BCFA in a Nutshell Study Guide for Exam 143-425

Brocade University

Revision 0214

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

All or some of the products detailed in this presentation may still be under development and certain specifications, including but not limited to, release dates, prices, and product features, may change. The products may not function as intended and a production version of the products may never be released. Even if a production version is released, it may be materially different from the pre-release version discussed in this presentation.
Nothing in this presentation shall be deemed to create a warranty of any kind, either express or implied, statutory or otherwise, including but not limited to, any implied warranties of merchantability, fitness for a particular purpose, or non-infringement of third-party rights with respect to any products and services referenced herein.

Revision 0214

BCFA in a Nutshell Gen 5 Edition

Brocade
Certified Fabric Administrator

Objective: The BCFA Nutshell guide is designed to help you prepare for the Brocade Certified Fabric Administrator certification, exam number 143-425.

Audience: The BCFA Nutshell self-study guide is intended for those who have successfully completed the CFA 200: Brocade Core Gen 5 SAN Administration 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 200: Brocade Core Gen 5 SAN Administration 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.

Joe Cannata
Certification Manager

Table of Contents

List of Figures

List of Tables

1 — Fibre Channel Concepts

After reviewing this section be sure you can perform the following:

- Demonstrate knowledge of the address components of a 24-bit address
- Identify valid port types
- Describe the election process, responsibilities, and attributes of a principal switch
- Identify Fibre Channel well-known addresses
- Describe flow control concepts

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.

Each switch is responsible for assigning unique hexadecimal addresses. Addresses are 24 bits and use the following format:

Domain ID (8 bits)	0x01 - 0xEF
Area ID (8 bits)	0x00 - 0xFF
Node Address (8 bits)	AL_PA / NPIV / Shared Area[1]

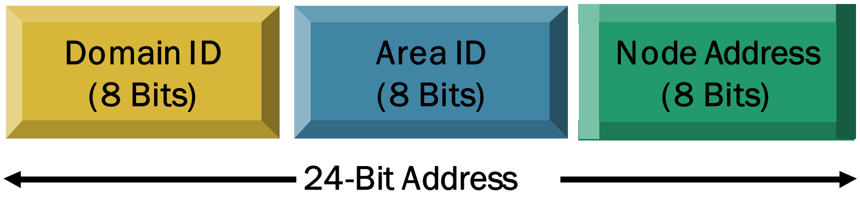

Figure 1: 24-bit Address

There are three different address types:

- Fabric DD AA 00

  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/NPIV DD AA PP

  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 arbitrated loop (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. More information on NPIV at the end of this module.

- Shared area DD AA 00/40/80/C0

  Shared Area IDs use the node address to allow more than 256 ports to be addressed in a single domain.

[1] There are other variations of addressing that are beyond the scope of this document.

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. Shared area addressing allows for more than 256 ports in a single domain. The FC8-48 has some ports that use shared areas[2]. Ports 16-47 of the FC8-48 blade use shared areas. Shared Area PIDs use a Node Address of either 0x00 or 0x80. An example of two shared areas on a FC8-48 blade in slot 1:

- 018000 – Port 16
- 018080 – Port 40

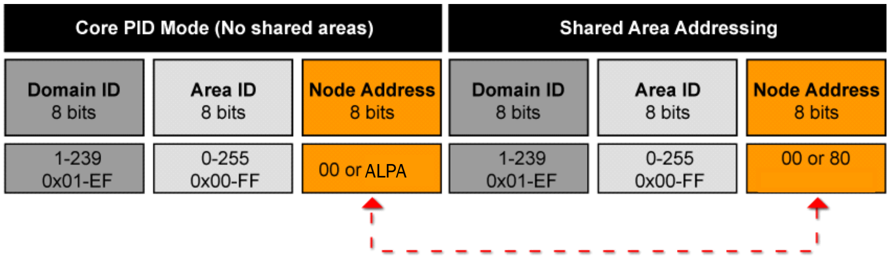

Core PID Mode (No shared areas)			Shared Area Addressing		
Domain ID 8 bits	Area ID 8 bits	Node Address 8 bits	Domain ID 8 bits	Area ID 8 bits	Node Address 8 bits
1-239 0x01-EF	0-255 0x00-FF	00 or ALPA	1-239 0x01-EF	0-255 0x00-FF	00 or 80

Figure 2: Overview of Shared Area Addressing

The ability to address more than 256 ports in a single switch required a change to the standard 24-bit addressing scheme. The second byte of a PID is referred to as the Area ID; with 8 bits, the Area ID can address ports 0–255. Brocade now uses the third byte of the PID to address ports 256-512. The third byte of a PID is referred to as the node address and was used to identify the loop address (ALPA) for a loop device. Since an FC8-48 port does not support loop devices the node address can be used to identify ports in the 256-512 range. This requires the Area ID to be shared (used twice).

[2] 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.

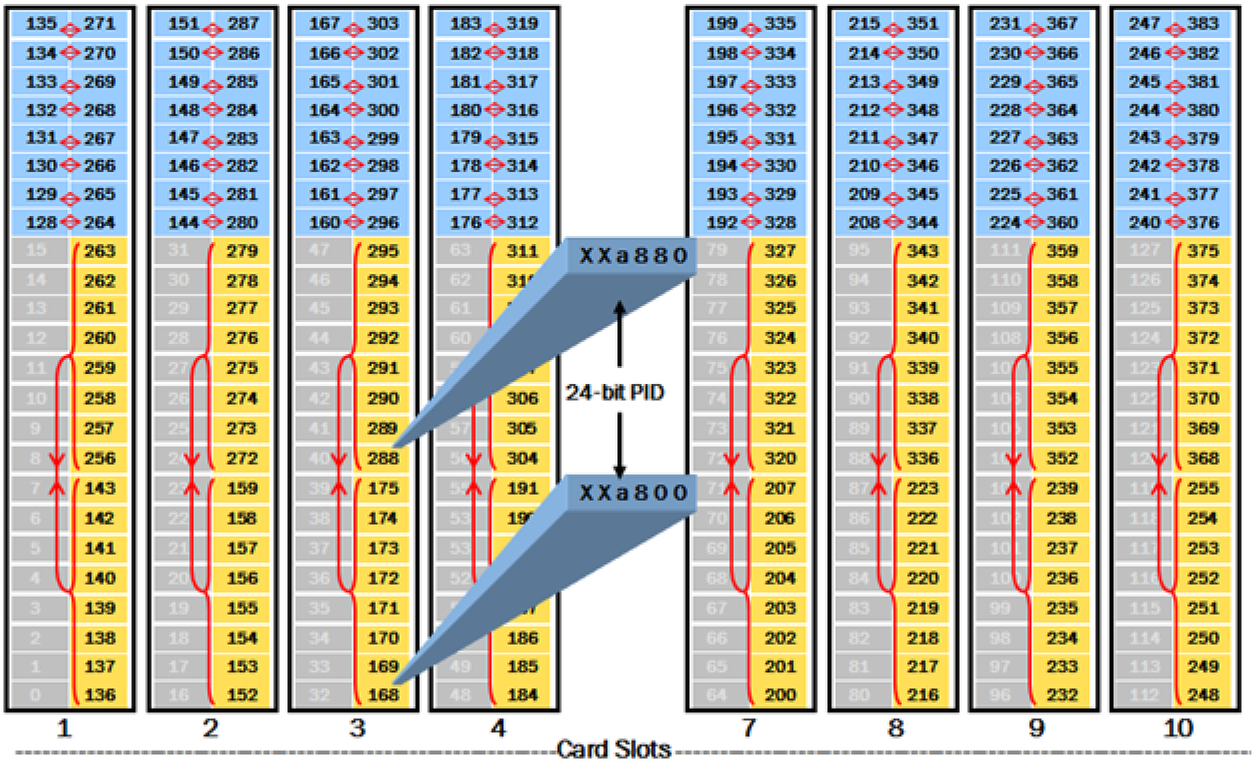

Figure 3: DCX Shared Area Addressing

The grey boxes represent port indexes 0-127. The area ID for these indexes is not shared. The blue and yellows boxes represent the port indexes on the Condor. With shared area IDs, the lower port number on the card has a node address of 0x00 while the higher port number on the card has a node address of 0x80.

Port Types

Device ports (Nx_Ports)

- N_Port: Node port, a fabric device directly attached
- NL_Port[3]: 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[3]: 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[4]: A type of E_Port used to connect to an FC router fabric
- VE_Port: Virtual E_Port (used in FCIP fabrics)
- VEX_Port[4]: VEX_Ports are no different from EX_Ports, except underlying transport is IP rather than FC
- D_Port: A configured port used to perform diagnostic tests on a link with another D_Port

For a switch port that goes through the fabric initialization process, it arrives at an ending status of F_Port, FL_Port or E_Port.

[3] Loop ports are not supported on Condor3 ASICs.

[4] EX and VEX_Ports allow communication between devices in independent fabrics without having to merge the fabrics. This is done through the use of FC-FC Routing. To learn more about FC-FC routing and fabric extension solutions please refer to CFP 300 course.

Fabric Initialization Process

When a port initializes it really is the beginning phase of a fabric initialization. The port could be an E_Port which would extend the fabric by adding an additional switch. If the port is an F_Port or an N_Port the fabric is still growing and because the added device causes a Registered State Change Notification (RSCN), this process initializes the fabric.

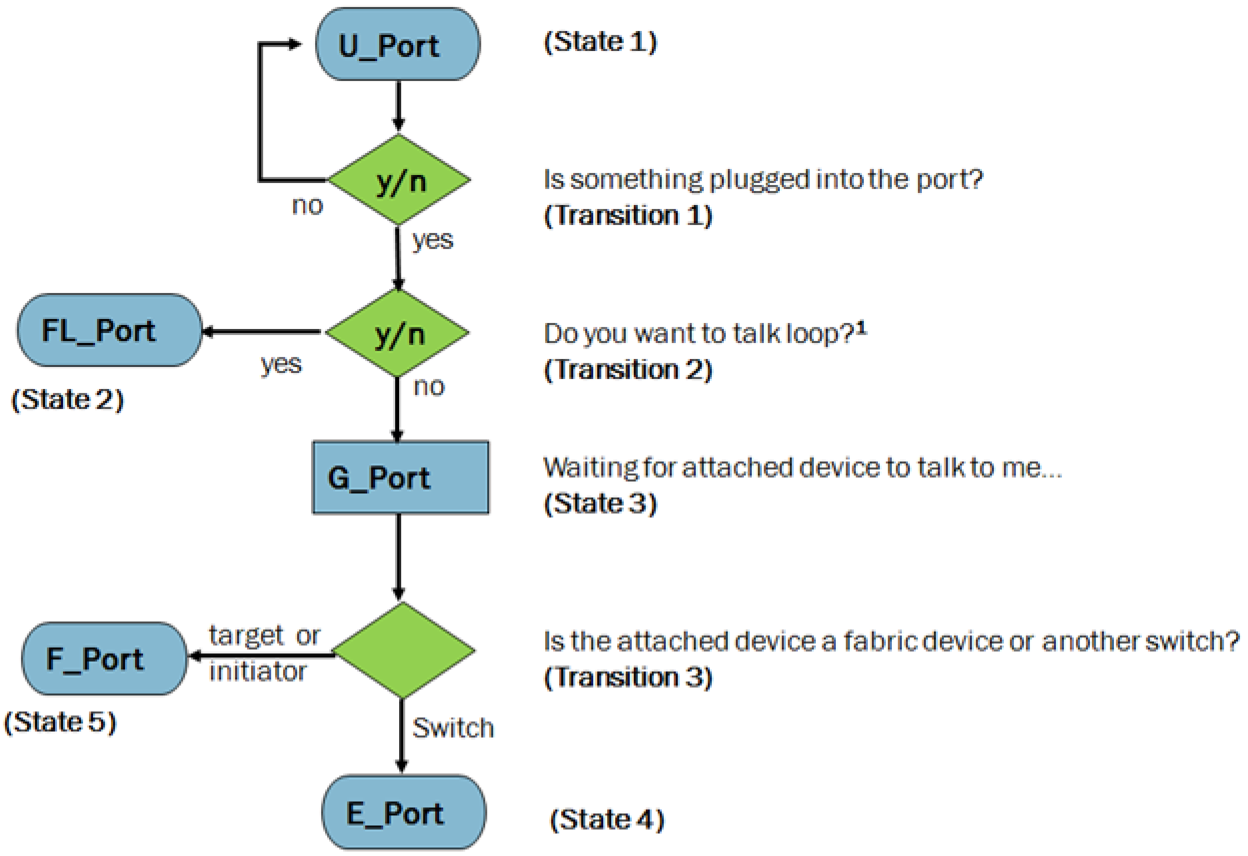
U_Port
(State 1)
y/n
no
yes
Is something plugged into the port?
(Transition 1)
FL_Port
y/n
yes
no
(State 2)
Do you want to talk loop?¹
(Transition 2)
G_Port
Waiting for attached device to talk to me...
(State 3)
F_Port
target or initiator
(State 5)
Is the attached device a fabric device or another switch?
(Transition 3)
Switch
E_Port
(State 4)

Figure 4: Fabric Initialization Process

TABLE 1 Fabric Initialization Process

Step	Process
1.	A Universal Port (U_Port) is the initial state of a port. (State 1)
2.	Is something connected (sending a light/electrical signal) to the port? If yes, continue. (Transition 1)
3.	U_Port starts mode detection process by transmitting at least 12 LIP (F7) Primitive Sequences. (Transition 2)[1]
4.	If at least 3 consecutive LIP Primitive Sequences are received, then the port enters OPEN_INIT state and attempts FC-AL loop initialization. (State 2)
5.	If LIP Primitive Sequences are not received, the U_Port attempts OLD_PORT initialization by taking the link down then transmitting NOS primitives. If Link Initialization Protocol fails after 1 retry or LIP received after 1 second, go to FC-AL initialization. (Transition 2)

TABLE 1 Fabric Initialization Process (Continued)

Step	Process
6.	When operating in the FL_Port mode, a U_Port will try the loop initialization procedure three times. If these fail, the port will be marked as faulty. To ensure N_Port, re-initialize the port and the switch port will cut the laser forcing a loss of signal state for at least 20 µs. Then the switch port will bring back the laser and issue NOSs. (Transition 2)
7.	If the attached device is not loop it continues into the G_Port stage. A device can be a switch, target (usually storage), or an initiator (usually a host). (State 3)
8.	If the attached device is a target or initiator it changes its port state from G_Port to an F_Port. (State 5)
9.	If the attached device is a switch then it changes its port state from a G_Port to an E_Port. (State 4).

[1] The flow chart outlines the Fabric Initialization process supported by Fabric OS, but not all Brocade hardware, such as the Condor3, support loop.

F_Ports and E_Ports continue to login in to the fabric through additional processes not discusses in this document.

Note

The firmware automatically attempts to re-initialize a faulty port every two seconds.

Fibre Channel uses a number of ordered sets (4 bytes) to perform the following control and signaling functions:

- Frame delimiters identify the start and end of frames (SOF, EOF).
- Primitive signals indicate events or actions:
- Replenishing used flow control buffer credits (R_RDY, VC_RDY).
- Fill words between frames when nothing else needs to be sent (IDLE, ARB).
- Primitive sequences are used for link initialization, recovering a link from a detected error and signaling a port offline (LIP, NOS, OLS, LR, LRR, etc). All primitive sequences require a minimum of three consecutive occurrences of the same ordered set before the primitive sequence is recognized as valid and action is taken.

Principal Switch

There two main functions of a principal switch:

- Manage build fabrics, such as ensuring unique domain IDs throughout the fabric
- Synchronize time throughout the fabric (you can also use Network Time Protocol (NTP) to synchronize time with the principal or within the fabric)

Fabric Shortest Path First (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 is used as the principal ISL.

Upstream means traffic going out an E_Port towards the principal switch. Downstream means traffic going out an E_Port away from the principal switch. These designations are seen in the `switchshow` output.

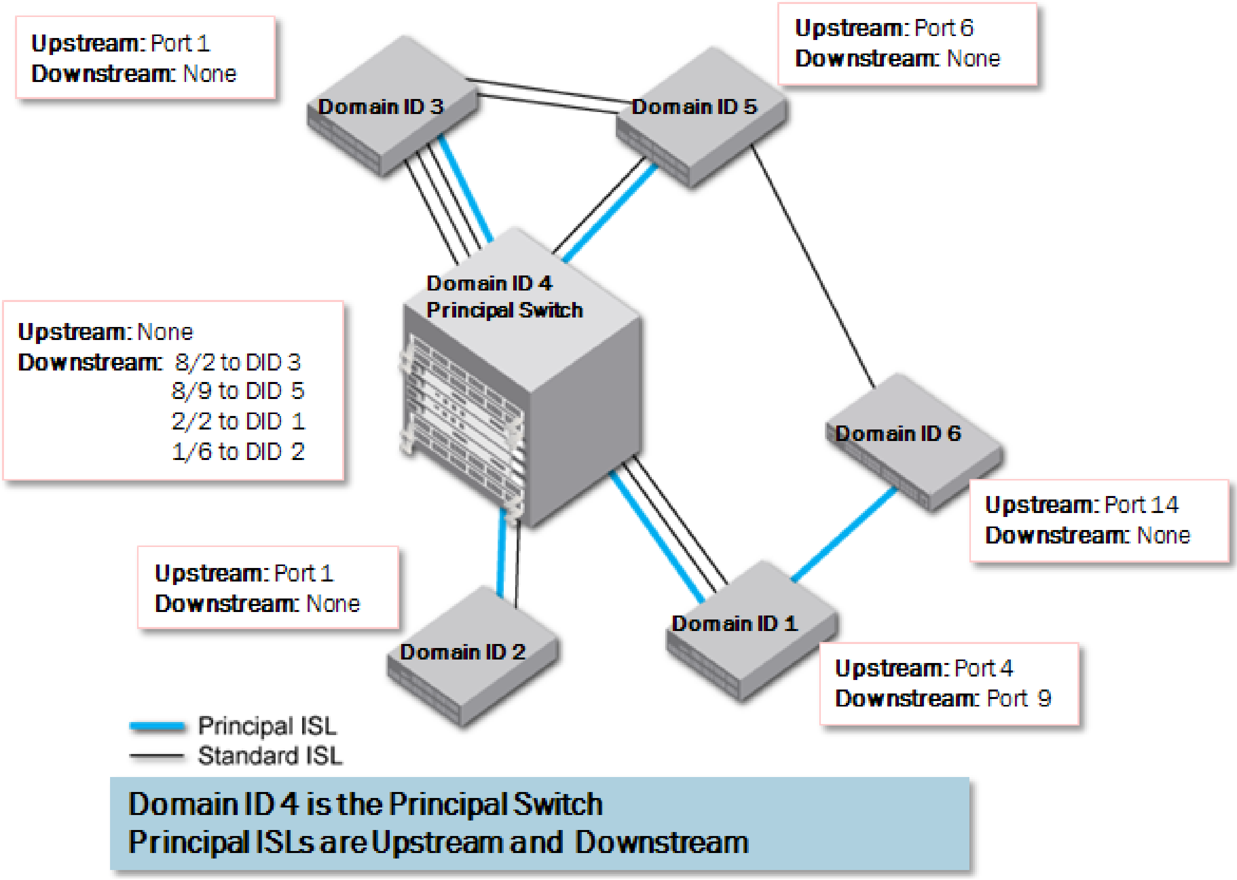

Figure 5: Principal Switch Path

Principal switch selection process:

- If none of the switches has a priority setting, switch with the lowest WWN becomes principal
- If switches have a priority setting, then only those switches participate in the selection
- A switch with the lowest priority becomes the principal switch
- If more than one switch has lowest priority, the switch with the lowest WWN becomes the principal switch

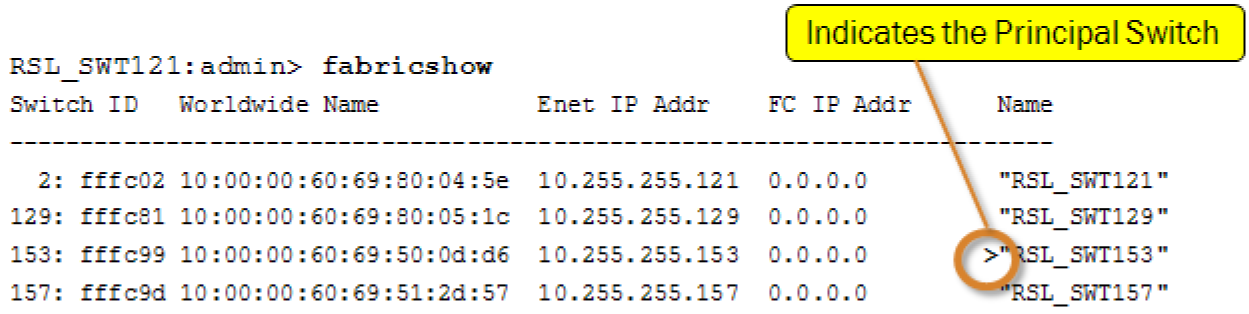

```
RSL_SWT121:admin> fabricshow
Switch ID   Worldwide Name           Enet IP Addr    FC IP Addr      Name
-------------------------------------------------------------------------
  2: fffc02 10:00:00:60:69:80:04:5e  10.255.255.121  0.0.0.0        "RSL_SWT121"
129: fffc81 10:00:00:60:69:80:05:1c  10.255.255.129  0.0.0.0        "RSL_SWT129"
153: fffc99 10:00:00:60:69:50:0d:d6  10.255.255.153  0.0.0.0       >"RSL_SWT153"
157: fffc9d 10:00:00:60:69:51:2d:57  10.255.255.157  0.0.0.0        "RSL_SWT157"
```

Indicates the Principal Switch

Figure 6: Principal Switch Indicator

Synopsis:

```
fabricprincipal --help|-h
fabricprincipal [--show|-q]
fabricprincipal --enable [-priority|-p priority] [-force|-f]
fabricprincipal --disable
fabricprincipal [-f] mode
```

Example of setting the preferred principal switch priority using:

```
SW1:admin> fabricprincipal --enable -p 0x01 -f
```

`-p` Sets the principal selection priority for the switch

`-f` Forces a fabric rebuild immediately after enabling on a switch

Example of disabling the mode setting:

```
switch:admin> fabricprincipal 0
Principal Selection Mode disabled
```

Domain IDs

Although domain IDs are assigned dynamically when a switch is enabled, you can change them manually so that you can control the ID number or resolve a domain ID conflict when you merge fabrics.

If a switch has a domain ID when it is enabled, and that domain ID conflicts with another switch in the fabric, the conflict is automatically resolved if the other switch's domain ID is not persistently set. The process can take several seconds, during which time traffic is delayed. If both switches have their domain IDs persistently set, one of them needs to have its domain ID changed to a domain ID not used within the fabric.

The default domain ID for Brocade switches is 1.

When Insistent Domain ID (IDID) is enabled, the switch's ID does not change even after a reboot. If you want to add the switch with IDID enabled into an existing fabric, you need to verify that each switch has a unique domain ID or disable IDID using the `configure` command.

Well-Known Addresses

Every switch has reserved a 24-bit address used for ‘Well-Known Addresses’. Table 2 lists the services and their associated addresses.

TABLE 2 Well-Known Addresses

Address	Description
FFFFFE	**Fabric Login Server:** 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 24-bits and are a combination of the domain ID plus the port area number of the port where the node is attached
FFFFFD	**Fabric Controller:** Provides state change notifications to registered nodes when a change in the fabric topology occurs
FFFFFC	**Directory Server:** The Directory Server, or Name Server, is where fabric/public nodes register and query to discover other devices in the fabric
FFFFFA	**Management Server:** Assists in the autodiscovery of switch-based fabrics and their associated topologies, such as exchanging fabric names
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.
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.
FFFFFF	**Broadcast Server:** When a frame is transmitted to this address, the frame is broadcast to all operational N and NL_Ports.

Buffer Credit Management

Buffer-to-buffer (BB) credit management affects performance over distances; therefore, allocating a sufficient number of buffer credits for long-distance traffic is essential to performance.

To prevent a target device (either host or storage) from being overwhelmed with frames, the Fibre Channel architecture provides flow control mechanisms based on a system of credits. Each of these credits represents the ability of the device to accept additional frames. If a recipient issues no credits to the sender, no frames can be sent. Pacing the transport of subsequent frames on the basis of this credit system helps prevent the loss of frames and reduces the frequency of entire Fibre Channel sequences needing to be retransmitted across the link.

Buffer credits are exchanged during the FLOGI/PLOGI process. After a device has issued a FLOGI and received a fabric address it will then register with the name server using a PLOGI and will include the number of buffer credits the device supports.

Because the number of buffer credits available for use within each port group is limited, configuring buffer credits for extended links may affect the performance of the other ports in the group used for core-to-edge connections. You must balance the number of long-distance ISL connections and core-to-edge ISL connections within a switch.

The optimal number of buffer credits is determined by the distance (frame delivery time), the processing time at the receiving port, link signaling rate, and size of the frames being transmitted. As the link speed increases, the frame transmission time is reduced and the number of buffer credits must be increased to obtain full link utilization, even in a short-distance environment.

Use Web Tools, Network Advisor, or the `portcfglongdistance` command to specify an Extended Fabric Distance level:

- Level 0 static mode (L0) is the normal mode for a port.
- Level E static mode (LE) reserves a static number of buffer credits that supports distances up to 10 km. The number reserved depends on the port speed.
- Dynamic long distance Mode (LD) calculates buffer credits based on the distance measured during port initialization.
- An upper limit is placed on the calculation by providing a desired distance value.
- Static long distance mode (LS) calculates a static number of buffer credits based on a desired distance value.

Note

L0 and LE modes do not require a license. Use of LD and LS modes requires an Extended Fabric License

Port-level Credit Recovery

This provides a more robust credit recovery for lost buffer credits or frames on long distance E_Ports on between two non-Condor3 ASIC switches or backbones. If a lost buffer credit or frame is detected, the switch performs a Link Reset (LR) to recover the lost credit. This is done by sending an LR to the target switch which sends back an LRR (Link Reset Response). Because this happens on an E_Port, the link does not reset, only the BB (frame and credit loss) counters at both ends of the link are reset.

- Only supported on long distance E_Port links
  - Only LE, LD, and LS long distance modes are supported
  - R_RDY mode is supported
  - EX, VE, and VEX ports are not supported
- Tracks both buffer credits and frames sent
- E_Port Credit Recovery is enabled by default2
  - Use the following CLI command to disable this feature:

    ```
    portcfgcreditrecovery --disable [slot/port]
    ```

  - This feature is only supported using ARBs as fill words. If IDLEs are being used as fill words, this feature must be disabled.

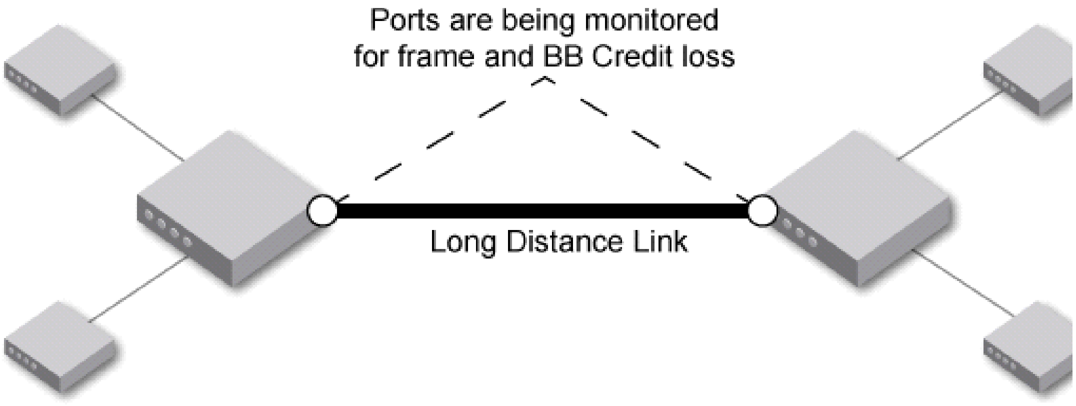

Figure 7: BB Credit Loss Monitoring

2 — Product Hardware Features

After reviewing this section be sure you can perform the following:

- Describe how to obtain switch environmental data
- Identify Brocade hardware components

Obtaining Chassis and Component Status

There are three tools you use to view various information on a chassis, including port information:

- Brocade Network Advisor
- Fabric Watch
- Command Line Interface (CLI)

The commands in the following table provide status and environmental information about the chassis and its components. These commands provide information only, and they do not interrupt traffic flow. For more information about these commands, refer to the *Fabric OS Command Reference*.

TABLE 3 Environmental Status and Maintenance Commands

Command	Information Displayed
sensorShow	Temperature readings for the port blades
	Temperature readings for the CP blades
	Status and RPM of all operational fans
	Status of all operational power supplies
tempShow	Temperature readings for the port blades
	Temperature readings for the CP blades
psShow	Status of all operational power supplies
fanShow	Status and RPM of all operational fans
chassisShow	Serial number, time awake, and additional information about each component
slotShow	Slot occupancy
errShow errDump	System error log. Refer to the *Fabric OS Message Reference* for more information on the messages in this log

The switchstatuspolicyshow Command

Use the switchstatuspolicyshow command to view the current policy parameters set for the switch. These policy parameters determine the number of failed or non-operational components allowed before triggering a status change in the switch. For port-related contributors, the numbers are expressed as a percentage of physical ports present in the switch at any given time.

To display the switch policy parameters on a Brocade DCX 8510-8:

```
switch:admin> switchstatuspolicyshow
The current overall switch status policy parameters:
                Down       Marginal
----------------------------------
PowerSupplies   0          0
Temperatures    0          0
Fans            1          0
WWN             0          0
CP              0          0
Blade           0          0
CoreBlade       0          0
Flash           0          0
MarginalPorts   0.00%[0]   0.00%[0]
FaultyPorts     0.00%[0]   0.00%[0]
MissingSFPs     0.00%[0]   0.00%[0]
ErrorPorts      0.00%[0]   0.00%[0]
Number of ports: 4
```

Optics Overview

Transceivers are used to transmit data over fiber or copper cabling. Brocade switches, HBAs, and FAs require Brocade-branded optics. Most Fibre Channel transceivers are tri-mode:

- 16 Gbps SFP+ supports 16, 8, and 4 Gbps speeds
- 8 Gbps SFP+ supports 8, 4, and 2 Gbps speeds
- 4 Gbps SFPs support 4, 2, and 1 Gbps speeds

Exceptions are the 10 Gbps SFP+ and 4x16 Gbps QSFP1 which only synch at their respective speeds. QSFPs (Quad-SFPs) are a new transceiver design that allows for four separate data paths through the transceiver and across an industry standard cable. Currently these are only being used in the 16 Gbps Brocade backbones (DCX 8510-8 and -4).

Note

The mSFP transceivers are used only with the FC8-64 port blade. Narrower OM-3 LC cables are used to connect the FC8-64.

Brocade DCX 8510-8 Features

The following are key features of the Brocade DCX 8510-8:

- Up to 384 16 Gbps end-user ports in a single chassis, enabling high density SAN configurations with reduced footprint
- Support for 2, 4, 8, 10, and 16 Gbps auto-sensing Fibre Channel ports. Trunking technology groups up to eight ports to create a high performance of up to 128-Gbps aggregate bandwidth ISL trunks between switches
- The Brocade DCX 8510-8 supports 10 Gbps FC-type SFPs in 16 Gbps port blades only and 10 GbE SFPs in the FX8-24 and FCOE10-24 application blades. The two types of SFPs are not interchangeable
- The 10 Gbps ports can be configured manually on only the first eight ports of the 16 Gbps port blades
- Support for all of the application, port blade, and control processor (CP) blades supported in the Brocade 8510-4 (with the exception of the Core Switch Blade), thereby providing flexible system configurations and fewer types of new blades
- Up to six chassis can be connected with the use of 4x16 Gbps quad SFP (QSFP) inter-chassis links (ICLs)
- Support for high-performance port blades running at 2, 4, 8, 10, or 16 Gbps, enabling flexible system configuration
- Redundant and hot-swappable control processor and core switch blades, power supplies, blower assemblies, and WWN cards that enable a high availability platform and enable nondisruptive software upgrades for mission-critical SAN applications
- Universal ports that self-configure as E_Ports (10 Gbps ports are E_Ports only), F_Ports, EX_Ports and M_Ports (mirror ports)
- Diagnostic port (D_Port) functionality
- In-flight data cryptographic (encryption/decryption) and data compression capabilities through the 16 Gbps port blades
- Fibre Channel over IP (FCIP) functionality through the FX8-24 blade

16 Gbps Port Blades

16 Gbps FC ports support E, F, EX, Diagnostic, and Mirror ports. They do not support FL_Ports. End-user ports can operate at 2, 4, 8, 10, and 16 Gbps using 8, 10, 16 Gbps SFPs. 10 Gbps ports can be configured as E_Ports only.

Diagnostic port (D_Port) for Condor3-based ports support electrical and optical loopback (16 Gbps Brocade branded optics only), cable length detection, and spinfab-like cable saturation tests (enough to saturate a 16 Gbps link).

10 Gbps FC using 10GE license:

- Existing slot-based 10GbE FCIP license (introduced in FOS v6.3 for FX8-24 blades) is extended to enable Condor3 FC ports running at 10 Gbps rate. These ports need to work with DWDM equipment with plain transponder cards
- When applied to a 16 Gbps blade, license allows all ports in the first octet to be enabled as 10 Gbps FC

Port Blade Compatibility

The figure below shows the compatibility of the various port blades with the DCX chassis.

Figure 8: Port Blade Compatibility Matrix

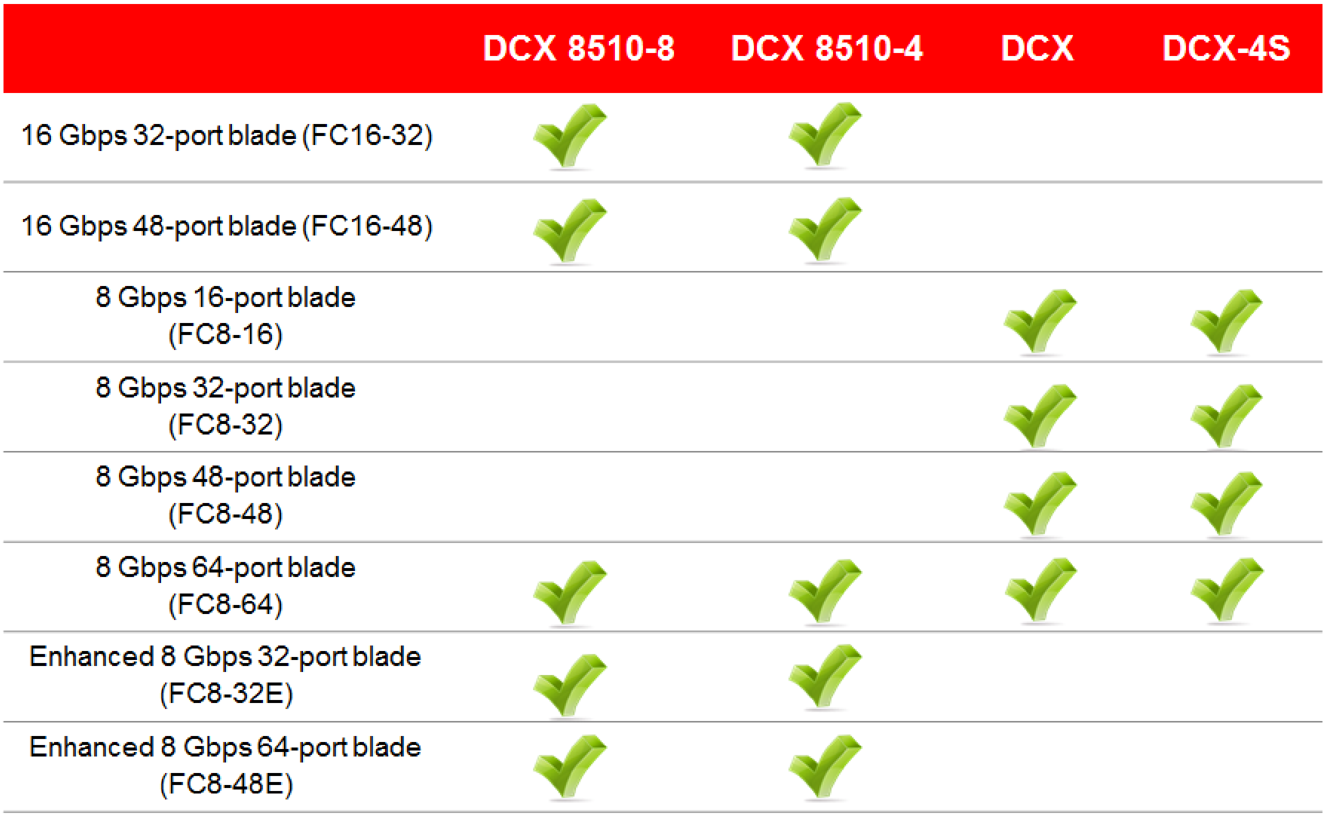

	DCX 8510-8	DCX 8510-4	DCX	DCX-4S
16 Gbps 32-port blade (FC16-32)	✓	✓		
16 Gbps 48-port blade (FC16-48)	✓	✓		
8 Gbps 16-port blade (FC8-16)			✓	✓
8 Gbps 32-port blade (FC8-32)			✓	✓
8 Gbps 48-port blade (FC8-48)			✓	✓
8 Gbps 64-port blade (FC8-64)	✓	✓	✓	✓
Enhanced 8 Gbps 32-port blade (FC8-32E)	✓	✓		
Enhanced 8 Gbps 64-port blade (FC8-48E)	✓	✓		

Figure 9: Extension and Application Blade Compatibility Matrix

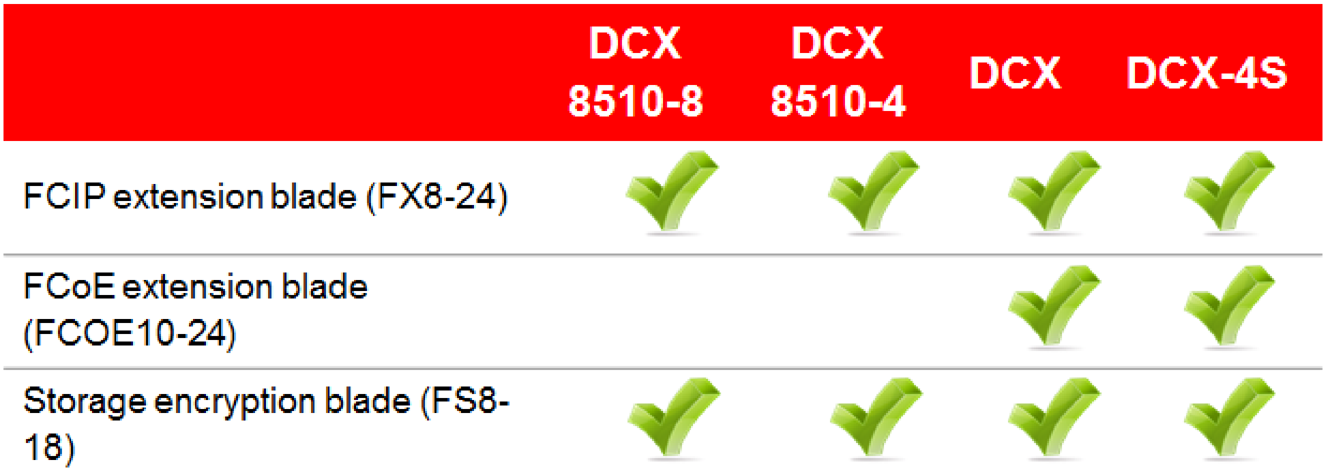

	DCX 8510-8	DCX 8510-4	DCX	DCX-4S
FCIP extension blade (FX8-24)	✓	✓	✓	✓
FCoE extension blade (FCOE10-24)			✓	✓
Storage encryption blade (FS8-18)	✓	✓	✓	✓

8 Gbps Inter-Chassis Link Overview

Inter-Chassis Links (ICLs) provide dedicated connections between DCX and DCX-4S chassis. ICL connectors are located on the CR8 and CR4S-8 routing blades. Allows up to three domains to be connected and requires Fabric OS v6.3 (earlier versions of firmware only support two domains).

For FICON purposes, the ICL connection is not considered a hop.

ICL link features for DCX and DCX-4S:

- Speed locked at 8 Gbps
- Copper-based proprietary cable and connector
- No SFPs
- Each cable provides up to 16 x 8 (128) Gbps bandwidth on the DCX and up to 8 x 8 (64) Gbps bandwidth on the DCX-4S
- Licensed feature
- ICL cables are 2 meters in length
- Allows for ISL connections without consuming user ports

16 Gbps Inter-Chassis Link Overview

16 Gbps ICL ports are used to connect DCX 8510-8/8510-4 chassis together without consuming user ports. Each connection uses a special SFP (called a Quad SFP or QSFP) that carries four 16 Gbps connections. These cables are industry standard optical cables, as opposed to the copper cables used on the DCX and DCX-4S. Using optical cables provides for increased speeds and reliability over longer distances.

- Uses a QSFP to connect
- Uses standard optical cables for distances up to 50 meters
- Each cable provides 4x16 Gbps connections
- Multiple cables from the same blade must be used to establish trunks

Brocade 6510

A 48-port 16 Gbps Gen 5 Fibre Channel switch based on the Condor3 ASIC.

TABLE 4 B6510 Hardware Attributes

Attributes	Comment
1U form factor	
FRUs: Two 125 W power supply/fan assemblies	Power supplies are available in front-to-back or back-to-front airflow options.
Two reversible airflow options	Port side to non-port side and non-port side to port side
USB port	
48 FC ports	
Ports on Demand (12-port increments)	The Brocade 6510 comes with 24 licensed ports. Additional ports are available in 12 port increments (36 and 48 ports).
16 Gbps FC port speed	Supports 2/4/8/10/16 Gbps speeds with the following optics: • 16 Gbps optics: 4/8/16 Gbps • 8 Gbps optics: 2/4/8 Gbps • 10 Gbps FC optics: 10 Gbps Support for 10 Gbps FC is by using a 10 Gbps FC optic. 10 Gbps Ethernet optics do not work.
E, EX, F, Diagnostic, and Mirror ports	
1 Condor3 ASIC	The Condor3 supports the Data In-Flight Encryption and Compression features.
7712 user BB credits per ASIC	The Condor3 has a total of 8192 buffer credits, 7712 are available to the user in the Brocade 6510.
1:1 subscription ratio	
6 x 8-port trunk groups	
Support for Integrated Routing (IR)	Brocade 6510 Integrated Routing support requires an Integrated Routing license.EX_Ports can be enabled or disabled on a per-port basis as needed.
Virtual Fabrics	Up to 4 logical switches.
Access Gateway-capable	

NOTE
The two rack kit options for the Brocade 6510 use rails that are slimmer than standard rails to accommodate the slightly wider chassis. Be sure to use one of these kits. Do not use standard rails to install the Brocade 6510 in a rack, they will not fit with the switch.

Brocade 6520

A 96-port 16 Gbps Gen 5 Fibre Channel switch based on the Condor3.

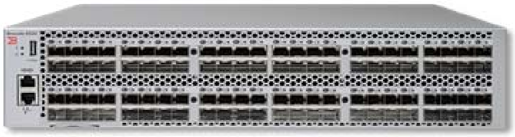

TABLE 5 B6520 Hardware Attributes

Attributes	Comment
2U form factor	
FRUs: Two power supplies, three fan assemblies	Power supplies are available in front-to-back or back-to-front airflow options.
Two reversible airflow options	Port side to non-port side and non-port side to port side
USB port	
96 FC ports	
Ports on Demand (24-port increments)	The Brocade 6520 comes with 48 licensed ports. Additional ports are available in 24 port increments (72 and 96 ports).
16 Gbps FC port speed	Supports 2/4/8/10/16 Gbps speeds with the following optics: • 16 Gbps optics: 4/8/16 Gbps • 8 Gbps optics: 2/4/8 Gbps • 10 Gbps FC optics: 10 Gbps Support for 10 Gbps FC is by using a 10 Gbps FC optic. 10 Gbps Ethernet optics do not work.
E, EX, F, Diagnostic, and Mirror ports	
6 Condor3 ASICs	The Condor3 supports the Data In-Flight Encryption and Compression features.
7712 user BB credits per ASIC	The Condor3 has a total of 8192 buffer credits, 7712 are available to the user in the Brocade 6510.
1:1 subscription ratio	
12 x 8-port trunk groups	
Support for Integrated Routing (IR)	Brocade 6510 Integrated Routing support requires an Integrated Routing license.EX_Ports can be enabled or disabled on a per-port basis as needed.
Virtual Fabrics	Up to 4 logical switches.
Access Gateway-capable	

Brocade 1860 Fabric Adapter

Industry's first multiprotocol server connectivity product. It is available in single and dual-port models. Using AnyIO, it combines:

- 16 Gbps Host Bus Adapter (HBA)
- 10 Gigabit Ethernet (GbE) Network Interface Card (NIC)
- 10 Gbps Converged Network Adapter (CNA)

Ports can be configured in three modes, HBA, NIC, or CNA and configuration is on a per-port basis. One port can run as an HBA and the other configured as a CNA on dual-port models. You can use Brocade Network Advisor and Host Connectivity Manager to install and configure the fabric adapter.

The Brocade 1860 Fabric Adapter can be managed using any one of these tools:

- Host Connectivity Manager
- Brocade Command Utility
- Brocade Network Advisor

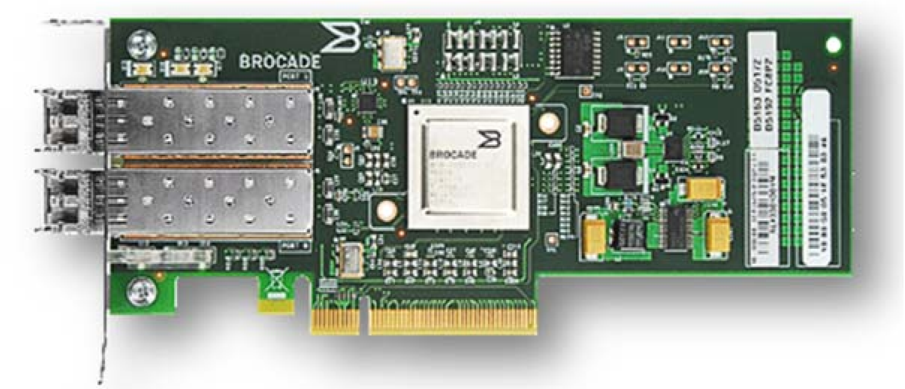

Figure 10: Brocade 1860 Fabric Adapter

Brocade AnyIO Technology

Brocade AnyIO technology is a unique capability that allows a single Brocade 1860 Fabric Adapter to support either native 16 Gbps Fibre Channel or 10 GbE on a port-by-port, user selectable basis. Brocade AnyIO technology enables a dual-port adapter to run 16 Gbps Fibre Channel on one port, and 10 GbE on the other port. In addition, the Brocade 1860 can run TCP/IP, Fibre Channel over Ethernet (FCoE), and iSCSI simultaneously on the same 10 GbE port.

A Brocade 1860 adapter port can be configured in any of the following modes:

- HBA mode — Appears as a 16 Gbps Fibre Channel HBA to the operating system (OS). This mode utilizes the Brocade Fibre Channel storage driver. An 8 or 16 Gbps Fibre Channel SFP can be installed for the port. The port provides Host Bus Adapter (HBA) functions on a single port so that you can connect your host system to devices on the Fibre Channel SAN. Ports with 8 Gbps SFPs configured in HBA mode can operate at 2, 4, or 8 Gbps. Ports with 16 Gbps SFPs configured in HBA mode can operate at 4, 8, or 16 Gbps.
- NIC mode — Appears as a 10 GbE NIC to the OS. It supports 10 GbE with DCB, iSCSI, and TCP/IP simultaneously. This mode utilizes the Brocade network driver. A 10 GbE SFP or direct attached SFP+ copper cable must be installed for the port.
- CNA mode — Appear as 10 Gbps FCoE ports when discovered in HCM and as as “10 GbE NIC” to the operating system. This mode provides all functions of Ethernet or NIC mode, plus adds support for FCoE features by utilizing the Brocade FCoE storage driver. A 10 GbE SFP or direct attached SFP+ copper cable must be installed for the port. Ports configured in CNA mode connect to an FCoE switch. These ports provide all traditional CNA functions for allowing Fibre Channel traffic to converge onto 10 Gbps DCB networks. The ports even appear as network interface controllers (NICs) and Fibre Channel adapters to the host. FCoE and 10 Gbps DBS operations run simultaneously.

3 – Installation and Configuration

After reviewing this section be sure you can perform the following:

- Describe the various types of Brocade documentation
- Describe product setup functions and features

Documentation

There are many different ways to gain information about any of Brocade products. Following are the most common:

- Administrative guides

  There is an administrative guide for each version of OS that Brocade produces, including some specialized features such as Access Gateway, FICON, FCIP, Web Tools, and Brocade Network Advisor. The administrative guides provide feature and configuration information, such as configuring the different features you can use in Adaptive Networking, Virtual Fabrics, and Zoning.

- Reference manuals

  There is a reference manual created for each piece of hardware Brocade produces, including for Fabric OS, such as the Fabric OS Command Reference. Reference manuals include configuration techniques specific to a hardware device and its components. The hardware reference manuals will also contain port numbering templates for various switches.

- Compatibility matrix

  The compatibility matrix summarizes equipment known to be compatible with the Brocade product family. Products named in the compatibility tables reflect equipment tested at Brocade or tested externally, such as tape libraries, SFPs, server adapters, and fabric management applications.

- Release notes

  Release notes are provided for each release and patch of OS that Brocade produces. It includes useful information such as the latest features released with the version of software, open and closed defects, blade support matrix, recommended migration paths for the released OS version, and standards compliance.

Basic Configuration Tasks

Each colored rectangle represents a task in setting up a Brocade SAN switch.

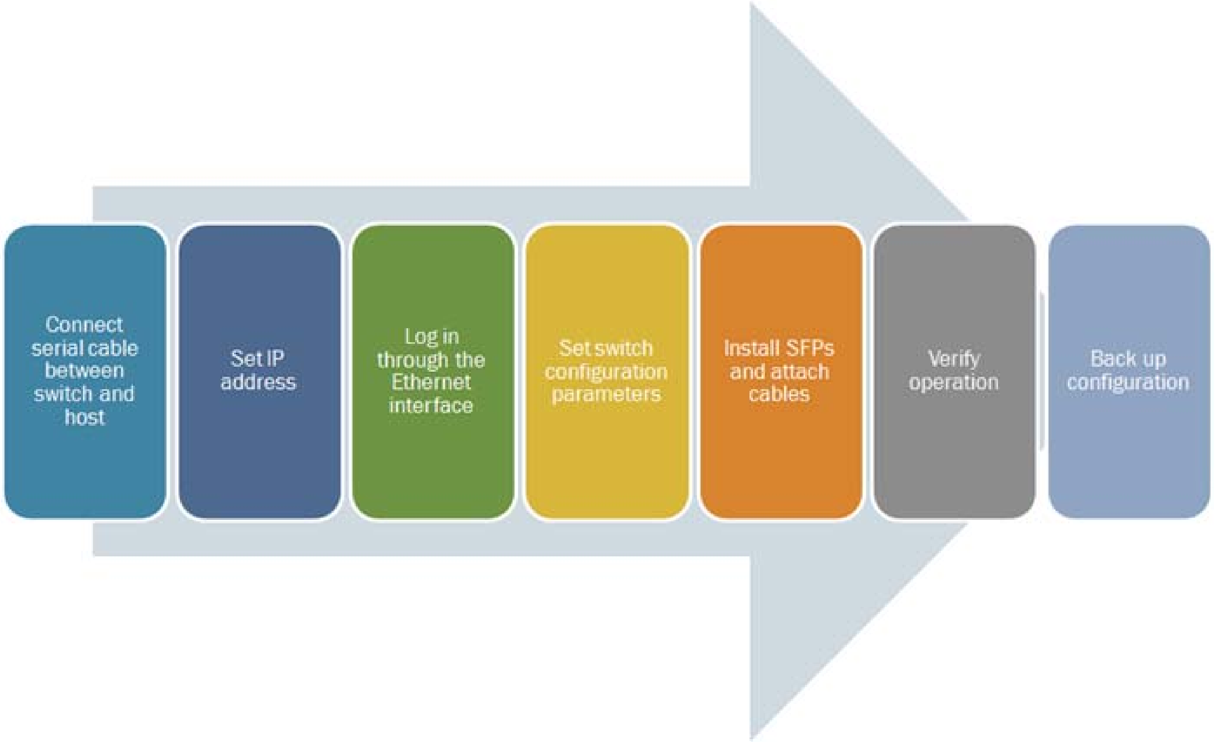
Connect serial cable between switch and host
Set IP address
Log in through the Ethernet interface
Set switch configuration parameters
Install SFPs and attach cables
Verify operation
Back up configuration

Figure 11: Overview of Basic Configuration Tasks

Connect Serial Cable Between Switch and Host

When a new switch has arrived for installation into a fabric, it is suggested to use a serial cable to configure the switch with an IP address. After the IP address is configured, the serial connection to the switch may be dropped and an SSH, telnet, or Web Tools session may be used for further switch configuration because of its convenience and speed.

To complete this portion of the setup you will need the following tools:

- A PC with:
  - Terminal emulator software
  - An available COM port
- A UNIX system with:
  - Tip utility
  - An available serial port
- Cable:
  - The required serial cable is provided with the switch

Use the following parameters to configure the serial connection:

- Bits per second: 9600
- Data bits: 8
- Parity: None
- Stop bits: 1
- Flow control: None

Installation steps

1. Insert the serial cable provided to an RS-232 serial port on the workstation. FOS switches use a straight-through cable.
2. Verify the switch has power and is past the POST stage.
3. Enter the ipaddrset command to set the IP address, subnet mask, and default gateway.

Set the IP Address

IP addresses are assigned to the management interface of a switch, director, or backbone and used to remotely manage the switch through telnet or SSH. A switch (e.g. B300, B5100) only has a single management interface and only uses a single IP address. Backbones require three IP addresses: one for chassis/switch management and one for each CP blade. The IP addresses used for the CP blades always connect to the blades they are assigned, the chassis management IP always connect to the active CP. Additionally, the DCX backbones have two Ethernet management interfaces on each CP (eth0 and eth3). These interfaces use port bonding to create a logical interface (bond0) to which the IP address of the CP is assigned. Ethernet bonding provides link (physical) layer redundancy using an active/standby model. By default all traffic is transmitted over the active interface, eth0. If eth0 experiences a link failure (e.g. cable unplugged) then eth3 becomes active. For more information on port bonding refer to the `ifmodeshow` command in the *Fabric OS Command Reference*.

- Use the `ipaddrset` command to configure the switch IP address
  - Default IP Address: 10.77.77.77
  - Default netmask: 255.255.255.0
- Obtain addressing information for your network
  - IP address and netmask
  - Default gateway
- Backbones require more than one IP address on the same subnet
  - One IP address required for switch management
  - This IP address is dynamically assigned to the active CP
  - One IP Address required per CP
  - Default IP Addresses for directors: 10.77.77.77 (switch management),10.77.77.75 (cp0), 10.77.77.74 (cp1)

Log in Through the Ethernet Interface

- Multiple concurrent Telnet sessions are allowed
- Use `killtelnet` to terminate a Telnet connection
- Login using a standard Telnet or SSHv2 client
- Telnet access may be disabled to force administrators to connect through an encrypted SSHv2 session using the `ipfilter` command

Set Switch Configuration Parameters

Following are tasks included in setting the switch configuration parameters:

- Setting the Domain ID
- Set the Fabric-wide Clock
- Set Switch Time Zone
- Set the Switch Name
- Set the Chassis Name
- Set syslog Server
- Set the Fabric Name
- Set the Default Port Names
- Port Settings

Set Command Line Session Timeout

When changing the timeout value you can use the `login` command to restart the login session and use the new timeout value.

- 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
- Default timeout on switches is 10 minutes.

The following example displays the current setting, type `timeout` with no arguments:

```
SW1:admin> timeout
Current IDLE Timeout is 0 minutes
SW1:admin> timeout 15
IDLE Timeout Changed to 15 minutes
The modified IDLE Timeout will be in effect after NEXT login

SW1 login: admin
Password:
SW1:admin> timeout
Current IDLE Timeout is 15 minutes
```

Set Banners

The following are two different types of banners you can set in Fabric OS.

- Message of the Day (MOTD) sets the banner on the chassis and displays before you login

  Example Example: `motd --set`

  ```
  SW1:admin> motd --set “Access by unauthorized personnel is prohibited."
  ```

- Login banner uses the `bannerset` command to set the banner on the chassis and it displays after you successfully log in

  Example Example: `bannerset`

  ```
  SW1:admin> bannerset “You have successfully logged into the switch.”
  ```

Activate Licensed Features

Activating your licenses enable Fabric OS features on the blade, switch, or chassis. Licenses are based on the switch license ID. A license string is up to 32 case-sensitive characters and a single license key may activate one feature or a bundle of features.

The following are useful license commands:

- `licenseidshow`
- `licenseshow`
- `licenseadd`
- `licenseremove`
- `licenseslotcfg` (Only used on backbones to install specific slot-based licenses)

Example of the `licenseshow` command output:

```
sw1:admin> licenseshow
c9SdQeRedATeRK:
    Fabric license
bRyRc9RRzbci3Sd6:
    Full Ports on Demand license - additional 16 port upgrade license
PWrGHFKXFHQ4EAMSPWNJFXXSfEYKY7C9BJ9MH:
    Enhanced Group Management license
RR3rEXQKTXHHmR3gLYRgF3ttSXS7KM9rBJtSK:
    8 Gig FC license
QDQfZS3QWaPmrSDHfrMRGXYFrffr3F9LB7tYN:
    Extended Fabric license
    Fabric Watch (Fabric Vision capable) license
    Performance Monitor (Fabric Vision capable) license
    Trunking license
    Integrated Routing license
    Adaptive Networking - obsolete license
7aACCMPLDAfrRrFJXHZEGHS3FfLf9HAtB7LTA:
    Server Application Optimization - obsolete license
```

In the output above you will notice that some of the licenses show as obsolete. These are examples of features that required a license but are now a part of the base Fabric OS and the license is no longer required.

Data In-Flight Encryption and Compression

- Both features are disabled by default and can be optionally enabled on a per port basis.
- No license needed to enable encryption or compression
- Supported only on E_Ports: ISL trunks (maximum 2 ports/trunk), QoS and long distance supported
- Supported in Virtual Fabrics (with XISLs) and non-VF modes
- Not supported on EX_Ports, ICLs, F_Ports, M_Ports, LISLs, or Access Gateway and FIPS modes
- Not supported with R_RDY flow control
- Encryption and compression combined introduce 5.5 µsec of latency for each end of the link
- Two ports on Brocade 6510 and 4 ports per FC16-32 and FC16-48 blades with encryption, compression or both
- The authentication policy must be enabled prior to enabling encryption

TABLE 6 Available Brocade Licenses

License	Description
10 Gigabit FCIP/Fibre Channel License (10G license)	Allows 10 Gbps operation of FC ports on the Brocade 6510 switch or the FC ports of FC16-32 or FC16-48 port blades installed on a Brocade DCX 8510 enterprise-class platform. • Enables the two 10GbE ports on the FX8-24 extension blade when installed on the Brocade DCX,DCX-4S, DCX 8510-4, or Brocade DCX 8510-8 enterprise-class platform. • Allows selection of the following operational modes on the FX8-24 blade: - 10 1GbE ports and 1 10GbE port, or - 2 10GbE ports • License is slot based when applied to a Brocade enterprise-class platform. It is chassis based when applied to a Brocade 6510 switch. NOTE: After installing the license you need to enable the 10 GbE ports using the portcfgoctetspeedcombo command.
7800 Upgrade License	Enables full hardware capabilities on the Brocade 7800 base switch, increasing the number of Fibre Channel ports from four to sixteen and the number of GbE ports from two to six. • Supports up to eight FCIP tunnels instead of two. • Supports advanced capabilities like tape read/write pipelining. NOTE: The Brocade 7800 switch must have the Upgrade License to add FICON Management Server (CUP) or Advanced Accelerator for FICON.
Adaptive Networking with QoS	Enables QoS SID/DID Prioritization and Ingress Rate Limiting features. These features ensure high priority connections by obtaining the bandwidth necessary for optimum performance, even in congested environments. • Available on all 8 and 16 Gbps platforms.

TABLE 6 **Available Brocade Licenses (Continued)**

License	Description
Advanced Extension License	Enables 2 advanced extension features: FCIP Trunking and Adaptive Rate Limiting. • FCIP Trunking feature allows all of the following: • Multiple (up to 10) IP source and destination address pairs (defined as FCIP Circuits) using multiple (up to 10) 1 GbE or 10 GbE interfaces to provide a high bandwidth FCIP tunnel and failover resiliency. • Support for 4 of the following QoS classes: Class-F, high, medium and low priority, each as a TCP connection. • Adaptive Rate Limiting feature provides a minimum bandwidth guarantee for each tunnel with full usage of available network bandwidth without any negative impact to throughput performance under high traffic load. • Available on the Brocade 7800 switch, and the Brocade DCX and DCX-4S and the Brocade DCX 8510 family for the FX8-24 on an individual slot basis.
Advanced FICON Acceleration	Allows use of specialized data management techniques and automated intelligence to accelerate FICON tape read and write and IBM Global Mirror data replication operations over distance, while maintaining the integrity of command and acknowledgement sequences. • Available on the Brocade 7800 switch, and the Brocade DCX and DCX-4S and the Brocade DCX 8510 family for the FX8-24 on an individual slot basis.
Brocade Advanced Performance Monitoring	Enables performance monitoring of networked storage resources. • Includes the Top Talkers feature.
Brocade Extended Fabrics	Provides greater than 10km of switched fabric connectivity at full bandwidth over long distances (depending on the platform this can be up to 3000km). NOTE: This license is not required for long distance connectivity using licensed 10G ports.
Brocade Fabric Watch	Monitors mission-critical switch operations. • Includes Port Fencing capabilities.
Brocade ISL Trunking	Provides the ability to aggregate multiple physical links into one logical link for enhanced network performance and fault tolerance. • Includes Access Gateway ISL Trunking on those products that support Access Gateway deployment. To enable this feature once the license is installed, run either of the following command sequences: - `switchdisable;switchenable` (enables the entire switch, is disruptive) - `portdisable;portenable` (enables on the specifed ports, is disruptive on the ports only)
Brocade Ports on Demand	Allows you to instantly scale the fabric by provisioning additional ports using license key upgrades. NOTE: Applies to the Brocade 300, 5000, 5100, 5300, 6510, and VA-40FC switches.
DataFort Compatibility License	Provides ability to read, write, decrypt, and encrypt the NetApp DataFort-encrypted Disk LUNs and Tapes to all of the following: • Brocade Encryption Switch • Brocade enterprise platforms with FS8-18 blade Includes metadata, encryption and compression algorithms. NOTE: Availability is limited. Contact your vendor for details.
Encryption Performance Upgrade License	Provides additional encryption bandwidth on encryption platforms. For the Brocade Encryption Switch, two Encryption Performance Upgrade licenses can be installed to enable the full available bandwidth. On a Brocade enterprise platforms, a single Performance License can be installed to enable full bandwidth on all FS8-18 blades installed in the chassis.

TABLE 6 **Available Brocade Licenses (Continued)**

License	Description
Enhanced Group Management	Enables full management of the device in a data center fabric with deeper element management functionality and greater management task aggregation throughout the environment. This license is used in conjunction with Brocade Network Advisor application software. This license is applicable to all of the Brocade 8 and 16 Gbps FC platforms. NOTE: This license is enabled by default on all 16G FC platforms, and on DCX and DCX-4S platforms that are running Fabric OS v7.0.0 or later. This license is not included by default on 8G FC fixed port switches (5300, 5100, VA-40FC, 300 and 8G FC embedded switches).
FCoE License	Included with the Brocade 8000 switch; enables Fibre Channel over Ethernet (FCoE) functions.
FICON Management Server (Also known as "CUP", Control Unit Port)	Enables host-control of switches in mainframe environments.
High Performance Extension over FCIP/FC (formerly known as "FC-IP Services")	Includes the IPsec capabilities. Applies to FR4-18i blade.
ICL 16-link License	Provides dedicated high-bandwidth links between two Brocade DCX chassis, without consuming valuable front-end 8 Gbps ports. Each chassis must have the ICL license installed in order to enable the full 16-link ICL connections. Available on the DCX only.
ICL 8-Link License	Activates all eight links on ICL ports on a Brocade DCX-4S chassis or half of the ICL bandwidth for each ICL port on the Brocade DCX platform by enabling only eight links out of the sixteen links available. This allows you 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. This license is also useful for environments that wish to create ICL connections between a DCX and a DCX-4S; the latter cannot support more than 8 links on an ICL port. Available on the Brocade DCX and DCX-4S platforms only
Inter Chassis Link (2nd POD) License	Provides dedicated high-bandwidth links between two Brocade DCX 8510-8 chassis, without consuming valuable user ports. Each chassis must have an ICL license installed in order to enable all available ICL connections. (Available on DCX 8510-8 only.)
Inter Chassis Link (1st POD) License	Activates half of the ICL bandwidth on a DCX 8510-8, or all the ICL bandwidth on a DCX 8510-4, allowing you to purchase less bandwidth and upgrade to a 2nd POD license at a later time. This license is useful for environments that wish to create ICL connections between a DCX 8510-8, and a DCX 8510-4; the latter platform supports only half the number of ICL links that the former platform supports. Available on the Brocade DCX 8510-8 and 8510-4 platforms only.
Enterprise ICL (EICL) License	The EICL license is required on a Brocade DCX 8510 chassis when that chassis is connected to four or more Brocade DCX 8510 chassis via ICLs.
Integrated Routing	Allows any ports in a Brocade 5100, 5300, 6510, 7800, and VA-40FC switches, the Brocade Encryption Switch, or the Brocade DCX, DCX 8510 family, and DCX-4S platforms to be configured as an EX_Port supporting Fibre Channel Routing (FCR).
Server Application Optimization	Optimizes application performance for physical servers and virtual machines. • Extends virtual channels across server infrastructure. • Enables configuration, prioritization, and optimization of application specific traffic flows. NOTE: This license is not supported on the Brocade 8000. For more information on this license, refer to the *Brocade Adapters Administrator's Guide.*

Port Speeds

Individual port speeds can be set by the administrator using the following command:

```
portcfgspeed <slot/port>,<speed_level>
```

Set the speed level for all ports on a switch[5]

```
0   -  Auto Negotiate (Hardware)
1   -  1Gbps
2   -  2Gbps
4   -  4Gbps
8   -  8Gbps
10   - 10Gbps
16   - 16Gbps
ax   - Auto Negotiate (Hardware) + retries
s    - Auto Negotiate (Software)
```

The SFP must be able to negotiate the hard-coded port speed, otherwise the port will not come up. The `switchshow` and `portshow` command outputs display Mod_Inv (Mod_Inv status can also result from using a non-Brocade SFP in a 8 or 16 Gbps capable port).

Configuring 10 Gbps Fibre Channel

A Condor3 ASIC has six octets, each of which contains eight ports. You can configure up to three different speed combinations. When you configure a given port, the combination applies to all ports in the octet. You can specify the octet by any port within the octet. To change the first octet, for example, you can specify any port from 0 through 7 as a port argument value. The following speed combinations are supported:

- Autonegotiated or fixed port speeds of 16 Gbps, 8 Gbps,4 Gbps, and 2 Gbps (combo option 1)
- Autonegotiated or fixed port speeds of 10 Gbps, 8 Gbps,4 Gbps, and 2 Gbps (combo option 2)
- Autonegotiated or fixed port speeds of 16 Gbps and 10 Gbps (combo option 3)

When configuring 10 Gbps Fibre Channel the octet speed combo must be set to 2 or 3 to support 10 Gbps speeds.

- Port must be disabled first

  ```
  portcfgoctetspeedcombo [<slot>/]<port> <combo>
  ```

- Port must be configured for 10 Gbps operation

  ```
  portcfgspeed [<slot>/]<port> 10
  ```

If the current combo configuration does not support 10 Gbps or is the 10G license is not present the command fails with an error message.

Install SFPs and Attach Cables

This is a physical process. You can use the `portshow` and `portcfgshow` commands to determine operational status.

[5] On Condor2 platforms 8 Gbps switches need a Brocade-branded 4 Gbps SFP for a port to run at 1 Gbps otherwise it is not supported. On Condor3 platforms, 2 Gbps support is allowed on 8 Gbps SFPs. 4 Gbps support is allowed on 8 Gbps and 16 Gbps SFPs. The default setting for `portcfgdefault` is 0. AX uses normal hardware autonegotiation with software retries (includes the s option) and allows for negotiation with certain problematic HBAs. If you know that you are having a hardware problem, choose the s option and it will autonegotiate using software only

Verify Operation

Display the overall status of switch using the `switchstatusshow` command. Display current policy settings with the `switchstatuspolicyshow` command.

- Marginal Status
  - Yellow color when displayed in Web Tools
  - Entry in error log, viewed with `errshow`, flagged as marginal
- Down Status
  - Red color when displayed in Web Tools
  - Entry in error log, viewed with `errshow`, flagged as faulty

Display temperature, power supply, and fan status using the `sensorshow` command.

Verifying Switch Operation

The `switchshow` command can be used to verify the switch is operating correctly and display information about the switch status.

```
R13-ST01-B5100:admin> switchshow
switchName:R13-ST01-B5100
switchType:66.1
switchState:Online
switchMode:Native
switchRole:Subordinate
switchDomain:2
switchId:fffc02
switchWwn:10:00:00:05:1e:0c:c9:c3
zoning:ON (Zone_Cfg)
switchBeacon:OFF
FC Router:OFF
FC Router BB Fabric ID:1
Address Mode:0


Index Port Address Media Speed State     Proto
==============================================
   0   0   020000   --     N8   Online      FC F-Port 10:00:00:00:c9:24:76:16
   1   1   020100   --     N8   Online      FC F-Port 10:00:00:00:c9:29:06:4d
   2   2   020200   --     N8   No_Module   FC
   3   3   020300   --     N8   No_Module   FC
   4   4   020400   id     N8   Online      FC E-Port  10:00:00:05:33:5b:7b:c7
"B6510_LS1" (downstream)(Trunk master)
   5   5   020500   id     N8   Online      FC E-Port  (Trunk port, master is Port  4 )
   6   6   020800   id     N8   Online      FC E-Port  (Trunk port, master is Port  4 )
   7   7   020900   id     N8   Online      FC E-Port  (Trunk port, master is Port  4 )
```

Port Status

Use the `portshow` command to determin port status:

```
portshow [slotnumber/]portnumber

SW1:admin> portshow 2
portName: Bay1
portHealth: No Fabric Watch License

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x20b03      PRESENT ACTIVE F_PORT G_PORT LOGICAL_ONLINE LOGIN
NOELP ACCEPT FLOGI
portType:  11.0
POD Port: Port is licensed
portState: 1    Online
portPhys:  6    In_Sync
portScn:   32   F_Port
port generation number:    0
portId:    010100
portIfId:    4302000d
(truncated output)
```

Back Up Configuration

Use the `configupload` command to backup the configuration file on your switch[6].

1. Upload the Virtual Fabric (VF) configuration using the `configupload -vf` command.
2. Upload the default configuration file using the `configupload –all` command.

NOTE
If no VFs are defined, use only the `configupload –all` command

Because some configuration parameters require a reboot to take effect after you download a configuration file, you must reboot to be sure that the changed parameters are enabled. Before the reboot the changed parameter is listed in the configuration file but it is not effective until after the reboot. On dual CP platforms, you must reboot both CPs simultaneously for changes to take effect. To restore a configuration file back onto the switch:

1. Disable the switch.
2. Download the VF configuration using the `configdownload –vf` command.

NOTE
The order of the commands is very important: The VF section must be done first.

The switch auto reboots, even if VF was already enabled2

[6] You can use the `configUpload -vf` or `configDownload -vf` command to restore configurations to a logical switch. The `-vf` option only restores the Virtual Fabrics configuration information on to a switch of the same model.

3. Download the normal configuration file with `configdownload -all`.
4. If no VFs are defined use only the `configdownload -all` command.
5. These two commands must be run together to ensure compatibility.

Example syntax:

```
configupload (defaults to interactive mode)
configupload -p ftp | -ftp [<host>,<user>,<path>[,<passwd>]]
configupload -p scp | -scp [<host>,<user>,<path>]
configupload -all -p ftp | -ftp [<host>,<user>,<path>[,<passwd>]]
configupload -all -p scp | -scp [<host>,<user>,<path>]
configupload -fid # -p ftp | -ftp [<host>,<user>,<path>[,<passwd]]
configupload -fid # -p scp | -scp [<host>,<user>,<path>]
configupload -chassis -p ftp | -ftp [<host>,<user>,<path>[,passwd>]]
configupload -chassis -p scp | -scp [<host>,<user>,<path>]
configupload -switch -p ftp | -ftp [<host>,<user>,<path>[,passwd>]]
configupload -switch -p scp | -scp [<host>,<user>,<path>]
configupload [-force] -local|-USB|-U [<filename>]
```

If `<path>` is not specified then the config filename defaults to config.txt. Otherwise the specified path is used as in the following example.

`/usr/home/myconfig.txt` (config file is /usr/home/myconfig.txt)

`[no path supplied]` is the default path on the FTP server and the filename defaults to config.txt.

NOTE
Remote file may get overwritten if same filename is used

Firmware Download

The firmware download process is the same for all of Brocade's SAN switches. Firmware is stored in flash on two separate partitions, primary and secondary. Firmware is first downloaded to the secondary partition and the switch rebooted from the updated code. This allows an opportunity to assess the new firmware and ensure that there are no problems. Once the switch has booted successfully from the update code it is then copied to the remaining flash partition. When performing a firmware upgrade on any DCX switch the default is to upgrade both CPs.

You can use the Web Tools, Brocade Network Advisor, and the Command Line Interface (CLI) to upgrade the firmware on your switch or backbone.

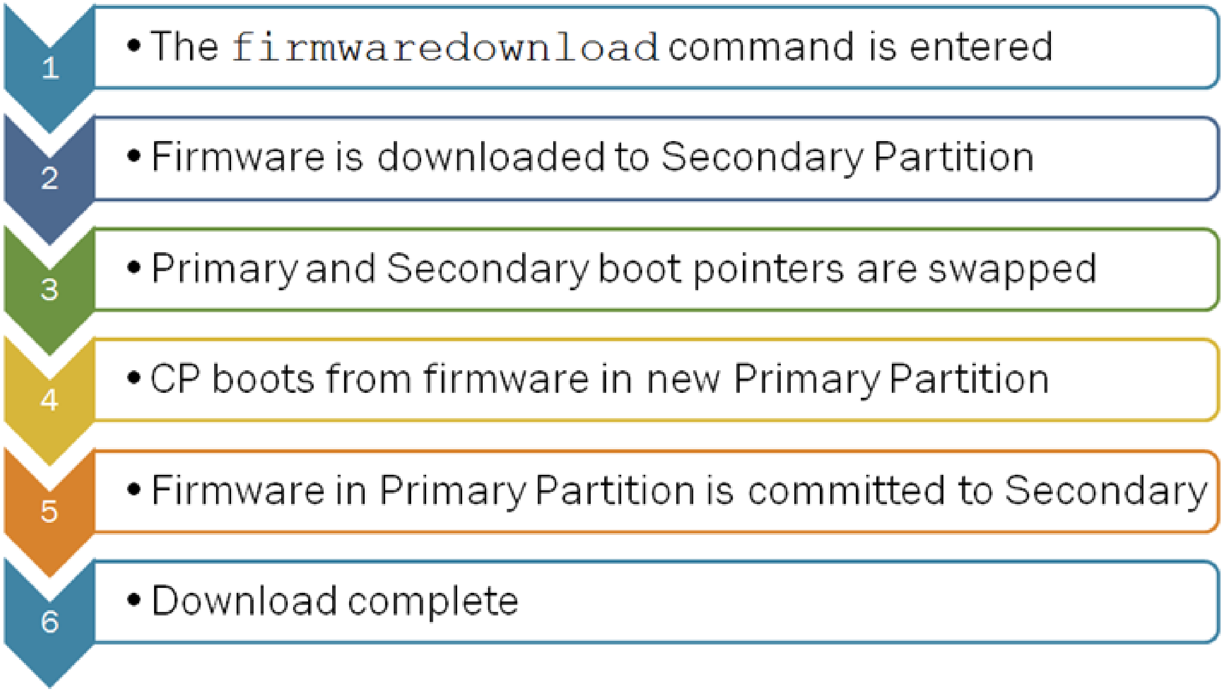
1
• The firmwaredownload command is entered
2
• Firmware is downloaded to Secondary Partition
3
• Primary and Secondary boot pointers are swapped
4
• CP boots from firmware in new Primary Partition
5
• Firmware in Primary Partition is committed to Secondary
6
• Download complete

Figure 12: Firmware Download Internal Process for SAN Switches

NOTE
Firmware upgrades that span multiple releases must be done one major release at a time, e.g. v6.4.x to v7.0.x

Firmware Download Requirements

1. Read the release notes for the new firmware to find out if there are any updates related to the firmware download process[7].
2. Connect to the switch and use the `firmwareshow` command to verify the current version of Fabric OS.
3. Brocade does not support upgrades from more than one previous release.

   For example, upgrading from Fabric OS v6.3.0 to v6.4.0 is supported, but upgrading from Fabric OS v6.2.0 or a previous release directly to v7.0.0 is not.

[7] If this switch was purchased from an OEM vendor there may be upgrade requirements different from those listed here. Contact your vendor for additional information.

In other words, upgrading a switch from Fabric OS v6.3.0 to v7.0.0 is a two-step process—first upgrade to v6.4.0, and then upgrade to v7.0.0. If you are running a pre-Fabric OS v6.2.0 version you must upgrade to v6.2.0, then to v6.3.0,then to v6.4.0, and finally to v7.0.0.

4. Perform a `configupload`.
5. *Optional*: For additional support, connect the switch to a computer with a serial console cable.

   Ensure that all serial consoles (both CPs for directors) and any open network connection sessions, such as Telnet, are logged and included with any trouble reports.
6. Connect to the switch and use the `supportsave` command to retrieve all current core files prior to executing the firmware download

   This helps to troubleshoot the firmware download process if a problem is encountered.
7. Optional: Enter the `errclear` command to erase all existing messages in addition to internal messages
8. If you are going to use FTP, verify that the service is running on the host.

Backbone Firmware Download Process

In addition to following the requirements stated above, the DCX family has the additional following requirements prior to a firmware download.

1. Run the `hashow` command and verify that HA is enabled and both CPs are in sync

```
dcx1:admin> hashow
Local CP (Slot 6, CP0): Active, Cold Recovered
Remote CP (Slot 7, CP1): Standby, Healthy
HA enabled, Heartbeat Up, HA State synchronized
```

2. Use the `firmwareshow` command to verify the version of Fabric OS currently running

```
dcx1:admin> firmwareshow -v
Slot Name     Appl     Primary/Secondary Versions               Status
------------------------------------------------------------------------
  6  CP0      Fabric OS      v7.2.0                             ACTIVE *
                             v7.2.0
              Co-Fabric OS   v7.2.0
                             v7.2.0
  7  CP1      Fabric OS      v7.2.0                              STANDBY
                             v7.2.0
              Co-Fabric OS   v7.2.0
                             v7.2.0
```

3. Verify the management interface from each CP has a network connection.
4. If you are going to use FTP, verify that the service is running on the host.

During the firmware download process, the CPs reboot switching between active and standby. This management interface keeps the active and standby attached to the network. If the management interfaces are not connected, then the firmware download fails.

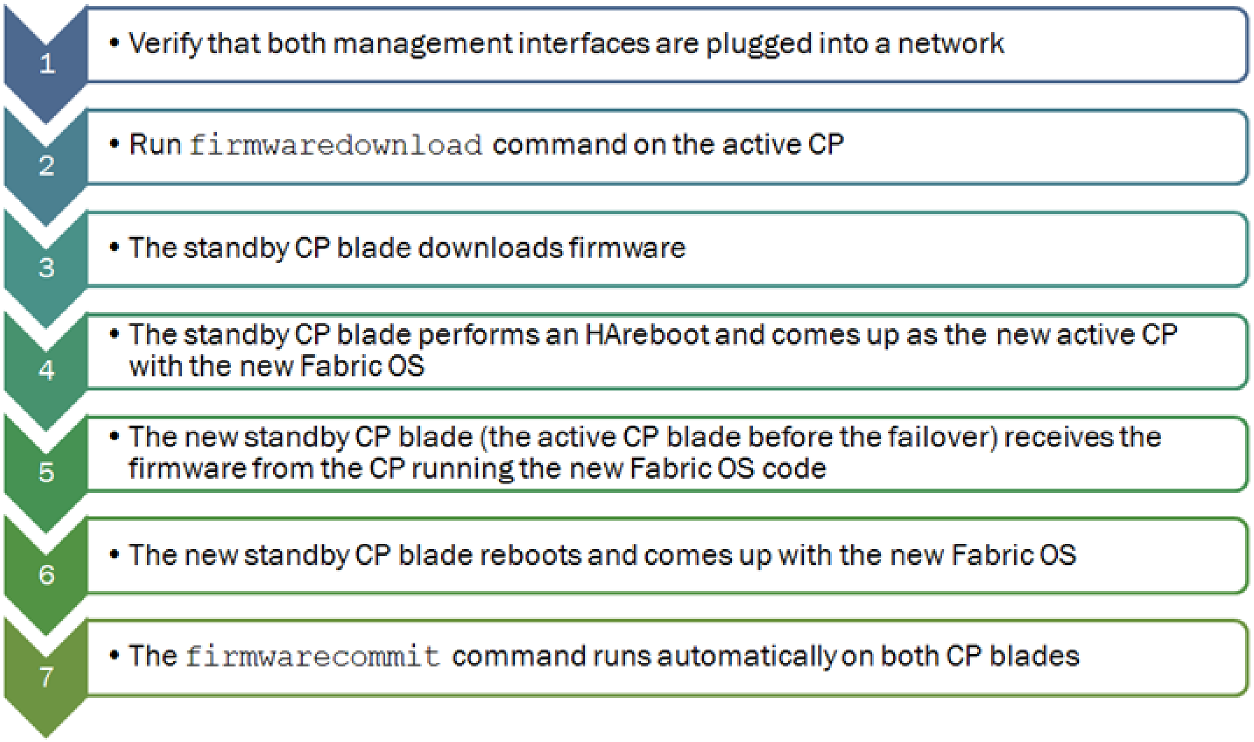

1
- Verify that both management interfaces are plugged into a network

2
- Run `firmwaredownload` command on the active CP

3
- The standby CP blade downloads firmware

4
- The standby CP blade performs an HAreboot and comes up as the new active CP with the new Fabric OS

5
- The new standby CP blade (the active CP blade before the failover) receives the firmware from the CP running the new Fabric OS code

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

Figure 13: DCX Family Firmware Upgrade Overview

USB Storage Device

All Brocade SAN switches support a firmware download from a Brocade branded USB device attached to the switch or active CP. Before the USB device can be accessed by the `firmwaredownload` command, it must be enabled and mounted as a file system. The firmware images to be downloaded must be stored under the relative path from `/usb/usbstorage/brocade/firmware` or use the absolute path in the USB file system. Multiple images can be stored under this directory. There is a firmwarekey directory where the public key signed firmware is stored.

When the `firmwareDownload` command line option, `-U` (upper case), is specified, the firmware download process downloads the specified firmware image from the USB device. When specifying a path to a firmware image in the USB device, you can only specify the relative path to /firmware or the absolute path. Only a Brocade branded USB device is supported and unsupported USB devices will fail with a `device not found` error message.

Brocade USB devices are pre-formatted with the required directory structure:

```
dcx1:admin> usbstorage -l
firmware\               0B       2011  Apr 25 13:54
config\                 0B       2011 Apr 25 13:54
support\                0B       2011 Apr 25 13:54
firmwarekey\            0B       2011 Apr 25 13:54
Available space on usbstorage 100%
```

Password Rules

- Password rules are enforced only when defining new passwords
- Passwords that have already been defined are not checked for policy compliance
- Set password rules using the `passwdcfg --set` command
- Set password strength policy by specifying the minimum number of:
    - Lowercase letters `-lowercase`
    - Uppercase letters `-uppercase`
    - Digits (0-9) `-digits`
    - Punctuation characters[8] `-punctuation`
    - Minimum length[9] `-minlength`
- Limit password re-use by setting the password history policy
    - Passwords kept in history[10] `-history`
- Avoid stale passwords by setting a password expiration policy[11]
    - Minimum age `-minpasswordage`
    - Maximum age `-maxpasswordage`
    - Expiration warning (days)[12] `-warning`
- Set the account lockout policy[13] `--enableadminlockout`
- Password failures allowed `-lockoutthreshold`
- Set lockout duration (minutes) `-lockoutduration`

[8] All printable punctuation characters except colon ":" are allowed.

[9] The minimum password length may be set from 8 to 40 characters in length. The password length is the total number of lowercase, uppercase, digits, and punctuation characters. The total number of these characters may not exceed 40. Keep this in mind as you specify the minimum number of each type of character required.

[10] The password history policy is not enforced when an administrator sets a password for another user, but the password set by the administrator is recorded in the user's password history.

[11] The password expiration policy is not enforced for root and factory accounts.

[12] The user will begin seeing warning messages when they login a number of days prior to password expiration. They will be compelled to change their password when it has expired.

[13] The account lockout policy is not enforced for root, factory, and admin role accounts

4 – FCP Concepts

After reviewing this section be sure you can perform the following:

- Demonstrate knowledge of trunking concepts
- Differentiate between fabric routing policies

Trunking

The trunking feature optimizes the use of bandwidth by allowing a group of links to merge into a single logical link, called a trunk group. Traffic is distributed dynamically and in-order over this trunk group, achieving greater performance with fewer links. Within the trunk group, multiple physical ports appear as a single logical port, thus simplifying management. Trunking also improves system reliability by maintaining in-order delivery of data and avoiding I/O retries if one link within the trunk group fails.

Trunking is frame-based instead of exchange-based. Since a frame is much smaller than an exchange, this means that frame-based trunks are more granular and better balanced than exchange-based trunks and provide maximum utilization of links.

Requirements

- A 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` command
- Trunk ports must operate at a common speed and long distance setting
- Trunk ports must originate and end in a valid port group
- Trunking port groups include: ports 0-7, 8-15, and so on
- When trunking criteria is met, the trunk forms automatically

The 2, 4, 8, 10, and 16 Gbps Trunking Overview

- Automatically aggregates up to 8 ISLs when the switches are connected[14]
  - Condor2 ASICs provide up to 64 Gbps of aggregate bandwidth
  - Condor3 ASICs provide up to 128 Gbps of aggregate bandwidth
- All ports in a trunk group must operate at the same speed
  - Condor2 ASICs support multiple 2/4/8 Gbps trunks between switches
  - Condor3 ASICs support multiple 2/4/8/10/16 Gbps trunks between same switches

[14] Automatically creates ISL trunks using from 2 to 8 ISLs when the switches are connected and all trunking requirements are met.

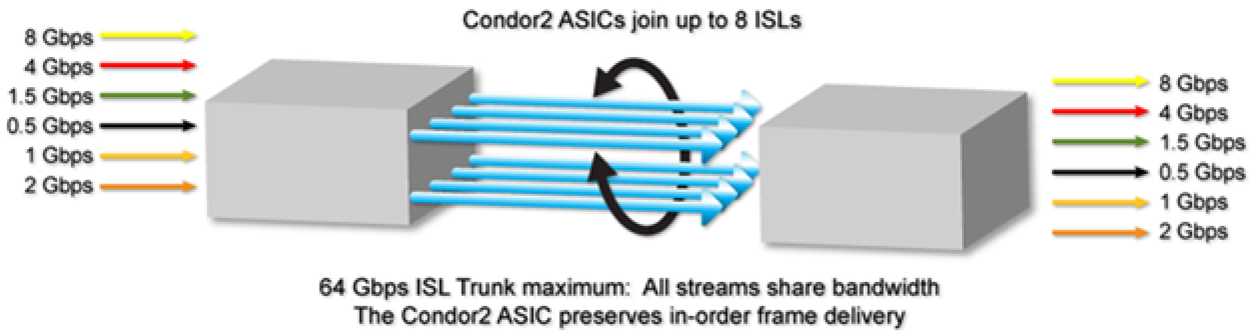
Condor2 ASICs join up to 8 ISLs
8 Gbps
4 Gbps
1.5 Gbps
0.5 Gbps
1 Gbps
2 Gbps
8 Gbps
4 Gbps
1.5 Gbps
0.5 Gbps
1 Gbps
2 Gbps
64 Gbps ISL Trunk maximum: All streams share bandwidth
The Condor2 ASIC preserves in-order frame delivery

Figure 14: Trunking

Masterless Trunking

The 4/8/16 Gbps ASICs use pseudo-trunking to prevent disruption when the trunk master is offline. The trunk master represents the group in the routing table. There is no build fabric when the trunk master goes offline. The `trunkshow` command displays the current master

```
SW1:admin> trunkshow
 1: 4 ->  8   10:00:00:05:1e:02:12:b1   deskew 15   MASTER
    0 ->  9   10:00:00:05:1e:02:12:b1   deskew 15
    1 ->  10  10:00:00:05:1e:02:12:b1   deskew 16
    5 ->  11  10:00:00:05:1e:02:12:b1   deskew 16
. . .<truncated output> . . .
```

When the MASTER is offline, `trunkshow` displays the new master

```
SW1:admin> trunkshow
 1: 1 -> 10   10:00:00:05:1e:02:12:b1   deskew 16   MASTER
    5 -> 11   10:00:00:05:1e:02:12:b1   deskew 16
    0 -> 9    10:00:00:05:1e:02:12:b1   deskew 15
```

The first ISL in the trunk to initialize is selected as the trunk master. The length of the cable is not a consideration when selecting the master.

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 ISL transmission capabilities. Differences in deskew can be caused by signal degradation which affects the transmission time of frames through the link. Differences in deskew can also be caused by excessive differences in cable length. Deskew values are displayed in the `trunkshow` command output.

The `trunkshow` Command

Use the `trunkshow` command to display trunking information of both E_Ports and EX_Ports.

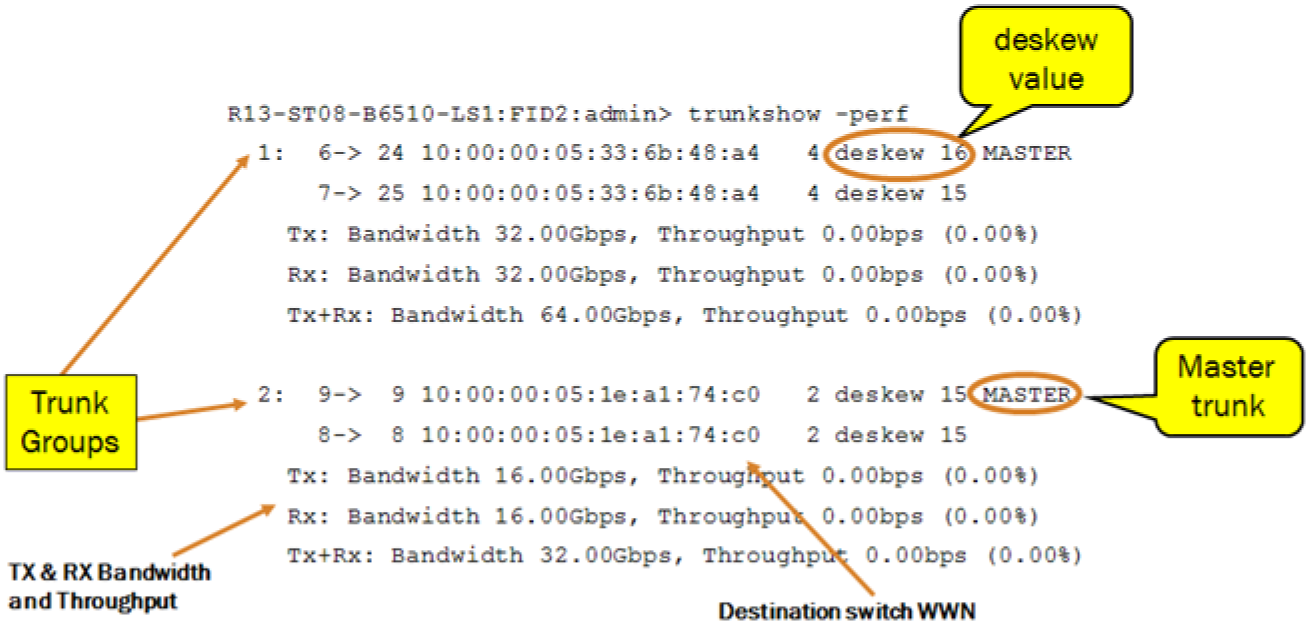

```
R13-ST08-B6510-LS1:FID2:admin> trunkshow -perf
  1:  6-> 24 10:00:00:05:33:6b:48:a4   4 deskew 16 MASTER
      7-> 25 10:00:00:05:33:6b:48:a4   4 deskew 15
    Tx: Bandwidth 32.00Gbps, Throughput 0.00bps (0.00%)
    Rx: Bandwidth 32.00Gbps, Throughput 0.00bps (0.00%)
    Tx+Rx: Bandwidth 64.00Gbps, Throughput 0.00bps (0.00%)

  2:  9->  9 10:00:00:05:1e:a1:74:c0   2 deskew 15 MASTER
      8->  8 10:00:00:05:1e:a1:74:c0   2 deskew 15
    Tx: Bandwidth 16.00Gbps, Throughput 0.00bps (0.00%)
    Rx: Bandwidth 16.00Gbps, Throughput 0.00bps (0.00%)
    Tx+Rx: Bandwidth 32.00Gbps, Throughput 0.00bps (0.00%)
```

Figure 15: Output from the `trunkshow` Command

ICL Trunking

ICL trunking is configured on an inter-chassis link (ICL) between two enterprise-class platforms and is applicable only to ports on the core blades.

ICL trunks automatically form on the ICLs when you install the Trunking license on each platform.

Trunk Monitoring

To monitor E_Port (ISL) and F_Port trunks, you can set monitors only on the master port of the trunk. If the master changes, the monitor automatically moves to the new master port. If a monitor is installed on a port that later becomes a slave port when a trunk comes up, the monitor automatically moves to the master port of the trunk.

For masterless trunking, if the master port goes offline, the new master acquires all the configurations and bottleneck history of the old master and continues with bottleneck detection on the trunk.

F_Port Trunking for Brocade Adapters

You can configure trunking between the F_Ports on an edge switch and the Brocade adapters. In addition to the trunk group requirements listed above, note the following requirements, which are specific to F_Port trunking for Brocade adapters:

- The edge switch must be running in Native mode. You cannot configure trunking between the Brocade adapters and the F_Ports of an Access Gateway module.
- A trunk are must be configured on the switch. You can configure only two F_Ports in one trunk group.

Routing Policies

Data moves through a fabric from switch to switch and from storage to server along one or more paths that make up a route. Each switch maintains its own routing policy and tables. Routing policies determine the path for each frame of data. The routing policy is unidirectional and responsible for selecting a route based on one of two user-configurable routing policies:

- Port-based routing

  The choice of routing path is based only on the incoming port and the destination domain. To optimize port-based routing, DLS round-robins the input ports across the available output ports to balance the load across the available output ports within a domain. Chosen routes are used until one of the devices in the fabric goes offline or the fabric changes.

  In the following `topologyshow` output, notice how the `In Ports` are in numerical order between the two paths. This tells you that the routing policy in place is port-based:

```
SW1:admin> topologyshow
2 domains in the fabric; Local Domain ID: 20
Domain:         10
Metric:         500
Name:           SW1
Path Count:     2
        Hops:                   1
        Out Port:               0
        In Ports:               8,10,12,14
        Total Bandwidth:        16.000 Gbps
        Bandwidth Demand:       400%
        Flags:                  D
2 domains in the fabric; Local Domain ID: 20
Domain:         10
Metric:         500
Name:           SW1
Path Count:     2
        Hops:                   1
        Out Port:               1
        In Ports:               9,11,13
        Total Bandwidth:        16.000 Gbps
        Bandwidth Demand:       300%
        Flags:                  D
```

- 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 path through the fabric. Exchange-based routing requires the use of the Dynamic Load Sharing (DLS) feature; when this policy is in effect, you cannot disable the DLS feature.

  Each switch has its own routing policy because different policies can exist in the same fabric. 8 and 16 Gbps ASICs use the FSPF protocol and either Port-based routing or Exchange-based routing, Exchange-based routing is Brocade's factory default setting.[15]

[15] Different OEMs may use different default settings. Please check with your switch vendor for settings.

In the following `topologyshow` output, notice how the `Out Ports` span more than one port and the `In Ports` are spread between these two `Out Ports`. This tells you that the routing policy in place is exchanged-based:

```
SW1:admin> topologyshow

2 domain(s) in the fabric; Local Domain ID: 3

Domain:2
Metric:500
Name:  R14-ST08-B6510
Path Count:  2

Hops:             1
Out Port:         8
In Ports:         14 15
Total Bandwidth: 8.000 Gbps
Bandwidth Demand:  200 %
Flags:            D

Hops:             1
Out Port:         9
In Ports:         14 15
Total Bandwidth: 8.000 Gbps
Bandwidth Demand:  200 %
Flags:            D
<truncated output>
```

Note

The switch must be disabled before changing the routing policies.

FSPF Link Cost

Fabric Shortest Path First (FSPF) detects link failures, determines the shortest route for traffic based on link cost, updates the routing table, provides fixed routing paths within a fabric, and maintains correct ordering of frames.

ISLs provide the physical pathway for routing frames from a Source ID (SID) to a Destination ID (DID) on a different domain (switch). When an ISL goes online or offline FSPF will update the routing tables to reflect the change.

As each host transmits a frame to the switch, the switch will read the SID and DID in the frame header. If the domain ID of the destination address is the same as the switch, intra-switch communication, the frame buffer is copied to the destination port and an R_RDY is sent to the host.

Since the SID and DID are in the first two words of the frame Brocade switches perform cut-through routing. The first two words of an incoming frame are read, if the DID is another port on the local domain the frame input is immediately transferred to the DID port. The entire frame does not need to be buffered in the switch; a frame may begin to emerge from the output port before it has been entirely received by the input port.

The routing table can be viewed using the `urouteshow` command. Static routes can be assigned using the `urouteconfig` command.

Example FSPF Path Cost

- The metric value is assigned to the transmit (Tx) side of each ISL port
- The Brocade default link cost (metric) value for 2, 4, 8, 10, and 16 Gbps links is 500
- Ports 2 and 5 have a cost of 500 from Domain 1 (D1) to Domain 3 (D3)
- Port 6 has a cost of 1000 from D1 to D3
- 500 (D1 to D2) + 500 (D2 to D3)
- Lowest cost paths are Ports 2 and 5
- FSPF configures the routing table in Domain 1 to only use the routes on ports 2 and 5 for frames with a destination of Domain 3
- If a switch is inserted in the path the metric cost changes

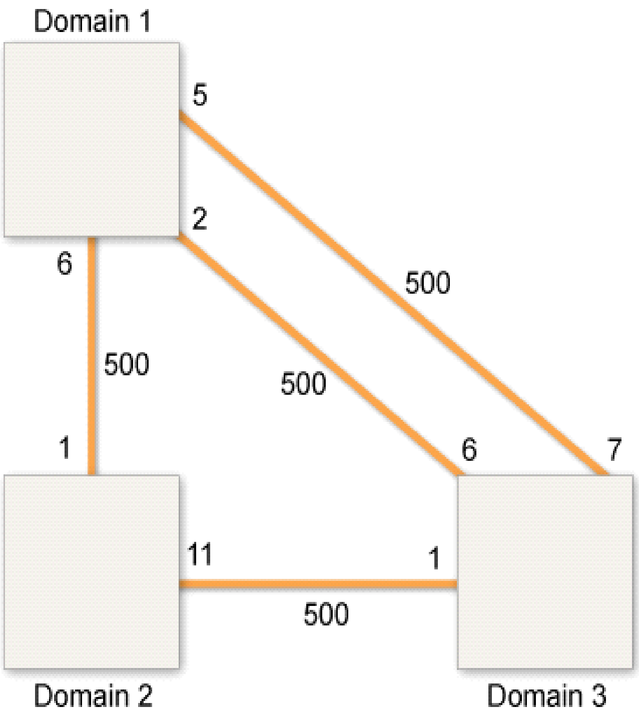

Between Domain 1 and Domain 3 in the figure above there are three paths: port 2 and port 5, each with a cost of 500, and port 6 with a cost of 1000. Only the lowest cost routes are in the routing table. In the figure above Domain 1 ports 2 and 5 would be in the routing table, port 6 would not.

Dynamic Load Sharing

The exchange-based routing policy depends on the Fabric OS Dynamic Load Sharing (DLS) feature for dynamic routing path selection. When using the exchange-based routing policy, DLS is enabled by default and cannot be disabled. In other words, you cannot enable or disable DLS when the exchange-based routing policy is in effect.

When the port-based policy is in force, you can enable DLS to optimize routing. When DLS is enabled, it shares traffic among multiple equivalent paths between switches. DLS recomputes load sharing when any of the following occurs:

- A switch boots up
- An E_Port goes offline and online
- An EX_Port goes offline
- A device goes offline

5 – Zoning

After reviewing this section be sure you can perform the following:

- Describe zoning concepts and implementation

Hierarchy of Zone Objects

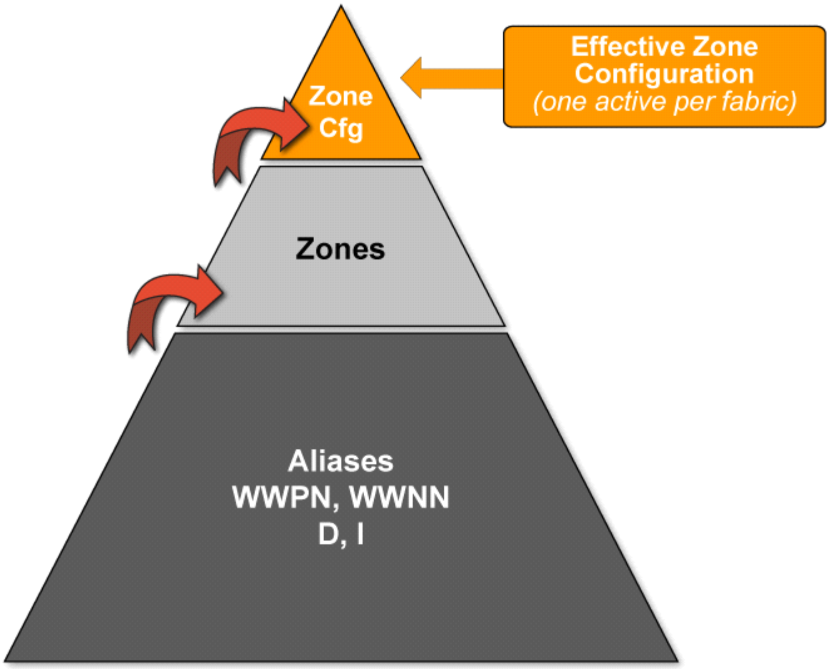

Member

- Alias is given a name, e.g. "Server_1", "Disk_Array_2".
- Physical Fabric port number or area number.
- Node World Wide Name - obtained using `nsshow` or `switchshow`.
- Port World Wide Name – obtained using `nsshow` or `portloginshow`.
- 64 characters maximum: A-Z, a-z, 0-9 and the "_" are allowed.

Zone

- Is given a name, e.g. "Red_Zone".
- Contains two or more members and uses a ";" as a separator.
- The same member can be in multiple zones.
- Zone definition is persistent; it remains until deleted or changed by an administrator.

Configuration[16]

- Is given a name, e.g. "Production_Cfg".
- Includes one or more zones.
- A configuration may be disabled or one configuration may be in effect from any switch in the fabric.
- An administrator selects which configuration is currently enabled.
- A configuration is saved when enabled and then distributed to the remaining switches in the fabric where it is enabled and saved.

[16] Also known as a Zone Configuration.

Zone Aliases

An alias is a name assigned to a device or group of devices. By creating an alias, you can assign a familiar name to a device, or you can group multiple devices into a single name. This can simplify cumbersome entries and it allows an intuitive naming structure such as using NT_Storage to define all NT storage ports in the fabric.

Zone aliases simplify repetitive entry of zone objects, such as PWWNs. The use of aliases is optional and involves the following components:

- Naming
  - Must begin with an alpha character
  - Can include numeric and underscore characters
  - Up to 64 characters

    Zone and configuration names are also limited to 64 characters maximum.
  - Case sensitive (DISK1 and Disk1 are unique names)
- Members
  - Domain, Index

    Zone objects identified by “Domain, Index” are specified as a pair of decimal numbers where “Domain” is the Domain ID of the switch and “Index” is the index number for the port on that switch.
  - Node World Wide Name - from `nsshow`
  - Port World Wide Name - from `nsshow`, `portloginshow` or `switchshow`

    Zone objects identified by World Wide Name (WWN, either node or port) are specified as a 16 digit hexadecimal number separated by colons, for example 10:00:00:90:69:00:00:8a. When a node name is used to specify a zone object, all ports on that device are in the zone. When a port name is used to specify a zone object, only that single port is in the zone.
- Sample naming convention

  For example, the name “Eng” could be used as an alias for PWWN: 10:00:00:80:33:3f:aa:11. Other possible examples are: Eng_Host1, Eng_Disk1, Eng_Disk2, Mkt_Host1, Mkt_Disk1, Mkt_Disk2 Zone_Eng, Zone_Mkt

Alias objects only appear in the defined configuration since they are used to assign a meaningful name to a device or group of devices

When a zoned host receives the list of network targets (referenced by domain,index, PWWN, or NWWN) from the Name Server, the host sends a PLOGI request to the destination addresses. If the PLOGI frame is allowed to pass and the target address replies with an accept to the PLOGI request, the switch and the zoning configuration have completed their responsibility of networking the source and destination.

NOTE
Limiting the number of LUNs and target IDs that the host can access is the responsibility of the management software being used at the storage end.

Default Zoning

In early versions of Fabric OS, when zoning was not implemented or the `cfgdisable` command was issued, all devices in the fabric could access each other

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 using the `defzone --allaccess` command (default setting)
- Can disable all device access using the `defzone --noaccess` command

The default zone setting is in effect when a user-specified zone configuration is not enabled and not in effect when a user-specified zone configuration is enabled.

The default zone feature can enable or disable device access within a fabric depending on the setting you specify (`--allaccess` or `--noaccess`). Default zones are based on the FC-GS standard.

The `defzone` command configures a default zone configuration and displays the current configuration. The command has no optional parameters, and takes one of three required arguments:

- `--allaccess`: Enables all device-to-device access within the fabric. This is the default behavior, and matches the default behavior in a non-zoned fabric.
- `--noaccess`: Create a default zone that disables all device-to-device access within the fabric.
- `--show`: Display the current default zone.

Names beginning with d__efault__ are reserved for default zoning and two underscores are used in each instance.

The setting of the default zone command is stored in the zoning transaction buffer. Normally, the `cfgsave` command is used to commit the zoning transaction to the entire fabric. Using the `cfgenable` or `cfgdisable` command initiates the commit since each command does an implied `cfgsave`. Because the setting is stored in the zoning transaction buffer, a `cfgtransabort` could be used to abort the `defzone` command.

Types of Zones

Zoning does more than define which devices can access each other. There are now several different types of zoning that can be implemented:

- "Normal" zones

  In a single fabric, "normal" zoning refers to defining zones with the purpose that only devices within the same zone can access each other. This was the initial purpose for zoning. Zones now exist for other purposes.

- LSAN zones

  An LSAN is a zone that spans routed fabrics. It is a logical storage area network that spans multiple physical fabrics, allowing devices in autonomous edge fabrics to communicate with each other. It defines device communication between autonomous fabrics but only allows designated devices in those fabrics to communicate and can be defined in edge fabrics and backbone fabrics.

- TI zones

  Traffic Isolation zones use a special zoning command, `zone`, and are intended to control the routing of frames between zone members, not to control access to devices (uses `zone --create` not `zonecreate` command).

  A normal zone must be in effect granting access between devices before a TI zone will be effective. TI zones will only appear in the defined zoning configuration, not in the effective zoning configuration and can only be created using D,I (domain,index) notation. TI zones must include E_Ports and F_Ports in order to create a complete, dedicated, end-to-end route from initiator to target and ports can only be members of a single TI zone.

- QoS zones

  QoS enables the setting of traffic priorities between specific hosts and targets. Prioritization is accomplished by the use of QoS zones, which will appear as normal zones. When creating a QoS zone they must be created using WWN notation and all normal zoning rules apply. To distinguish QoS zones from normal zones, special prefixes are used in the zone names:

  - QOSH_ to set high priority
  - QOSL_ to set low priority

  Default setting is medium priority and is used when no QoS zones are specified or when QoS cannot be enforced.

- Broadcast Zones

  Controls which devices receive broadcast frames. A broadcast zone restricts broadcast packets to only those devices that are members of the broadcast zone.

Dynamic Fabric Provisioning using FA-PWWN

Dynamic fabric provisioning simplifies and accelerates new server deployment and improves operational efficiency by using a fabric-assigned PWWN or FA-PWWN. An FA-PWWN is a “virtual” port WWN that can be used instead of the physical PWWN to create zoning and LUN mapping or masking. When the server is later attached to the SAN, the FA-PWWN is then assigned to the server.

The FA-PWWN feature allows you to do the following:

- Replace one server with another server, or replace failed HBAs or adapters within a server, without having to change any zoning or LUN mapping or masking configurations.
- Easily move servers across ports or Access Gateways by way of reassigning the FA-PWWN to another port.
- Use FA-PWWN to represent a server in boot LUN zone configurations so that any physical server that is mapped to this FA-PWWN can boot from that LUN, thus simplifying boot over SAN configuration.

For the server to use this feature, it must be using a Brocade HBA/Adapter with HBA driver version 3.0.0.0 or later. Some configuration of the HBA must be performed to use FA-PWWN.

Configurations

A zone configuration is a group of zones that is enforced whenever that zone configuration is enabled. A zone can be included in more than one zone configuration.

To define a zone configuration, specify the list of zones to be included and assign a zone configuration name. Zoning may be disabled at any time. When a zone configuration is in effect, all zones that are members of that configuration are in effect.

- Defined configuration: The complete set of all zone objects that have been defined in the fabric.
- Effective configuration: A single zone configuration that is currently in effect. The effective configuration is built when an administrator enables a specified zone configuration. This configuration is “compiled” by checking for undefined zone names, or zone alias names, or other issues.
- Saved configuration: A copy of the defined configuration plus the name of the effective configuration which is saved in flash memory by the cfgsave command. There may be differences between the saved configuration and the defined configuration if the system administrator has modified any of the zone definitions and has not saved them.

Enabling Zoning

Use the `cfgenable` command to enable a zone configuration. The specified zone configuration is built by checking for undefined zone names, zone alias names, or other inconsistencies by expanding zone aliases, removing duplicate entries, and then installing the current configuration.

If the build fails, the previous state is preserved (zoning remains disabled, or the previous configuration remains in effect). If the build succeeds, the new configuration replaces the previous configuration.

Disabling Zoning

Use the `cfgdisable` command to disable the current zone configuration. The fabric returns to non-zoning mode, in which all devices see each other.

Clearing Zoning

Use the `cfgclear` command to clear all zone information in the defined configuration. All defined zone objects are deleted. After using the `cfgclear` command, use the `cfgsave` command to commit the defined and effective configuration to flash memory for all the switches in the fabric. To completely clear the zoning database, use the commands in the following order:

1. Enter the `cfgdisable` command.
2. Enter the `cfgclear` command.
3. Enter the `cfgsave` command.

Saving Zoning

Using the cfgsave command saves the current zone configuration. When the current zone configuration is saved, the defined configuration and the name of the effective configuration are written to flash memory in all switches in the fabric. Prior to being saved the changes made are stored in a transaction buffer on the local switch and will be lost if the switch is rebooted.

Because the saved configuration is reloaded at power on, only valid configurations are saved. Saving the configuration verifies that the enabled configuration is valid by performing the same tests as enabling the configuration. If the tests fail, an error is displayed and the configuration is not saved. Tests might fail if a configuration has been modified since it was last enabled.

This command ends and commits the current transaction. If a transaction is open on a different switch in the fabric when this command is run, the transaction on the other switch is automatically aborted and a message is displayed on the other switches.

If the defined configuration is larger than the supported maximum zoning database size, the following message is issued:

```
Commit zone DB larger than supported - <zone db size> greater than <max zone
db size>
```

Zoning Transactions

When changes are made to the zone configuration on a switch they are stored in a transaction buffer until the configuration is either saved or aborted. You can use the cfgtransshow command to view the transaction token (if there is one) and cfgtransabort to abort any changes.

Other Zoning Commands

A full list of zoning commands can be displayed on any Fabric OS switch using the zonehelp command.

- aliadd - Add a member to a zone alias
- alicreate - Create a zone alias
- alidelete - Delete a zone alias
- aliremove - Remove a member from a zone alias
- alishow - Print zone alias information
- bootluncfg - Configure boot LUN for an HBA
- cfgactvshow - Display Effective zone configuration information
- cfgadd - Add a member to a configuration
- cfgclear - Clear all zone configurations
- cfgcreate - Create a zone configuration
- cfgdelete - Delete a zone configuration
- cfgdisable - Disable a zone configuration
- cfgenable - Enable a zone configuration
- cfgremove - Remove a member from a configuration
- cfgsave - Save zone configurations in flash
- cfgshow - Print zone configuration information

- `cfgsize` - Print size details of zone database
- `cfgtransabort` - Abort zone configuration transaction
- `cfgtransshow` - Print zone configurations in transaction buffer
- `defzone` - Activates or deactivates a default zone configuration
- `msfr` - Create a MSFR Zone
- `nsdevlog` - Manage Name Server device logs
- `nszonemember` - Display the information of all the online devices which are zoned with the given device
- `openfr` - Create a MSFR Zone
- `zone` - Configure zone objects
- `zoneadd` - Add a member to a zone
- `zonecreate` - Create a zone
- `zonedelete` - Delete a zone
- `zonehelp` - Print zoning help info
- `zoneobjectcopy` - Copies a zone object
- `zoneobjectexpunge` - Expunges a zone object
- `zoneobjectrename` - Rename a zoning Object
- `zoneobjectreplace` - Replace a zoning Object
- `zoneremove` - Remove a member from a zone
- `zoneshow` - Print zone information

Types of Zoning Enforcement

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

- Session-based hardware enforcement

  A session enforced zone is a zoning protection that guarantees that only members of the zone can complete PLOGI/ADISC/PDISC which prevents any unauthorized access by devices that are not a member of the zone. Enforcement to a zone with WWN members and domain, index will change from hardware to session enforcement. The ASIC will perform authentication using the name server to compare the SID/DID in the primitive commands with the current zone configuration. If the current zone configuration does not permit the devices to communicate, the switch issues a reject to the SID, effectively blocking communications.

  - Name Server restricts PLOGIs which means that devices that are session enforced cause any PLOGIs to the device to be rejected.

- Frame-based hardware enforcement

  Frame-based hardware enforced zoning (also known as hardware enforced zoning) is used by zones with all members defined by their domain ID, index or all members defined by their WWN. This the strongest form of enforcement and will block all frames that compromise the zone from a device that is not a member of a zone. Destination ASIC checks SID on every frame against CAM table entries. Overlapping zones (zone members that appear in two or more zones) are permitted and hardware enforcement will continue as long as the overlapping zones have either all WWNs or domain ID, index entries.

Using all WWNs in a zone allows for the node to attach to any port in the fabric and have hardware enforcement. Using all domain ID, index members restricts the movement of devices in the fabric until a zone update is made.

- 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 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.

Figure 16 on page 51 illustrates the results of Hardware and Session enforced overlapping zones.

The Blue Zone is Hardware enforced because all devices have been specified by domain,port. The Green Zone is also Hardware enforced because all devices have been specified by WWNs.

The Red Zone is Session enforced because a mix of domain,port and WWNs have been specified in the zone. The Orange Zone is also Session enforced because of a mix of ports and WWNs in the same zone.

The Green zone is defined with all WWNs (WWN1, WWN2 and WWN5) and meets the rules for Hardware enforcement. The Red zone is defined with a mix of port and WWNs and meets the rules for Session enforcement.

The target device WWN1 is defined in both the Red and Green zones. When a device is defined in overlapping zones, where one is Hardware enforced and the other is Session enforced, the device becomes Session enforced in all zones. What is important to note is the host (WWN2) is still Hardware enforced, even though the target device (WWN1) is now Session enforced. Under these conditions, zoning enforcement is determined at the device level, not the zone level.

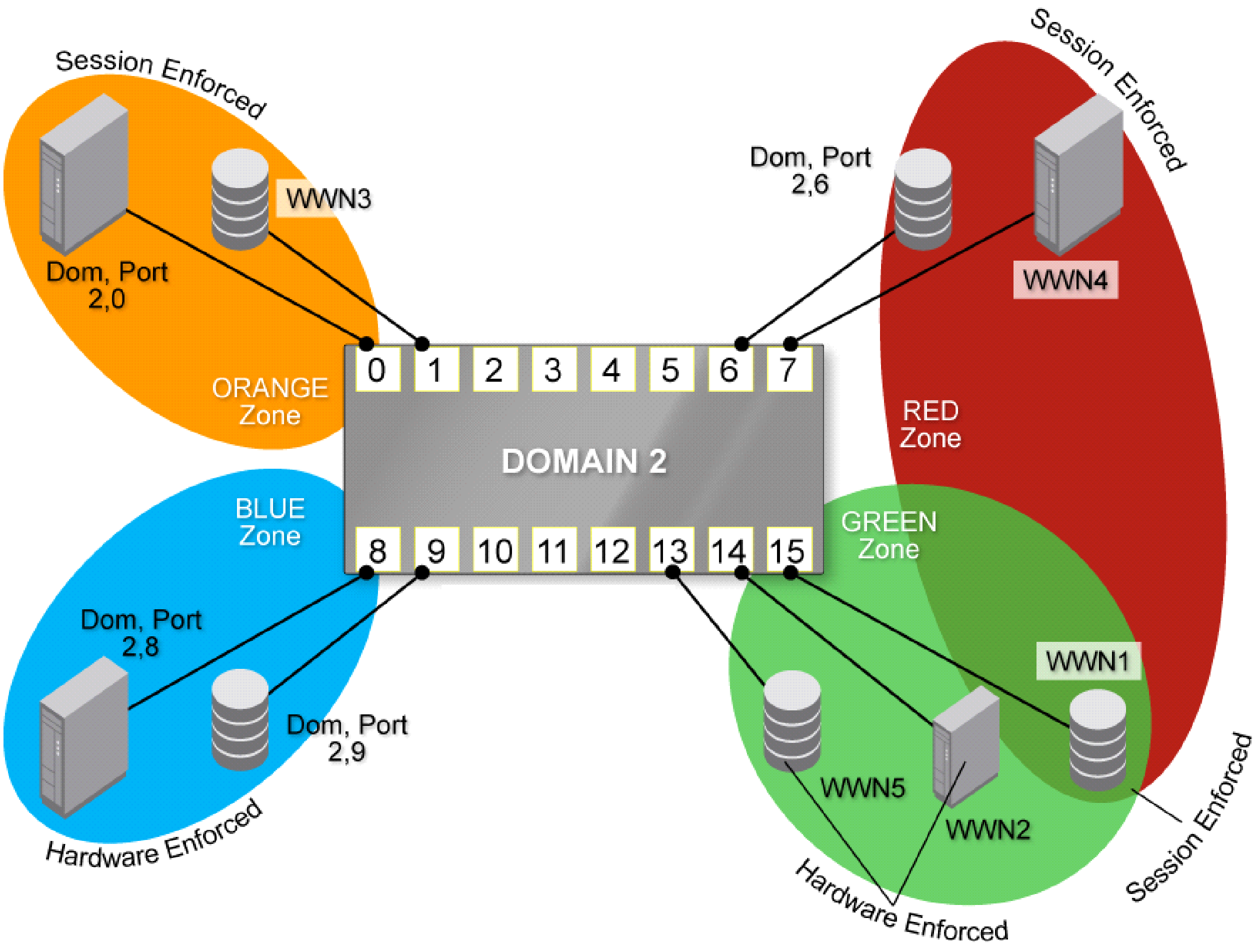

Figure 16: Zoning Enforcement

Note

The Red and Orange Zones also illustrate that the type of device (initiator vs. target) has no bearing on the type of enforcement.

6 – Basic Management

After reviewing this section be sure you can perform the following:

- Identify basic switch management interfaces
- Demonstrate basic knowledge of management and reporting tools

IP Filter Policies

The IP Filter policy is a set of rules applied to the IP management interfaces as a packet filtering firewall. The firewall permits or denies the traffic to go through the IP management interfaces according to the policy rules. IP packets that are not in the rules are denied.

Example: To limit the permitted management stations.

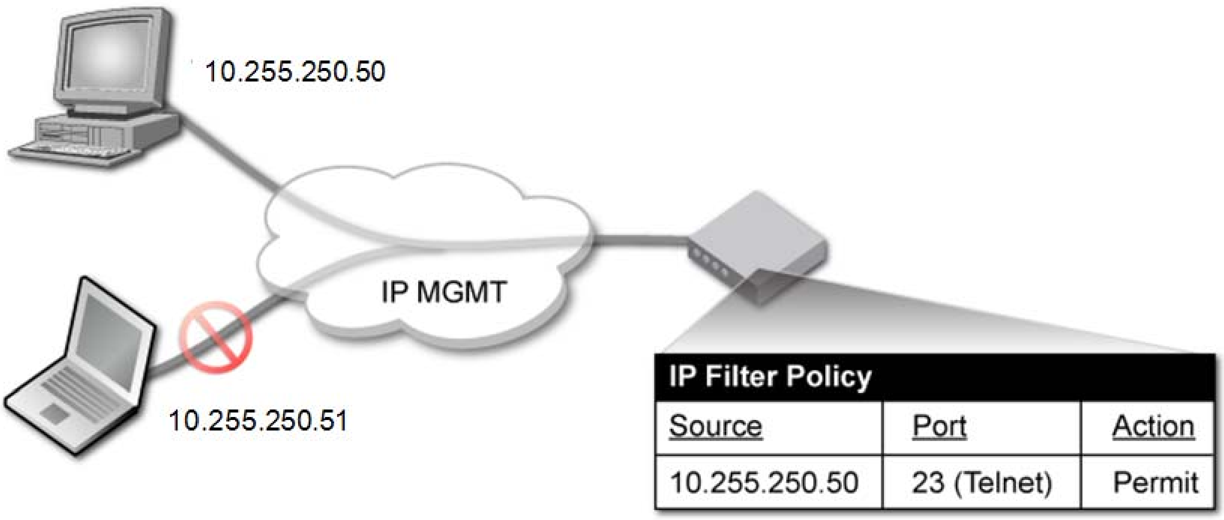

IP Filter Policy

Source	Port	Action
10.255.250.50	23 (Telnet)	Permit

Figure 17: IP Filter Policy Example

Use the `ipfilter --create` command to create an IP Filter rule and the `ipfilter --delete` command to delete an IP Filter rule. When deleting an IP Filter rule be sure to connect to the switch from the serial port or other port that will not be effected by the rule being deleted.

Protocol Overview

Security protocols provide endpoint authentication and communications privacy using cryptography. Typically, you are authenticated to the switch while the switch remains unauthenticated to you. This means that you can be sure with what you are communicating. The next level of security, in which both ends of the conversation are sure with whom they are communicating, is known as two-factor authentication. Two-factor authentication requires public key infrastructure (PKI) deployment to clients.

TABLE 7 Secure protocol support

Protocol	Description
HTTPS	HTTPS is a Uniform Resource Identifier scheme used to indicate a secure HTTP connection. Web Tools supports the use of hypertext transfer protocol over secure socket layer (HTTPS).
IPsec	Internet Protocol Security (IPsec) is a framework of open standards for providing confidentiality, authentication and integrity for IP data transmitted over untrusted links or networks.
LDAPS	Lightweight Directory Access Protocol over SSL uses a certificate authority (CA). By default, LDAP traffic is transmitted unsecured. You can make LDAP traffic confidential and secure by using Secure Sockets Layer (SSL) / Transport Layer Security (TLS) technology in conjunction with LDAP.
SCP	Secure Copy (SCP) is a means of securely transferring computer files between a local and a remote host or between two remote hosts, using the Secure Shell (SSH) protocol. This protocol can be used in downloading firmware, uploading/downloading configuration files, and installing certificates.
SNMP	SNMP is used in network management systems to monitor network-attached devices for conditions that warrant administrative attention. Supports SNMPv1, v2, and v3.
SSH	Secure Shell (SSH) is a network protocol that allows data to be exchanged over a secure channel between two computers. Encryption provides confidentiality and integrity of data. SSH uses public-key cryptography to authenticate the remote computer and allow the remote computer to authenticate the user, if necessary.
SSL	Fabric OS uses secure socket layer (SSL) to support HTTPS. A certificate must be generated and installed on each switch to enable SSL. Supports SSLv3, 128-bit encryption by default.

Secure Sockets Layer Protocol

Secure sockets layer (SSL) protocol provides secure access to a fabric through Web-based management tools like Web Tools. SSL support is a standard Fabric OS feature.

Switches configured for SSL grant access to management tools through hypertext transfer protocol over SSL links (which begin with https://) instead of standard links (which begin with http://).

SSL uses public key infrastructure (PKI) encryption to protect data transferred over SSL connections. PKI is based on digital certificates obtained from an Internet Certificate Authority (CA) that acts as the trusted key agent.

Certificates are based on the switch IP address or fully qualified domain name (FQDN), depending on the issuing CA. If you change a switch IP address or FQDN after activating an associated certificate, you may have to obtain and install a new certificate. Check with the CA to verify this possibility, and plan these types of changes accordingly.

Secure Copy

The secure copy protocol (SCP) runs on port 22. It encrypts data during transfer, thereby avoiding packet sniffers that attempt to extract useful information during data transfer. SCP relies on SSH to provide authentication and security.

Secure Shell protocol

To ensure security, Fabric OS supports secure shell (SSH) encrypted sessions. SSH encrypts all messages, including the client transmission of the password during login. The SSH package contains a daemon (sshd), which runs on the switch. The daemon supports a wide variety of encryption algorithms, such as Blowfish-Cipher block chaining (CBC) and Advanced Encryption Standard (AES).

NOTE
To maintain a secure network, you should avoid using Telnet or any other unprotected application when you are working on the switch.

The File Transfer Protocol (FTP) is also not secure. When you use FTP to copy files to or from the switch, the contents are in clear text. This includes the remote FTP server's login and password. You can use SFTP to work around this issue.

Commands that require a secure login channel must originate from an SSH session. Starting an SSH session and then using the login command to start a nested SSH session results in the rejecting of commands that require a secure channel.

Brocade Network Advisor

Brocade Network Advisor is the industry’s first unified network management platform for SAN and IP networks. It provides a single, easy-to-use solution across Brocade Fibre Channel SANs, Layer 2/3 IP networks, Layer 4-7 application delivery networks, Fibre Channel over Ethernet (FCoE) networks, wireless networks, and Multiprotocol Label Switching (MPLS) networks. Brocade Network Advisor is a comprehensive tool for configuring, managing, monitoring, and reporting on Brocade data center, enterprise campus, and service provider networks with robust Role-Based Access Control (RBAC), automation, operational simplicity, and end-to-end network visibility. In addition, it integrates seamlessly with industry-leading orchestration products from Brocade Partners to provide a best-in-class solution.

TABLE 8 Brocade Network Advisor Key Features

Feature	Description
Call Home Support	Automatically collects system information and sends notifications for faster fault diagnosis, isolation, and remote support operations.
Change Management	Tracks device configuration changes and image management, which enables viewing, retrieving, and restoring configuration files.
CLI Configuration Wizard	Configure and deploy CLI-based templates across one or more IP devices.
Configuration Management	Easy-to-use Device Configuration wizard to configure groups of devices.
Event Management	Helps in troubleshooting network-related issues and captures SNMP traps and Syslog event messages for reporting, analysis, monitoring, and remediation.
Intelligent Dashboard	At-a-glance summary of all discovered Brocade and third-party IP devices.
Partner Integration	Seamless integration with Microsoft SCOM, VMware vCenter, IBM Tivoli, and IBM Systems Director.
Performance Monitoring	Real-time network monitoring and accounting capabilities for all IP switches and routers; Detailed SAN monitoring of port and link utilization.

TABLE 8 **Brocade Network Advisor Key Features (Continued)**

Feature	Description
Security Management	Robust security administration by integrating with leading solutions such as RADIUS, Active Directory, and TACACS+.
User Management	Flexible definitions of administrator roles and responsibilities with RBAC for SAN and IP management.

MAPS and Flow Vision

Fabric Vision License

Enables MAPS (Monitoring and Alerting Policy Suite), Flow Vision, and D_Port to non-Brocade devices. MAPS enables rules-based monitoring and alerting capabilities, provides comprehensive dashboards to quickly troubleshoot problems in Brocade SAN environments. Flow Vision enables host to LUN flow monitoring, application flow mirroring for offline capture and deeper analysis, and test traffic flow generation function for SAN infrastructure validation. D_Port to non-Brocade devices allows extensive diagnostic testing of links to devices other than Brocade switches and adapters. (Functionality requires support by attached device, availability TBD).

Fabric Vision license also enables Fabric Watch and Advanced Performance Monitoring functionalities without requiring Brocade Fabric Watch or Brocade Advanced Performance Monitoring license (with FOS v7.2 and later only).

- Fabric Vision capability is also available if you have both the Fabric Watch and Performance Monitor licenses installed

MAPS

The Monitoring and Alerting Policy Suite (MAPS) is an optional storage area network (SAN) health monitor supported on all switches running Fabric OS 7.2.0 or later that allows you to enable each switch to constantly monitor itself for potential faults and automatically alerts you to problems before they become costly failures.

MAPS tracks a variety of SAN fabric metrics and events. Monitoring fabric-wide events, ports, and environmental parameters enables early fault detection and isolation as well as performance measurements.

- Mutually exclusive with Fabric Watch, when MAPS is enabled on a switch Fabric Watch and all monitors are automatically disabled
- Can send email, SNMP, and RASLog alerts; fence a port, or change the switch status
- Can be managed using CLI or Network Advisor
- Network Advisor can be used to distribute rules across an entire fabric

Flow Vision

Flow Vision is a Fibre Channel SAN network diagnostic tool supported on all platforms supported by Fabric OS 7.2 and later, that provides you with a comprehensive vision of fabric traffic flows and with the ability to non-disruptively create and capture copies of traffic flows for later analysis. Flow Vision also provides a test flow generation capability that you can use to pre-test a SAN infrastructure for robustness. This test flow generation capability is also useful for testing the internal connections on a switch before deploying the switch into a production environment. You cannot run Flow Vision and Advanced Performance Monitor (APM), or Port Mirroring at the same time on a chassis (across logical switches).

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, and voltage), Security, resource and FRU
- Port status: Monitors F and 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. It tracks the number of occurrences of each condition and each counter is compared with an upper boundary and lower boundary.

The `thMonitor` Command

The `thmonitor` command enables Fabric Watch threshold monitoring for 10 and 16 Gbps SFPs and for 16 Gbps QSFPs. The monitoring of 10 and 16 Gbps SFPs and 16 Gbps QSFPs must be explicitly enabled. Fabric Watch does not monitor these components by default. It takes precedence over pause and continue configuration for advanced SFP thresholds.

Example of **`thmonitor`** command

```
Switch1:admin> thMonitor --enable brcdsfp
Brcd SFP Threshold Monitoring is enabled

Switch1:admin> thMonitor --show
Brcd SFP Threshold Monitoring is enabled
```

Host Connectivity Manager

Host Connectivity Manager (HCM) is an HBA management tool that is loaded on the host containing the HBAs. HCM has the following capabilities:

- Collect event logs
- Manage HBA firmware upgrades
- Access the Boot BIOS
- View port statistics including error statistics
- Manage remote hosts (import HBAs from other hosts)
- Perform diagnostics
- Configure device persistence

SAN Health

Brocade SAN Health is a free utility that can be used to take inventory, diagram, and collect information about your SANs.

- Taking inventory of devices, switches, firmware versions, and SAN fabrics
- Capturing and displaying historical performance data
- Comparing zoning and switch configurations against best practices
- Assessing performance statistics and error conditions
- Producing detailed graphical reports and diagrams

7 – Basic Troubleshooting

After reviewing this section be sure you can perform the following:

- Demonstrate knowledge of how to collect information and perform basic troubleshooting under different scenarios

Switch and Backbone `supportsave`

Run Fabric OS `supportsave` as soon as you experience a problem. Run `supportsave` prior to beginning any troubleshooting because critical problem determination data may be captured if `supportsave` is run right away. In some cases, you may need to duplicate the problem and then capture a `supportsave`.

- If unable to resolve then run `supportsave` again for recent activity.
- If problem escalation is required, send the escalation team all relevant supportsave files.

If you maintain your `supportsave` then it can give you a baseline when troubleshooting. You need to maintain a baseline `supportsave` and take it after every configuration change.

Brocade Network Advisor can do this automatically or you can use the `supportftp` command to enable or disable auto file transfer.

The `supportsave` command supports the following protocols for transferring files from the switch:

- FTP
- SCP
- SFTP

Adapter supportSave

The supportsave data can be gathered for a Brocade HBA, CNA or FA. The `supportsave` utility is supported for Brocade adapters only and found in:

- Brocade Network Advisor (host must be discovered)
- Host Connectivity Manager (HCM)
- Brocade Command Utility (BCU) script (using a command string)

Bottleneck Detection

A *congestion* bottleneck is a port that is unable to transmit frames at the offered rate because the offered rate is greater than the physical data rate of the line.

Example: Attempt to transfer 16 Gbps of data on an 8 Gbps ISL

A *latency* bottleneck is defined as a port which is unable to transmit frames at the offered rate because credits are not returned fast enough from the attached device (receiver).

Slow drain spreads into the fabric and can slow down unrelated flows in the fabric.

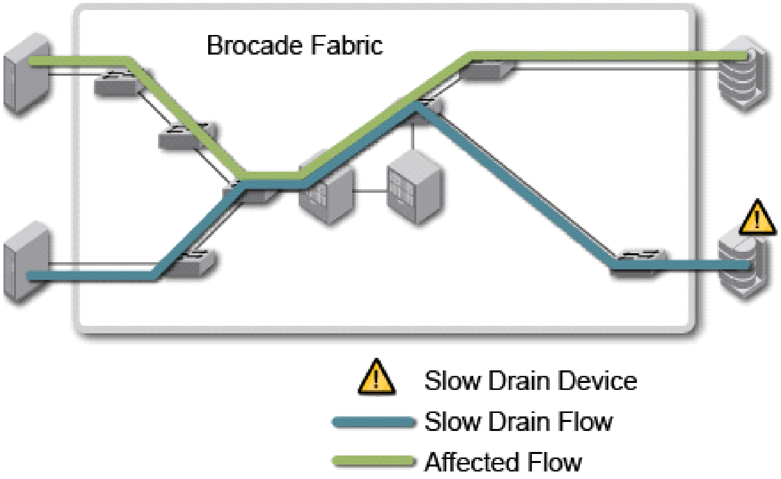

Figure 18: Bottlenecks

Use the `bottleneckmon` command to detect latency and congestion bottlenecks on F[L]_Ports and E_Ports. The `bottleneckmon` command detects slow drain devices (F_Ports) at the egress side of the port.

NOTE
Not recommended on links that are normally above 85% utilization. No license is required.

ClearLink Diagnostic Port

Starting with Fabric OS v7.0.0, you can convert a Fibre Channel port, including ISLs and loopback ports, into a Diagnostic Port (D_Port). This port lets you isolate the inter-switch link (ISL) to diagnose link level faults. The D_Port does not carry any fabric traffic, and is designated to run only specific diagnostics tests on it. The creation of a D_Port is subject to Virtual Fabric restrictions that may be in place. The ports must be 10G or 16G Brocade-branded SFPs.

You must configure both ends of the link between a given pair of switches, and you must disable the port before you can configure a D_Port. Re-enabling the D_Ports automatically starts the diagnostics when the ports come online, and includes the following tests:

- Electrical loopback (16G SFPs only)
- Optical loopback (16G SFPs only)
- Link traffic (16G SFPs and 10G SFPs)
- Link latency and distance measurement (16G SFPs and 10G SFPs)

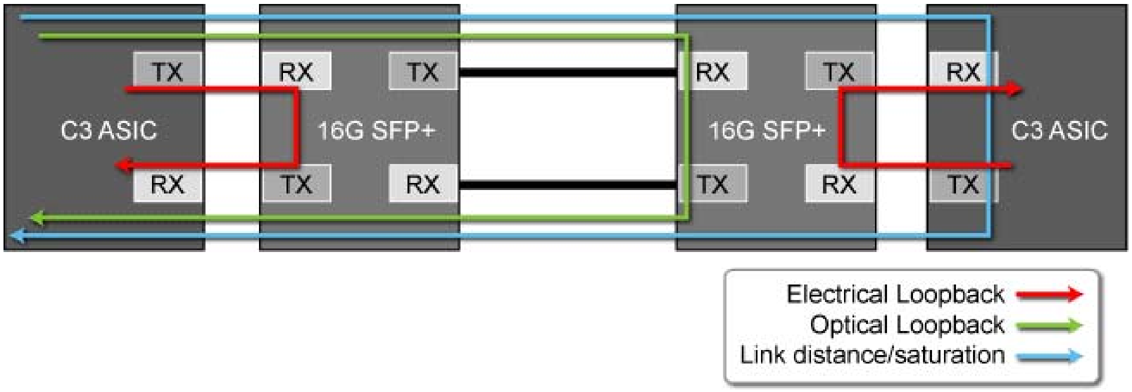

Figure 19: D_Port Tests

Use the D_Port tests for the following situations:

- Testing a new ISL before adding it to the fabric
- Testing a trunk member before joining it with the trunk
- Testing long distance cables and SFPs
- Testing loopback ports

D_Port configuration is not supported on inter-chassis link (ICL) ports, EX_Ports, F_Ports, or N_Ports (both in Access Gateway and Host Bus Adapters). D_Port configuration fails if the port is configured in R_RDY mode, encryption mode, or compression mode. Links with D_Port configuration mismatch become segmented or disabled.

When large number of D_Ports are configured:

- The test is run on one port per blade at a time
- Other ports wait until the test completes
- No tests begin until the fabric is stable

Zoning

Analyze zones to verify correct devices are communicating using the following tools:

- Brocade Network Advisor
- `nszonemember`
- `fcping`
- `zoneshow`
- SAN Health

Segmented fabrics

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

- Incompatible fabric parameters
- Incompatible zoning configuration
- Domain ID conflict
- Fabric ID conflict
- 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.

The following fabric parameters must be identical on each switch for a fabric to merge:

- R_A_TOV
- E_D_TOV
- Data field size
- Sequence level switching
- Disable device probing
- Suppress class F traffic
- Per-frame route priority
- Long distance fabric

Zones and Fabric Segmentations

When a new switch is added to the fabric it automatically takes on the zoning configuration in formation from the fabric. When merging fabrics that already have existing zone configurations several rules must be followed to avoid segmentations.

- Local and adjacent configurations: If the local and adjacent zone database configurations are the same, they will remain unchanged after the merge.
- Effective configurations: If there is an effective configuration between two switches, the effective zone configurations must match.
- Zone object naming: If a zoning object has the same name in both the local and adjacent defined configurations, the object types and member lists must match. When comparing member lists, the content and order of the members are important.
- Objects in adjacent configurations: If a zoning object appears in an adjacent defined configuration, but not in the local defined configuration, the zoning object is added to the local defined configuration. The modified zone database must fit in the nonvolatile memory area allotted for the zone database.
- Local configuration modification: If a local defined configuration is modified because of a merge, the new zone database is propagated to other the switches within the merge request.
- TI zones: If there is an effective configuration between two switches and TI zones are present on either switch, the TI zones are not automatically activated after the merge. Check the TI zone enabled status using the `zone --show` command, and if the status does not match across switches, issue the `cfgenable` command.

Port Initialization Problems

Running `switchshow` on Switch2 shows the storage failed to log into the fabric. Notice that in Figure 20 port 14 indicates an initialization problem.

The G_Port being online indicates a problem. The device connected to that port has a good link (it shows `Online`) but did not successfully get far enough into the process to become either an E_Port or an F_Port. If the device did not come up as a G_Port and was still not physically connected, it would come up with one of the following port states:

- `No_Light` (not receiving)
- `No_Sync` (not synchronizing)
- `In_Sync` (receiving light and in synchronization but unable to go further in initialization process)
- `Laser_Flt`
- `Port_Flt`
- `Diag_Flt` (diagnostics failed during boot up process)
- `Testing` (which would explain why you do not see the device)

You want to see the port in the `Online` state.

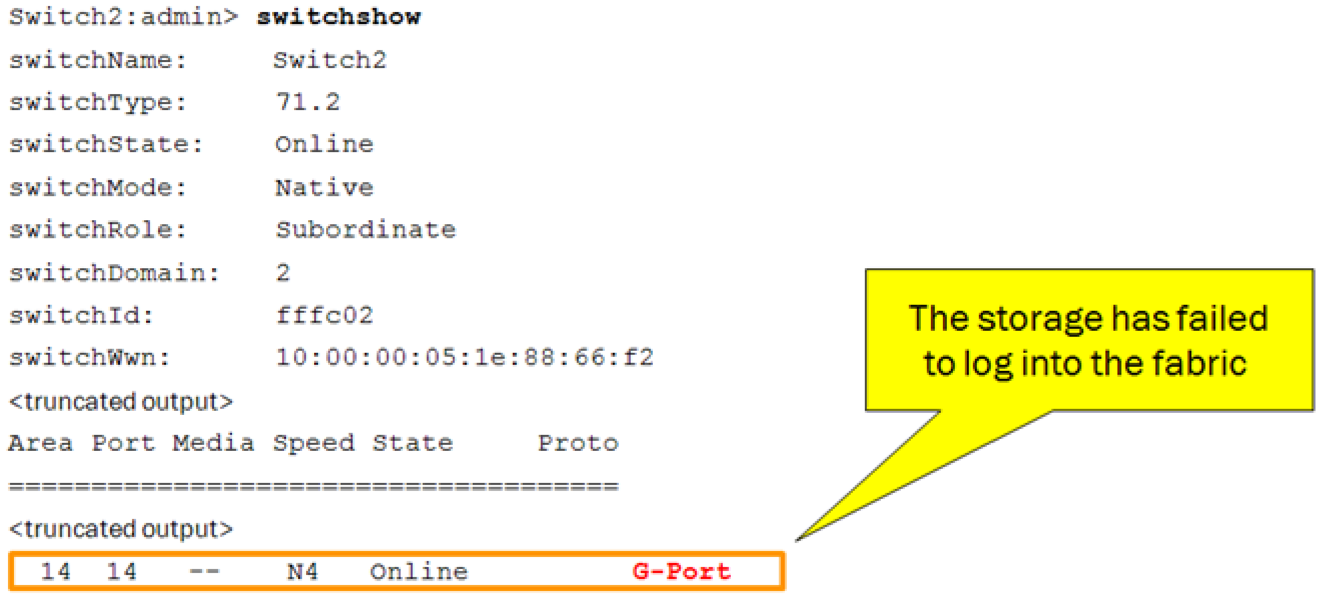

```
Switch2:admin> switchshow
switchName:      Switch2
switchType:      71.2
switchState:     Online
switchMode:      Native
switchRole:      Subordinate
switchDomain:    2
switchId:        fffc02
switchWwn:       10:00:00:05:1e:88:66:f2
<truncated output>
Area Port Media Speed State      Proto
=======================================
<truncated output>
  14   14    --      N4    Online             G-Port
```

The storage has failed to log into the fabric

Figure 20: G_Port

A `switchshow` displaying a G_Port indicates that port 14 made an incomplete connection to the fabric. A port stuck in a G_Port state indicates the link is in the active state but the attached device is not transmitting a FLOGI. Possible causes for this are:

- Issue with the driver on the attached device
- Possible, but much less likely, is a physical issue

NOTE
If it has been verified that the storage device is connected to the switch then use diagnostic tools, such as D_Port and SuperPing (`fcping`), to test the ISL connections between host and storage.

Marginal Links

A marginal link occurs when the connections between the switch and the edge device are not operating up to specifications identified by hardware reference manuals. 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, and edge device SFP).

Troubleshooting a marginal link can involve replacing cables and SFPs, inspecting error counters, and running diagnostics on a link, a port, or an end-to-end path.

1. Enter the `porterrshow` command.
2. Determine whether there is a relatively high number of errors (such as `CRC ERR` or `ENC_OUT` errors), or if there are a steadily increasing number of errors to confirm a marginal link. Sample the data every 5 minutes until you see the counters increment.
    - The `frames tx` and `rx` are the number of frames being transmitted and received.

- The `crc_err` counter are 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 an 8 Gbps blade, use the `portcfgfillword` command to reduce EMI.
- The `crc_g_eof` counter are 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, since 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` are errors that occur outside the frame and usually indicating a bad primitive. To determine if you are having a cable problem, take snapshots of the port errors by using the `porterrshow` command in increments of 5 to 10 minutes. If you notice the `crc_err` counter go up, you have a bad or damaged cable, or a bad or damaged device in the path.

  NOTE
  ICLs see `enc_out` errors when ports on one side of the link are disabled.
- The `disc_c3` errors are 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.

3. If you suspect a marginal link, isolate the areas by moving the suspected marginal port cable to a different port on the switch. Reseating of SFPs may also cure marginal port problems.
   - If the problem stops or goes away, the switch port or the SFP is marginal (proceed to step 6).
   - If the problem does not stop or go away, see Step 7.
4. Run `portloopbacktest` on the marginal port. You need an adapter to run the loopback test for the SFP. Otherwise, run the test on the marginal port using the loopback mode lb=5.
5. Check the results of the loopback test and proceed as follows:
   - If the loopback test failed, the port is bad. Replace the port blade or switch.
   - If the loopback test did not fail, the SFP was bad.
6. Replace the SFP on the marginal port.
7. 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.

The `fcping` Command

When you use `fcping` with a source and a destination, the command performs a zoning check between the two ports. In addition, two Fibre Channel 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 identifier. The ELS Echo request elicits an ELS Echo response from a port identifier in the fabric and is useful for validating link connectivity.

The source and destination port identifiers can be specified as a 24-bit Fibre Channel port identifier (PID), a port World Wide Name, or a node World Wide Name. The two port identifiers are then used to determine if the identifiers are zoned together.

Taking the Test

After the Introduction Screen, once you click on **Next**, you will see the following 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 credential, 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. Brocade also reserves the right to use any technologies and methods for verifying the identity of candidates. Such technology may include, without limitation, personally identifiable information, challenge questions, identification numbers, photographic information, and other measures to protect against fraud and abuse.
>
> 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.
>
> A. YES, I AGREE
>
> B. NO, I DO NOT AGREE