric OS Support Data

- For long-term support of a Fabric OS switch, begin gathering switch support data
- There are several different types of switch support data that can be collected; most of which is included in a supportsave capture
 - Switch error logs (RAS Logs)
 - Audit events
 - First Failure Data Capture (FFDC) files
 - Panic dump and core files
 - Trace dump files
 - Output from supportshow

Taking the Test

After the Introduction Screen, once you click on **Next**, you will see the non-disclosure agreement:



IMPORTANT: PLEASE READ THE FOLLOWING BROCADE NON-DISCLOSURE CONFIDENTIALITY AGREEMENT CAREFULLY BEFORE TAKING THIS EXAM.

The following Non-Disclosure Confidentiality Agreement (the "Agreement") sets forth the terms and conditions of your use of the exam materials as defined below.

The Disclosure to you of this Exam and any questions, answers, worksheets, computations, drawings, diagrams, or any communications, including verbal communication by any party, regarding or related to the Exam and such Exam Materials and any derivatives thereof is subject to the Terms and Conditions of this Agreement. You understand, acknowledge and agree:

- That the questions and answers of the Exam are the exclusive and confidential property of Brocade and are protected by Brocade intellectual property rights;
- That you may not disclose the Exam questions or answers or discuss any of the content of the Exam Materials with any person, without prior approval from Brocade;
- Not to copy or attempt to make copies (written, photocopied, or otherwise) of any Exam Material, including, without limitation, any Exam questions or answers;
- Not to sell, license, distribute, or give away the Exam Materials, questions, or answers.
- You have not purchased, solicited or used unauthorized (non-Brocade sanctioned) Exam Materials, questions, or answers in preparation for this exam.
- That your obligations under this Agreement shall continue in effect after the Exam and, if applicable, after termination of your Certification, regardless of the reason or reasons for terminations, and whether such termination is voluntary or involuntary.

Brocade reserves the right to take all appropriate actions to remedy or prevent disclosure or misuse, including, without limitation, obtaining an immediate injunction. Brocade reserves the right to validate all results and take any appropriate actions as needed. Neither this Agreement nor any right granted hereunder shall be assignable or otherwise transferable by you.

By clicking on the "A" button ("YES, I AGREE"), you are consenting to be bound by the terms and conditions of this agreement and state that you have read this agreement carefully and you understand and accept the obligations which it imposes without reservation. You further state that no promises or representations have been made to induce agreement and that you accept this agreement voluntarily and freely.

- Y. YES, I AGREE
- N. NO, I DO NOT AGREE

Next 

Figure 31: Sample NDA