



BCFD in a Nutshell Study Guide for Exam 143-270



Brocade University

Revision 0412

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

© 2012 Brocade Communications Systems, Inc. All Rights Reserved.

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

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

Revision 0412

BCFD in a Nutshell First Edition



Objective: The BCFD Nutshell guide is designed to help you prepare for the BCFD Certification, exam number 143-270.

Audience: The BCFD Nutshell self-study guide is intended for those who have successfully completed the Course# withFullName Training course, and who wish to undertake self-study or review activities before taking the actual BCFD exam. The BCFD guide is not intended as a substitute for classroom training or hands-on time with Brocade products.

How to make the most of the BCFD guide: The BCFD guide summarizes the key topics on the BCFD 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 BCFD guide, we strongly recommend you have successfully completed the Course# withFullName course.

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

Joe Cannata
Certification Manager

A handwritten signature in blue ink that reads "Joe Cannata".

Table of Contents

1 — Pre-Design Data Gathering

SAN Planning	1
Gathering Requirements	2
Virtualization Design Guidelines	3
Cable Considerations	4
Tools for Gathering Data	5
Determining SAN Bandwidth for Backups	6

2 — Design Assessment

Access Gateway Design Guidelines	7
In-Flight Encryption and Compression Design Guidelines	7
Virtual Fabrics Design Guidelines	7
Interoperability Between Fabric OS and M-EOS Fabrics Using an FC Router	8
Long Distance Fabrics Design Considerations	9
FCIP Platforms and Supported Features	9
ICL Connectivity for Brocade DCX 8510-8 and DCX 8510-4	10
FCIP Fastwrite and OSTP configurations	13
Redundancy and Resiliency	15

3 — Management and Monitoring Tools

Health Monitoring	16
Advanced Performance Monitors	17
End-to-End Performance Monitoring	17
Frame Monitoring	18
Top Talker	18
Brocade Network Advisor	19

4 — Hardware and Software Products and Features

FC-FC Routing Supported Configurations	20
Virtual Fabrics	20
SAN Configuration with Access Gateway	21
Diagnostic Port	21
QoS: SID/DID Traffic Prioritization	22
QoS Zone-based Traffic Prioritization	22
Traffic Isolation Zoning	22
FICON and FICON CUP in Virtual Fabrics	23
Fabric OS Encryption	24
FCoE overview	24

5 — Distance Solutions

Distance Extension Topologies	26
Buffer Credit Management.	26
Wave Division Multiplexing (WDM)	26
FCIP Design Best Practices	27
Compression	28

6 — Performance and Tuning Optimization

Trunking Overview.	30
F_Port Trunking for Access Gateway	31
EX_Port Trunking.	31
Configuring Encryption and Compression	31
Fibre Channel Link Lengths and Loss Budget	32
Sources of Congestion	33
Topologies	34

7 — Migration, Integration, and Validation

Migration Strategies	36
Integrating and Consolidating Fabrics.	37
Zoning Overview	37
Upgrading Firmware	38
Port Decommissioning	38
Access Gateway Supported Hardware	39

8 — Security

Device and Switch Access	40
Administrative Access.	41

Taking the Test	42
------------------------------	-----------

List of Figures

Core-Edge Multi-Chassis ICL Configurations	11
Minimum connections needed between two Brocade DCX 8510 chassis	11
ICL-based Full-Mesh Topology	12
Single Tunnel, FastWrite and OSTP enabled	13
Multiple Tunnels to Multiple Ports: FastWrite and OSTP enabled on a per-Tunnel/per-Port basis ..	14
Traffic Isolation Zone Creating a Dedicated Path Through the Fabric	23
FCIP Tunnel and Circuits	28
Four scenarios of tiered network topologies (hops shown in heavier, orange connections)	34

List of Tables

Licensing	10
Standard Link Lengths	32
Link Loss Budget	32

1 – Pre-Design Data Gathering

When you finish this section you should be able to perform the following tasks:

- Identify how to assess current and future SAN infrastructure requirements
- Identify the information that must be captured for applications, servers, storage, and network devices to design a fabric
- Gather requirements to design a fabric to meet Service Level Agreements (SLAs)

SAN Planning

SAN planning process is similar to any project planning. The length and duration is dependent on the business critically of the project. Planning can be broken down to the following phases:

- Phase I—Gathering Requirements
- Phase II—Developing Technical Specifications
- Phase III—Estimating Project Cost
- Phase IV—ROI or TCO Analysis (if needed)
- Phase V—Detailed SAN Design and Rollout Plan

Deploying new SANs or expanding existing SANs to meet additional workloads in the fabrics requires critical assessment of business and technology requirements. A fabric design should not only take into account existing requirements but should plan for a 4-6 year life cycle. Focus on planning will ensure that the deployed SAN meets all business objectives, including cost, availability, deployment simplicity, performance and managing future business growth. Tables in Appendix A are provided as a reference for documenting assets and metrics for the SAN project.

There are several factors that will heavily influence your hardware purchasing decision process:

- What level of availability does your data environment require?
- How many host and storage ports are required for your selected level of availability?
- What port count growth rate do you foresee for the future?
- Will your design be optimized for performance, cost, or both?
- Do advanced capabilities such as Adaptive Networking and Integrated Routing play a part in your design?
- Is this an open systems, FICON, or intermix environment?

In a typical design cycle, the customer should have a fairly clear answer to these questions before exploring hardware and software possibilities.

Traffic Locality

Designing device connectivity depends a great deal on the expected data flow between devices. For simplicity, communicating hosts and targets can be attached to the same switch.

Gathering Requirements

The SAN project team should interview all stakeholders (IT application owners, finance, corporate facilities, IT lab administrators, storage and network administrators, and end users) who have a vested interest in the project—and this applies equally to planning for both new and updated SANs.

- IT Application Owners: As critical stakeholders, IT application owners care because everyone is measured on application uptime. Application outages are something that users notice, and they can have severe financial impact for a business. With a redundant or a resilient infrastructure, hardware outages are transparent to the user, and only SAN administrators need to pay attention. Consider the following:
 - What is the business goal for this application? (Is it a database that multiple applications rely on for business transactions?)
 - What are the availability requirements?
 - Is the application latency sensitive?
 - Are there peak periods of utilization or other traffic patterns?
- Server and Storage Administrators: Once the application requirements have been defined, identify physical server and storage on which the application and data will reside to determine the overall high-level architecture of the SAN, especially if this includes existing equipment as well as new equipment. Consider the following:
 - Gather information about the servers on which the applications are running (blade or rack, CPU, memory, HBA/embedded FC switch, OS level, OS patch level, HBA driver version)?
 - What is the primary storage for the application, and is there enough storage capacity to support this application and data? What is the current cache utilization? Is there enough cache to meet required response times?
 - How many FC ports are there in the array? Will you need to purchase new equipment?
- SAN Administrator—General: A SAN administrator is responsible for the day-to-day operation of the network. The SAN design must be easy to monitor, manage, and maintain. If the current SAN is being expanded, adequate performance metrics should be collected to ensure that the existing design can be expanded to address new workloads. Consider the following:
 - Are there performance (bandwidth) or latency issues in the existing SAN?
 - Are procedures in place to address redistribution of capacity when switch port utilization exceeds 75 percent?
 - Is the current design two-tier (core-edge) or three-tier (edge-core-edge)?
 - Is the current SAN a redundant configuration?
 - Is there an identified server to capture logs from the fabric?
 - Is the traffic equally distributed across the ISLs or the trunks?

- SAN Administrator—Backup and Restore: Backup and restore continue to be the primary drivers for SANs. As data growth continues to increase, backup windows continue to shrink. What is often overlooked is the restore time, which for some customers can take days. Consider the following:
 - If the backup site is local, what is the window to complete the backup? If the backup site is remote, what is the window to complete the backup? How much of the bandwidth pipe is available?
 - Is there a dedicated backup server, or do other applications share the server? Is the backup SAN on a separate SAN or a shared network?
 - How often are full backups completed, and how long does it take? How often are backups checked for the integrity of the backup? How often do the backups fail to complete? What are the primary reasons (link down, tape drive failure, low throughput, other)? What is the restore time for Tier 1 and 2 applications?
- Facilities: Facility requirements are often overlooked as SANs grow due to business expansion or data center consolidation after mergers. Even when a SAN design meets application requirements, if physical plant, power, cooling, and cable infrastructure are not available, a logically designed SAN may have to be physically distributed, which can impact application performance and ongoing servicing. Consider the following:
 - Is there existing space for new SAN devices (servers, switches, and storage)? What is the physical real estate (floor space, number of racks, rack dimensions), and do the racks have internal fans for cooling?
 - What is the available power (AC 120/240), and what is the in-cabinet power and plug type? Is it the same as existing types, or do you need new power supplies?
 - What method of cooling is available (hot/cool aisle, other), and what is the worst-case temperature that the data center can tolerate?
 - What is the cable infrastructure (OM3, other), and are cables already installed?
 - Is there a structured cable plant with patch panels, and so forth? If so, how many patch panels will the data traverse?
- Finance: Once the technical specifications have been determined, a reasonable cost estimate can be calculated based on available equipment, new purchases required, manpower, and training. Financial metrics for a total cost analysis should include the following:
 - Lease versus buy
 - Budget for equipment
 - Budget for service and support (is 24x7 required?)
 - Budget for daily operation

Virtualization Design Guidelines

With the higher levels of I/O potentially occurring at each edge port in the fabric, you must ensure that there is sufficient bandwidth and paths across the fabric to accommodate the load. Consider a lot of trunked ISLs and lower subscription ratios on the ISLs, if at all possible. Remember that many flows are partially hidden due to the increased use of NPIV. Frame congestion is also a greater possibility. Many of the VMs may still be in clusters and may require careful configuration. Spread out the LUNs across a lot of storage ports.

Separate the hypervisors on separate directors and, certainly, keep them separate from storage ports. This allows you to very easily apply controls through Brocade Fabric Watch classes without affecting storage. Use the following questions to guide you in designing a virtualization environment:

- Which applications will run under a virtual machine (VM) environment?

- How many VMs per server?
- Migration of VMs under what conditions (business and non-business hours, need additional CPU or memory to maintain response times)?
- Is there a need for SSDs to improve read response times?

Determine what latencies are tolerable to both storage and hosts (VMs and storage), and consider setting Brocade FOS thresholds accordingly.

Port Fencing is a powerful tool. Once many applications are running in VMs on a single physical platform, take care to ensure that Port Fencing does not disable ports too quickly.

Design Guidelines

- If possible, try to deploy VMs to minimize VM migrations if you are using shared LUNs.
- Use individual LUNs for any I/O-intensive applications such as SQL Server, Oracle databases, and Microsoft Exchange.

Cable Considerations

Link loss is a combination of fiber attenuation related to the distance of the link and the connectors or splices in the link. There are a number of reasons why understanding and “budgeting” potential link loss in a data center network is important:

- Operating within certain parameters is critical for establishing a reliable and stable network.
- The link loss measurement can assist with the task of selecting cable type (multimode fiber or singlemode fiber), transmitter wavelength (850 nm, 1300 nm, and so on), and other network components.
- It can help you determine the maximum physical distance of a link.
- It is a means of estimating whether or not current and future equipment can be supported over a cabling system.

The following factors can impact signal power loss:

- Number of paired connectors in the link
- Number and type of splices in the fiber
- Stressing of the cable
- Extending the length of the cable in the link
- Misaligned optical couplers
- Change in temperature
- Degradation in end equipment signal transmitter
- Aging of passive components

Tools for Gathering Data

SAN Health

Brocade SAN Health is a free tool that allows SAN Administrators to securely capture, analyze, and report comprehensive information about their Brocade fabrics running on Fabric OS and M-EOS operating systems and Cisco fabrics running MDS. It is able to perform tasks such as:

- 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 reports (excel) and diagrams (Visio)

SupportSave

The output from the `supportsave` command contains different commands, such as `islshow` and `fabricshow`, that can help you diagram your fabric.

Brocade Network Advisor

Performance monitoring provides details about the quantity of traffic and errors a specific port or device generates on the fabric over a specific time frame. You can also use performance to indicate the devices that create the most traffic and to identify the ports that are most congested. Performance allows you to monitor your SAN using the following methods (requires a Licensed version):

- Display the connections which are using the most bandwidth on the selected device or one of the F_Ports on the device with a feature called Top Talkers.
- Gather and display real-time performance data (FC ports, ISL ports, Device ports, GE ports, FCIP tunnels, managed HBA ports, managed CNA ports, E_Port trunks, and 10 GE ports).

The Professional version only allows you to monitor your SAN by gathering and displaying real-time performance data (FC ports, ISL ports, Device ports, GE ports, FCIP tunnels, Managed HBA ports, Managed CNA ports, E port trunks, and 10 GE ports).

- Persist and display historical performance data (FC ports, ISL ports, Device ports, FCIP tunnels, and 10 GE ports) for selected fabrics or the entire SAN.
- Create custom port and time data filters for historical performance data that can be saved as a Favorite.
- Support End-to-End monitors for real-time and historical performance data.
- Enforce user-defined performance thresholds and notification when thresholds are exceeded.

NOTE

When the server is busy and you request performance statistics, an "insufficient resources" message displays. Wait a while before you request the performance statistics again.

- Display percentage utilization for FC and FCIP links.

Server Traffic Patterns

On the server side, there are Windows and Unix tools for collecting CPU, memory, and network utilization built into the OS. HBA vendors also provide tools to gather the following on a per-port basis:

- Percent of reads
- MB/s reads
- Percent of writes
- MB/s writes
- Worst-case latency (ms)
- HBA queue depth

When planning to migrate a group of locally hosted applications to a remote facility, get the average frame payload using a dedicated long distance FC link. This will aid you in determining the amount of bandwidth needed for current applications and how much more bandwidth you will need for future scalability.

Determining SAN Bandwidth for Backups

At a minimum, available bandwidth in the fabric should be able to support applications and backup throughput. For example, in an edge-core-edge topology, the ISL paths from the storage-core tape and host-core tape should be able to support total throughput of all active tape drives and all applications without congestion. These paths should be redundant so that the failure of an ISL will not cause congestion in the fabric, impacting application or backup performance.

The key drivers for data recovery include the following:

- How quickly access to data is restored, called the Recovery Time Objective (RTO)
- The point in time in which the last valid data transaction was captured, called the Recovery Point Objective (RPO)
- Where the recovered data is located

2 – Design Assessment

When you finish this section you should be able to perform the following tasks:

- Design a solution that meets the customer's requirements
- Demonstrate knowledge of resiliency, redundancy, high availability, and locality
- Describe the various documents required in the design assessment

Access Gateway Design Guidelines

Use the Brocade Access Gateway when you deploy bladed environments or have a lot of low-port-count switches and when you need to connect different servers in different fabrics from a single-bladed enclosure. The Access Gateway can be very valuable when you want to separate the management of blade enclosures so that the enclosure is completely managed by server administrators, and the fabric is handled by storage administrators. Management separation is provided through the NPIV connection, which allows the Access Gateway to be managed separately by tools such as integrated blade server enclosure management tools without any adverse effects on the fabric.

In-Flight Encryption and Compression Design Guidelines

Brocade 16-Gbps platforms support both in-flight compression and/or encryption at a port level for both local and long-distance ISL links. In-flight data compression is a useful tool for saving money when either bandwidth caps or bandwidth usage charges are in place for transferring data between fabrics. Similarly, in-flight encryption enables a further layer of security with no key management overhead when transferring data between local and long-distance data centers besides the initial setup.

- It is supported only on 16 Gbps E_Ports
- ISL ports must be set to LD (long-distance) mode when compression is used
- Twice the number of buffers should be allocated if compression is enabled for long distance
- If both compression and encryption are used, enable compression first
- When implementing ISL encryption, using multiple ISLs between the same switch pair requires that all ISLs be configured for encryption—or none at all
- No more than two ports on one ASIC can be configured with encryption, compression, or both
- Encryption is not compliant with Federal Information Processing Standards (FIPS)

Virtual Fabrics Design Guidelines

The Brocade FOS Virtual Fabrics (VF) feature provides a mechanism for partitioning and sharing hardware resources, with the intention of providing more efficient use, increased fault isolation, and improved scalability. Virtual Fabrics use hardware-level fabric isolation between Logical Switches and fabrics. Virtual Fabrics consist of one or more Logical Switches and physical switches (non-partitioned).

If no local switching is used, any set of ports in the chassis/fabric can be used to create a Virtual Fabric. If local switching is used, ports for the VF should be from the same port groups.

Use Case: FICON and Open Systems (Intermix)

Virtual Fabrics enable customers to share FICON and FCP traffic on the same physical platform. As chassis densities increase, this is a viable option for improved hardware utilization while maintaining director class availability. The primary reasons for moving to an Intermix environment are the following:

- ESCON-FICON migration
- Sharing of infrastructure in a non-production environment
- Reduced Total Cost of Ownership (TCO)
- Growth of zLinux on the mainframe

From a SAN design perspective, the following guidelines are recommended when considering FICON Intermix:

- Connect devices across port blades (connectivity from the same device should be spread over multiple blades)
- One-hop count still applies

Interoperability Between Fabric OS and M-EOS Fabrics Using an FC Router

You can set up your Fabric OS SAN and M-EOS SAN to route traffic through FC router. Unlike with earlier releases of Fabric OS, you cannot mix Fabric OS v7.0.0 with M-EOS switches in the same L2 fabric. In Fabric OS v7.0.0 and later releases, the only way you can interoperate between Fabric OS and M-EOS fabrics is through FC router, which must connect to the M-EOS fabric through an EX_Port.

Fabric OS provides the ability to configure any EX_Port to connect to an M-EOS fabric by using an E_Port without disrupting the existing services. All EX_Port functions, such as fabric isolation and device sharing, remain the same as when connecting to a Fabric OS fabric.

Supported Platforms for Fibre Channel Routing

Fibre Channel routing is supported on the following platforms:

- Enterprise-class platforms: Brocade DCX, DCX-4S, and DCX 8510 family
 - 8 Gbps port blades (FC8-16, FC8-32, FC8-48, FC8-64)
 - 16 Gbps port blades (FC16-32, FC16-48)
 - FX8-24 DCX Extension Blade
 - FR4-18i Router Blade (Brocade DCX and DCX-4S only, and only VEX_Ports)
- Brocade 5100 switch
- Brocade 5300 switch
- Brocade 6510 switch
- Brocade VA-40FC switch
- Brocade 7800 Extension Switch
- Brocade Encryption Switch

For the enterprise-class platforms, note the following restrictions:

- EX_Ports and VEX_Ports are supported only on the FX8-24 DCX Extension Blade, and the 8-Gbps and 16-Gbps port blades. Ports on the core blade cannot be configured as EX_Ports.

- VEX_Ports are supported on the FR4-18i Router Blade, but EX_Ports are not supported. The FR4-18i blade is not supported in the same chassis as the FX8-24 blade.
- The enterprise-class platforms have a limit of 128 EX_Ports for each chassis.

Long Distance Fabrics Design Considerations

The most effective configuration for implementing long-distance SAN fabrics is to deploy Fibre Channel switches at each location in the SAN. Each switch handles local interconnectivity and multiplexes traffic across long-distance dark fiber or wave division multiplexing (WDM) links while the Brocade Extended Fabrics software enables SAN management over long distances. Brocade Extended Fabrics is an optional licensed feature for Brocade SAN deployment over distance beyond 10 km. A Brocade Extended Fabrics license is required before you can implement long distance dynamic (LD) and long distance static (LS) distance levels. The LD and LS settings are necessary to achieve maximum performance results over Inter-Switch Links (ISLs) that are greater than 10 km. The Extended Fabrics feature enables the following:

- Fabric interconnectivity over Fibre Channel at longer distances
ISLs can use long distance dark fiber connections to transfer data. Wave division multiplexing, such as DWDM (Dense Wave Division Multiplexing), CWDM (Coarse Wave Division Multiplexing), and TDM (Time Division Multiplexing), can be used to increase the capacity of the links. As Fibre Channel speeds increase, the maximum distance decreases for each switch. The Extended Fabrics feature extends the distance the ISLs can reach over an extended fiber. This is accomplished by providing enough buffer credits on each side of the link to compensate for latency introduced by the extended distance.
- Simplified management over distance
Each device attached to the SAN appears as a local device, an approach that simplifies deployment and administration.
- Optimized switch buffering
When Extended Fabrics is installed on gateway switches (E_Port connectivity from one switch to another), the ISLs (E_Ports) are configured with a large pool of buffer credits. The enhanced switch buffers help ensure that data transfer can occur at near-full bandwidth to efficiently utilize the connection over the extended links. This ensures the highest possible performance on ISLs.

FCIP Platforms and Supported Features

Fibre Channel over IP (FCIP) enables you to use existing IP wide area network (WAN) infrastructure to connect Fibre Channel SANs. FCIP supports applications such as remote data replication (RDR), centralized SAN backup, and data migration over very long distances that are impractical or very costly using native Fibre Channel connections. FCIP tunnels are used to pass Fibre Channel I/O through an IP network. FCIP tunnels are built on a physical connection between two peer switches or blades. Fibre Channel frames enter FCIP through virtual E_Ports (VE_ports or VEX_Ports) and are encapsulated and passed to Transmission Control Protocol (TCP) layer connections. The TCP connections ensure in-order delivery of FC frames and lossless transmission. The Fibre Channel fabric and all Fibre Channel targets and initiators are unaware of the presence of the IP network.

There are three Brocade platforms that support FCIP:

- Brocade 7800 switch
- Brocade FX8-24 blade (DCX, DCX-4S, DCX 8510-8, and DCX 8510-4 chassis)
- Brocade FR4-18i blade (DCX, DCX-4S chassis)

Table 1 lists the licensing applicable to supporting FCIP platforms.

TABLE 1 Licensing

License	Description
10 Gigabit FCIP/Fibre Channel License (10G license)	<ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 10 1 GbE ports and 1 10 GbE port 2 10 GbE ports License is slot based when applied to a Brocade enterprise-class platform. It is chassis based when applied to a Brocade 6510 switch.
7800 Upgrade License	<ul style="list-style-type: none"> 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. <p>NOTE: The Brocade 7800 switch must have the Upgrade License to add FICON Management Server (CUP) or Advanced Accelerator for FICON.</p>
Advanced Extension License	<ul style="list-style-type: none"> Enables 2 advanced extension features: FCIP Trunking and Adaptive Rate Limiting. FCIP Trunking feature allows all of the following: <ul style="list-style-type: none"> Multiple (up to 4) IP source and destination address pairs (defined as FCIP Circuits) using multiple (up to 4) 1 GbE or 10 GbE interfaces to provide a high bandwidth FCIP tunnel and failover resiliency. Support for up to 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	<ul style="list-style-type: none"> 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.
Integrated Routing	<ul style="list-style-type: none"> Allows any ports in a Brocade 5100, 5300, 6510, 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). Eliminates the need to add an FR4-18i blade or use the 7500 for FCR purposes.

ICL Connectivity for Brocade DCX 8510-8 and DCX 8510-4

The Brocade DCX 8510-8 and DCX 8510-4 platforms use second generation ICL technology from Brocade with optical Quad Small Form Factor (QSFP). The Brocade DCX 8510-8 allows up to 32 QSFP ports, and the Brocade DCX 8510-4 allows up to 16 QSFP ports to help preserve switch ports for end devices. Each QSFP port has four independent 16-Gbps links, each of which terminates on a different ASIC within the core blade.

Each core blade has four ASICs. A pair of connections between two QSFP ports can create 32 Gbps of bandwidth. Figure 1 illustrates a core-edge design based on ICLs supporting 2,304 16-Gbps ports with a minimum of 256 Gbps of bandwidth between the chassis (12:1 oversubscription). As more ICLs are added, oversubscription can be reduced to 6:1. The ICL architecture for the DCX 8510s support up to 50 meters optical.




FIGURE 1 Core-Edge Multi-Chassis ICL Configurations

To connect one Brocade DCX 8510 to another, a minimum of four QSFP ports must be connected to the second chassis. This provides at least 256 Gbps of bandwidth between the Brocade DCX 8510s. The dual connections must also reside within the same ICL trunk boundary on the core blades for creating ICL trunks, as shown in Figure 7. If you want to add more ICL cables, one QSFP cable from each blade must be connected to the same blade of its neighbor. For example, the connection from Core Blade 1 must connect to Core Blade 1 on the neighboring chassis.

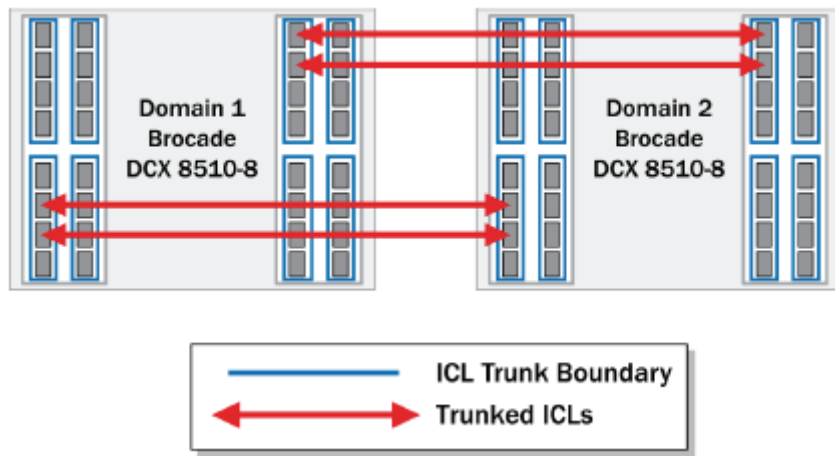


FIGURE 2 Minimum connections needed between two Brocade DCX 8510 chassis

Mesh Topology

A mesh design provides a single hop between source and destination, and initial Brocade Fabric OS 7.0 release supports a 3 chassis mesh design (same as existing 8 Gb platform) with 15/25/50 meter distance. In the configuration shown in [Figure 3](#), up to 1152 16 Gbps ports are supported using ICLs with a 12:1 oversubscription. As more ICLs are added, oversubscription can be reduced to 3:1.




FIGURE 3 ICL-based Full-Mesh Topology

FCIP Fastwrite and OSTP configurations

The FCP features used in FCIP Fastwrite and OSTP require a deterministic FC Frame path between initiators and targets when multiple tunnels exist. If there are non-controlled parallel (equal-cost) tunnels between the same SID/DID pairs, protocol optimization will fail when a command is routed over one tunnel and the response is returned over a different tunnel. To help understand the supported configurations, consider the configurations shown in [Figure 4 on page 13](#) and [Figure 5 on page 14](#). In both cases, there are no multiple equal-cost paths. In [Figure 4](#), there is a single tunnel with Fastwrite and OSTP enabled. In [Figure 5](#), there are multiple tunnels, but none of them create a multiple equal-cost path.

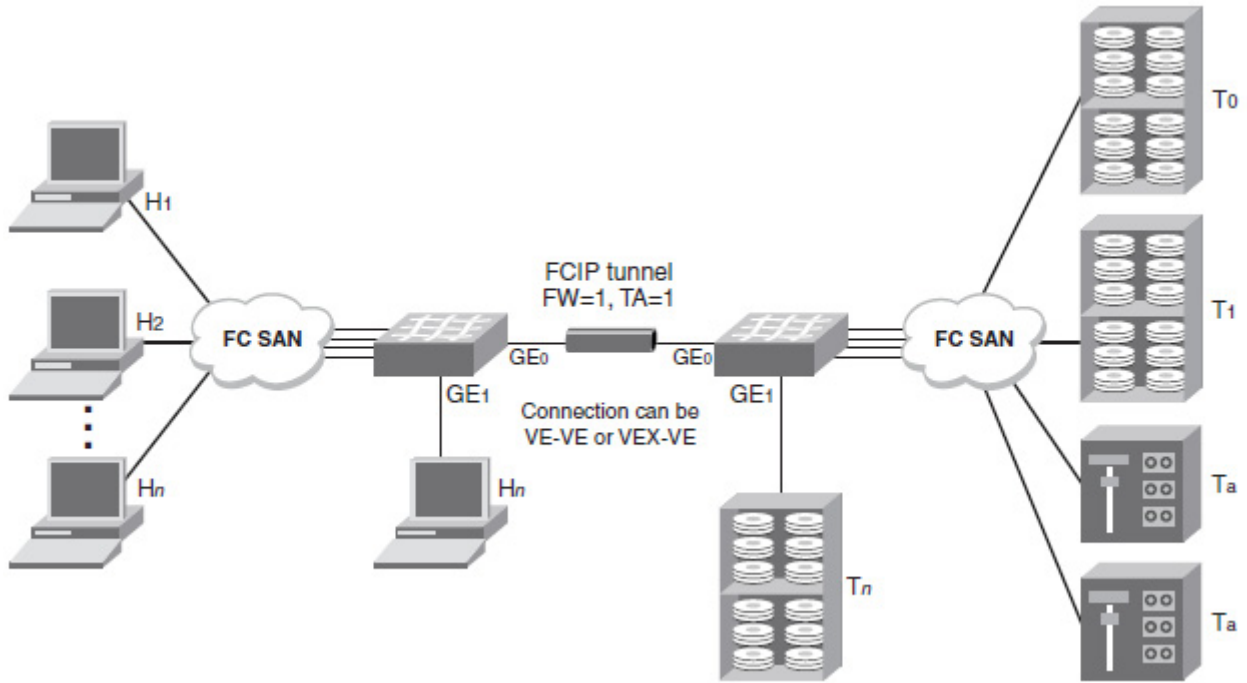


FIGURE 4 Single Tunnel, FastWrite and OSTP enabled




FIGURE 5 Multiple Tunnels to Multiple Ports: FastWrite and OSTP enabled on a per-Tunnel/per-Port basis

In some cases, Traffic Isolation zoning or LS/LF configurations may be used to control the routing of SID/DID pairs to individual tunnels and provide deterministic flows between the switches, allowing the use of multiple equal cost tunnels.

Redundancy and Resiliency

An important aspect of SAN topology is the resiliency and redundancy of the fabric. The main objective is to remove any single point of failure. Resiliency is the ability of the network to continue to function and/or recover from a failure, while redundancy describes duplication of components, even an entire fabric, to eliminate a single point of failure in the network. Brocade fabrics have resiliency built into Brocade Fabric OS® (Brocade FOS), the software that runs on all Brocade B-Series switches, which can quickly “repair” the network to overcome most failures. For example, when a link between switches fails, routing is quickly recalculated and traffic is assigned to the new route. Of course this assumes that there is a second route, which is when redundancy in the fabric becomes important.

The following recommendations for SAN design are to ensure application availability and resiliency:

- Redundancy built into fabrics to avoid a single point of failure
- Servers connected to storage via redundant fabrics
- MPIO-based failover from server to storage
- Redundant fabrics based on similar architectures
- Separate storage and server tiers for independent expansion
- At a minimum core switches should be of equal or higher performance compared to the edges.
- Define the highest performance switch in the fabric to be the principal switch.

3 – Management and Monitoring Tools

When you finish this section you should be able to perform the following tasks:

- Describe how to satisfy a specific monitoring requirement
- Demonstrate knowledge of Brocade management tools

Health Monitoring

Any mission-critical infrastructure must be properly monitored. The Brocade product line provides several mechanisms that can alert you when issues are detected in the fabric.

Brocade Fabric Watch

Fabric Watch is an optional health monitor that allows you to constantly monitor each director or switch for potential faults and automatically alerts you to problems long before they become costly failures.

Fabric Watch lets you define how often to measure each switch and fabric element and specify notification thresholds. Whenever fabric elements exceed these thresholds, Fabric Watch automatically provides notification using several methods, including e-mail messages, SNMP traps, and log entries.

Port Fencing

Port Fencing is a licensed feature that works with Fabric Watch to monitor ports for erratic behavior and disables a port if specified error conditions are met. You can customize the thresholds and configure the ports to report errors for one or more areas using the `portthconfig` command. After the ports are configured, you can enable port fencing for specific areas of the physical ports, E_Ports, FOP_Ports, and FCU_Ports using the `portfencing` command.

RAS Log

RAS log is the Brocade Fabric OS error message log. Messages are organized by Brocade FOS component, and each one has a unique identifier as well as severity, source and platform information and a text message.

RAS log is available from each switch and director using the `errdump` command and RAS log messages can be forwarded to a syslog server for centralized collection.

Audit Log

The Audit log is a collection of information created when specific events are identified on a Brocade platform. The log can be dumped using the `auditdump` command, and audit data can also be forwarded to a syslog server for centralized collection.

Information is collected on many different events associated with zoning, security, trunking, FCIP, FICON, and others. Each release of the Brocade Fabric OS provides more audit information.

Design Guidelines

- Transaction-based systems: Make sure that ISL/ICLs traversed by these systems to access their storage do not contain too many flows. The fan-in from the hosts/initiators should not exceed a ratio of 10 to 1. Also ensure that there is as little interference from other applications as possible, to ensure that latencies and congestion from other sources do not affect the overall performance of the applications.
- I/O-intensive applications: Bandwidth is the typical constraint for these systems. Modern fabrics typically provide more bandwidth than is needed, except for the most powerful hosts. Take care to ensure that these systems do not interfere with other applications, particularly if they are run at specific times or if batch runs are scheduled. When in doubt, add more paths (ISLs or trunks) through the fabric.
- Clusters: Clusters often have behavioral side effects that must be considered. This is particularly true during storage provisioning. It is possible, for example, for a cluster to inundate the fabric and storage arrays with LUN status queried and other short frame requests. This behavior can cause frame congestion in the fabric and can stress the control processors of the arrays. Make sure that you spread out the LUNs accessed by the hosts in the cluster across as many arrays as possible.
- Congestion: Traffic congestion (total link capacity regularly consumed) is remedied by adding more links or more members to a trunk. Frame congestion is typically addressed by dealing with the nodes causing the congestion.
- Misbehaving devices: It has been stated earlier that there is very little that a fabric can do to mitigate the effects of a badly behaving device, other than to remove it from the fabric. Brocade supports a Brocade FOS capability called Port Fencing, which is designed to isolate rogue devices from the network. Port Fencing works with Brocade Fabric Watch to disable a port when a specific threshold has been reached. Port Fencing, in combination with Bottleneck Detection, can be used for detecting and isolating high-latency devices from impacting the rest of the devices in the fabric.
- Initiator and targets: If possible, isolate host and storage ports on separate switches for much greater control over the types of controls that you can apply to misbehaving and high-latency devices. The effect on applications is typically much less severe if a host is disabled versus disabling a storage port, which may be servicing flows from many servers.

Advanced Performance Monitors

Advanced Performance Monitoring provides the following monitors:

- End-to-End monitors (EE monitors) measure the traffic between a host/target pair.
- Frame monitors measure the traffic transmitted through a port with specific values in the first 64 bytes of the frame.
- Top Talker monitors measure the flows that are major consumers of bandwidth on a switch or port.

End-to-End Performance Monitoring

Use End-to-End (EE) monitoring when you want to monitor throughput between a pair of devices. End-to-end performance monitoring counts the number of words in Fibre Channel frames for a specified Source ID (SID) and Destination ID (DID) pair.

To enable end-to-end performance monitoring, you must configure an EE monitor on a port, specifying the SID-DID pair (in hexadecimal). The monitor counts only those frames with matching SID and DID.

Each SID or DID has the following three fields:

- Domain ID (DD)
- Area ID (AA)
- AL_PA (PP)

For example, the SID 0x118a0f denotes DD 0x11, AA 0x8a, and AL_PA 0x0f.

An EE monitor includes these counts:

- RX_COUNT - words in frames received at the port
For frames received at the port with the EE monitor installed, the RX_COUNT is updated if the frame SID is the same as the SID in the monitor and the frame DID is the same as the DID in the monitor.
- TX_COUNT - words in frames transmitted from the port
For frames transmitted from the port with the EE monitor installed, TX_COUNT is updated if the frame DID is the same as the SID in the monitor and the frame SID is the same as the DID in the monitor.

Supported Port Configurations for EE Monitors

You can configure EE monitors on F_Ports and, depending on the switch model, on E_Ports. The following platforms support EE monitors on E_Ports:

- Brocade 6510
- Brocade DCX 8510 family

Identical EE monitors cannot be added to the same port. Two EE monitors are considered identical if they have the same SID and DID values after applying the end-to-end mask.

An EE monitor and a port Top Talker monitor cannot co-exist on the same port. Co-existence of EE monitors and Top Talker monitors on ports belonging to the same ASIC is not recommended because the statistics for the same flow going through ports on the same ASIC might be inaccurate.

Frame Monitoring

Frame monitoring counts the number of times a frame with a particular pattern is transmitted by a port and generates alerts when thresholds are crossed. Frame monitoring is achieved by defining a filter, or frame type, for a particular purpose. The frame type can be a standard type, or example, a SCSI read command filter that counts the number of SCSI read commands that have been transmitted by the port, or a user-defined frame type customized for your particular use.

The Advanced Performance Monitoring license is required to use the `fmmonitor` command. The monitoring functionality, however, also requires the Fabric Watch license. When you configure actions and alerts using the `fmmonitor` command, Fabric Watch uses these values and generates alerts based on the configuration. If you do not have a Fabric Watch license, these values are ignored.

Top Talker

The Top Talker feature is based on the Advanced Performance Monitor (APM) feature. Top Talker monitors determine the flows (SID/DID pairs) through one or more switch F_Ports on any switch in the fabric that are the major users of bandwidth (after initial stabilization). Top Talker monitors measure bandwidth usage data in real-time and relative to the port on which the monitor is installed. This feature is not supported on Access Gateways.

Top Talker Monitors discard bandwidth information collected during the initial stabilization. Initial stabilization is the time taken by a flow to reach the maximum bandwidth. This time varies depending on the number of flows in the fabric and other factors. The incubation period can be up to 14 seconds in the enterprise-class platforms, and up to 82 seconds in the fixed-port switches.

Applications can use the Top Talker data to do the following:

- Re-route the traffic through different ports that are less busy, so as not to overload a given port.
- Alert you to the top-talking flows on a port if the total traffic on the port exceeds the acceptable bandwidth consumption.

Brocade Network Advisor

Brocade Network Advisor is a management application which provides easy, centralized management of the network, as well as quick access to all product configuration applications. Using this application, you can configure, manage, and monitor your networks with ease.

Call Home Feature

Call Home notification allows you to configure the Management application Server to automatically send an e-mail or dial-in to a support center to report system problems on specified devices (Fabric OS and M-EOS switches, routers, and directors). If you are upgrading from a previous release, all of your Call Home settings are preserved.

Call Home supports multiple call home centers which allows you to configure different devices to contact different call home centers. When you make any call home configuration changes or a call home event trigger occurs, the Management application generates an entry to the Master Log.

You can configure Call Home for the following call home centers:

- Brocade E-mail (Windows and Unix)
- Brocade International (Windows only)
- Brocade North America (Windows only)
- EMC (Windows only)
- HP LAN (Windows only)
- IBM (Windows only)
- IBM E-mail (Windows and Unix)
- Oracle E-mail (Windows and Unix)

Configuration Repository Management

You can save entire configurations of switch configuration files and use them to ensure consistent switch settings in your fabric, propagate configuration settings to additional switches in the fabric, and troubleshoot the switches. For complete feature support, you must have the Enterprise Edition license installed.

4 – Hardware and Software Products and Features

When you finish this section you should be able to perform the following tasks:

- Demonstrate knowledge of migration and connectivity for B-Series and M-Series products
- Describe Brocade hardware products and their purposes
- Demonstrate knowledge of Brocade software features and purposes

FC-FC Routing Supported Configurations

The FC-FC routing service provides Fibre Channel routing (FCR) between two or more fabrics without merging those fabrics. For example, using FCR you can share tape drives across multiple fabrics without the administrative problems, such as change management, network management, scalability, reliability, availability, and serviceability, that might result from merging the fabrics.

The supported configurations are:

- FC router connected to a Brocade nonsecured edge fabric.
- FC router connected to a Brocade secured edge fabric.
- FC router connected to a McDATA Open Mode edge fabric.
- FC router connected to a McDATA Fabric Mode edge fabric.
- FC router connected to Brocade secured and nonsecured fabrics with EX_Port trunking enabled.
- FC router interoperating with legacy FC routers (Brocade 7500 switch or FR4-18i blade).



Note

When configuring a fabric on which Fabric OS is installed to connect to a Native McDATA fabric, you must configure an EX_Port of an FC router to a Native McDATA fabric configured in Fabric mode.

License Requirements for Fibre Channel Routing

Fibre Channel routing is a licensed feature that requires the Integrated Routing license. This license allows 8-Gbps and 16-Gbps FC ports to be configured as EX_Ports (or VEX_Ports) which support Fibre Channel routing.

Virtual Fabrics

The Brocade FOS Virtual Fabrics feature provides a mechanism for partitioning and sharing hardware resources, with the intention of providing more efficient use, increased fault isolation, and improved scalability. Virtual Fabrics use hardware-level fabric isolation between Logical Switches and fabrics. Virtual Fabrics consist of one or more Logical Switches and physical switches (non-partitioned).

Virtual Fabric Guidelines

If no local switching is used, any set of ports in the chassis/fabric can be used to create a Virtual Fabric. If local switching is used, ports for the VF fabric should be from the same port groups.

Use Case: FICON and Open Systems (Intermix)

Virtual Fabrics enable customers to share FICON and FCP traffic on the same physical platform. As chassis densities increase, this is a viable option for improved hardware utilization while maintaining director class availability. The primary reasons for moving to an Intermix environment are the following:

- ESCON-FICON migration
- Sharing of infrastructure in a non-production environment
- Reduced Total Cost of Ownership (TCO)
- Growth of zLinux on the mainframe

From a SAN design perspective, the following guidelines are recommended when considering FICON Intermix:

- Connect devices across port blades (connectivity from the same device should be spread over multiple blades)
- One-hop count still applies

SAN Configuration with Access Gateway

Switches in Access Gateway mode can connect to third-party fabrics with the following firmware versions:

- M-EOSc v9.6.2 or later and M-EOSn v9.6 or later.
- Cisco MDS Switches with SAN OS v3.0(1).

Consider the following points when connecting multiple devices to a switch in AG mode:

- AG does not support daisy chaining when two AG devices are connected to each other in a loop configuration.
- Loop devices and FICON channels/control unit connectivity are not supported.
- When a switch is in AG mode, it can be connected to NPIV-enabled HBAs, or F_Ports that are NPIV-aware. Access Gateway supports NPIV industry standards per FC-LS-2 v1.4.

Diagnostic Port

Fabric OS v7.0.0 allows you to 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 on a Brocade DCX 8510, and running Fabric OS v7.0.0 or later.

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)

This fabric-based physical layer validation enables the following:

- Local and long-distance measurements (5 meter granularity for 16 Gbps SFPs and 50 meters for 10 Gbps SFPs)
- Latency measurements
- Link performance
- Transceiver health check
- Transceiver uptime

QoS: SID/DID Traffic Prioritization

Quality of Service (QoS) traffic shaping feature which allows the prioritization of data traffic based on the SID/DID of each frame. Traffic prioritization allows you to categorize the traffic flow between a host and target as having a high, medium, or low priority. Fabric OS supports two types of prioritization:

- Class Specific Control (CS_CTL)-based frame prioritization
Each frame between a host and a target is assigned a specific priority, depending on the value of the CS_CTL field in the frame header.
- QoS zone-based traffic prioritization
All traffic between a host and a target is assigned a specific priority, depending on the name you define for the QoS zone.

QoS Zone-based Traffic Prioritization

QoS zone-based traffic prioritization allows you to categorize the traffic flow between a host and target as having a high or low priority. For example, you could assign online transaction processing (OLTP) to high priority and backup traffic to low priority.

Traffic Isolation Zoning

The Traffic Isolation Zoning feature allows you to control the flow of interswitch traffic by creating a dedicated path for traffic flowing from a specific set of source ports (N_Ports). For example, you might use Traffic Isolation Zoning for the following scenarios:

- To dedicate an ISL to high priority, host-to-target traffic.
- To force high volume, low priority traffic onto a given ISL to limit the effect on the fabric of this high traffic pattern.
- To ensure that requests and responses of FCIP-based applications such as tape pipelining use the same VE_Port tunnel across a metaSAN.

Traffic Isolation Zoning does not require a license.

Traffic isolation is implemented using a special zone, called a Traffic Isolation zone (TI zone). A TI zone indicates the set of N_Ports and E_Ports to be used for a specific traffic flow. When a TI zone is activated, the fabric attempts to isolate all inter-switch traffic entering from a member of the zone to only those E_Ports that have been included in the zone. The fabric also attempts to exclude traffic not in the TI zone from using E_Ports within that TI zone.




FIGURE 6 Traffic Isolation Zone Creating a Dedicated Path Through the Fabric

FICON and FICON CUP in Virtual Fabrics

For FICON CUP, FICON Management Server (FMS) mode is enabled on the logical switch and not the chassis. For example, in a Virtual Fabrics environment on the Brocade 5100 switch and the Brocade DCX and DCX-4S, one would set CUP on each logical switch (currently limited to four logical switches per chassis).

There are two different addressing modes to provide extended functionality in dynamically created partitions.

- Flat or 10-bit addressing mode
- 256-area addressing mode (uses default or zero-based modes)

Platforms Supporting FICON

FICON protocol is supported on the following Brocade models for this release:

- Brocade DCX and DCX-4S Backbones (FC8-16 and FC8-32 port blades), FR4-18i and FX8-24 FCIP blades, and the FC10-6 10 Gbps port blade for ISL connections. The DCX with the FC8-48 and FC16-48 port blades are only supported in a logical switch defined for zero-based addressing.
- Brocade 5100 and 5300 switches.
- Brocade 6510 switch.
- The Brocade 7800 Extension Switch and the Brocade FR4-18i (for the Brocade DCX and DCX-4S) and FX8-24 blade (for the Brocade DCX and DCX-4S) provide FICON over IP extension.

Fabric OS Encryption

Based on industry standards, Brocade encryption for data-at-rest provides centralized, scalable encryption and compression services that seamlessly integrate into existing Brocade Fabric OS® (FOS).

The Brocade fabric-based approach to data encryption scales to meet performance requirements, provides a centralized point of management for both disk and tape storage security as well as key management, and supports heterogeneous storage environments. Deployment is simple and non-disruptive: Organizations can encrypt data from any switch port without reconfiguring the fabric.

Key advantages of the Brocade FS8-18 include:

- The ability to encrypt data at wire speed
- Central management of storage and fabric-based security resources
- Concurrent support for both disk and tape encryption operations from a single device
- Transparent, online encryption of “cleartext” LUNs and rekeying of encrypted LUNs without disruption
- Data compression and integrity authentication for tape backup data
- Simplified, non-disruptive installation and configuration

FCoE overview

Fibre Channel over Ethernet (FCoE) enables you to transport FC protocols and frames over Converged Enhanced Ethernet (CEE) networks. CEE is an enhanced Ethernet that enables the convergence of various applications in data centers (LAN, SAN, and HPC) onto a single interconnect technology.

Supported Converged Products

The following hardware platforms and software management tools are supported for Converged Enhanced Ethernet:

- Brocade VDX 6730 Data Center Switch
The Brocade VDX® 6730 Data Center Switch is a 10 Gigabit Ethernet (GbE) fixed port switch with LAN and native Fibre Channel ports. It supports multiple connectivity options, including classic ToR server deployments, Ethernet fabrics, and Ethernet storage connectivity for Fibre Channel over Ethernet (FCoE), iSCSI, NAS, and bridging Fibre Channel Storage Area Networks (SANs) and Ethernet fabrics.
- Brocade VDX 6720 Data Center Switch
The Brocade VDX® 6720 Data Center Switch is a high-performance, ultra-low latency wire-speed 10 Gigabit Ethernet (GbE) fixed port switch. It supports several connectivity options, including classic ToR server deployments, Ethernet fabrics, and Ethernet storage connectivity.
- Brocade 8000 Switch
The Brocade 8000 is a top-of-rack link layer (Layer 2) DCB/FCoE switch with 24 10 Gigabit Ethernet (GbE) ports for LAN connections and eight Fibre Channel ports (with up to 8 Gbps speed) for Fibre Channel SAN connections. This reliable, high-performance switch provides advanced Fibre Channel services, supports Ethernet and DCB capabilities, and is managed by Brocade DCFM.

- Brocade 1860 Fabric Adapter

The Brocade 1860 Fabric Adapter is a new class of adapter that meets all the connectivity needs of cloud-enabled data centers while providing unmatched performance, application-aware services, unified management, and reduced cost and complexity. It is the simplest, most flexible, and most powerful server connectivity adapter designed to extend fabric services to Virtual Machines (VMs) and applications in highly demanding virtualized environments

- Brocade 1010 and 1020 CNAs

The Brocade 1010 (single port) and Brocade 1020 (dual port) Converged Network Adapters (CNAs) integrate 10 Gbps Ethernet Network Interface Card (NIC) functionality with Fibre Channel technology—enabling transport over a 10 Gigabit Ethernet (GbE) connection through the new Data Center Bridging (DCB) and Fibre Channel over Ethernet (FCoE) protocols, providing best-in-class LAN connectivity and I/O consolidation to help reduce cost and complexity in next-generation data center environments.

- Brocade Network Advisor

Brocade Network Advisor is the industry's first unified network management solution for data, storage, application delivery, wireless, and converged networks. It supports Fibre Channel SANs, FCoE, IP switching and routing (including Ethernet fabrics), and MPLS networks — providing end-to-end visibility across different network types through a seamless and unified user experience.

5 – Distance Solutions

When you finish this section you should be able to perform the following tasks:

- Given availability, performance, and distance requirements, design an appropriate long-distance solution using Fibre Channel
- Given a specific set of requirements, demonstrate the ability to design a SAN extension solution using FCIP

Distance Extension Topologies

For a complete DR solution, SANs are typically connected over metro or long-distance networks. In both cases, path latency is critical for mirroring and replication solutions. For native Fibre Channel links, the amount of time that a frame spends on the cable between two ports is negligible, since that aspect of the connection speed is limited only by the speed of light. The speed of light in optics amounts to approximately 5 microseconds per kilometer, which is negligible compared to typical disk latency of 5 to 10 milliseconds. The Brocade Extended Fabrics feature enables full-bandwidth performance across distances spanning up to hundreds of kilometers. It extends the distance ISLs can reach over an extended fiber by providing enough buffer credits on each side of the link to compensate for latency introduced by the extended distance.

Buffer Credit Management

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

Buffer Allocation

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.

Consider the following when allocating the appropriate amount of buffers:

- Determine the desired distance in kilometers of the switch-to-switch connection.
- Determine the speed that you will use for the long-distance connection.

Wave Division Multiplexing (WDM)

Wave Division Multiplexing (WDM) describes the concept of combining several streams of data onto the same physical fiber-optic cabling where light of different wavelengths does not interfere.

- Dense Wavelength Division Multiplexing (DWDM) is optimized for high-speed, high-capacity networks and longer distances.
- Coarse Wavelength Division Multiplexing (CWDM) provides the same optical transport and features of DWDM, but at a lower capacity, which allows for lower cost.

There are two basic types of WDM solutions:

- Transponder-based Solutions
Using 850nm or 1310 nm, it converts these signals using Optical-to-Electrical-to-Optical (OE-O) conversion WDM frequencies for transport across a single fiber.
- SFP-Based Solutions
These eliminate the need for transponders by requiring switch equipment to utilize special WDM transceivers (also known as colored optics), reducing the overall cost.

FCIP Design Best Practices

Best practice is to always rate limit the FCIP traffic on the 7800 or FX8-24 blade and never rate limit FCIP traffic in the IP network, which often leads to problems that are difficult to troubleshoot. The rate limiting technology on the 7800/FX is advanced, accurate, and consistent, so there is no need to double rate limit. If policy required you to double rate limit, then the IP network should set its rate limiting above that of the 7800/FX with plenty of headroom.

Brocade 7800 and FX8-24 have an exclusive feature called FCIP Trunking. FCIP Trunking offers the ability to perform the following functions:

- Bandwidth aggregation
- Lossless failover
- Granular load balancing
- In-order delivery
- Prevent IFCC on mainframes

A single tunnel defined by a VE_Port or VEX_Port may have one or more circuits associated with it. A circuit is an FCIP connection defined by a source and destination IP address and other arguments that define its characteristics, such as compression, IPsec, QoS, rate limit, and others. All the circuits terminate at the single VE/VEX_Port on each side; therefore, there are no multiple tunnels or ISLs, but only a single tunnel load balanced across multiple circuits. The one ISL that an FCIP Trunk forms is from VE_port to VE_Port or VEX_Port to VE_Port.

The circuits can have different characteristics. They can have different RTTs (Round Trip Times) and take different paths and different service providers. They can have different bandwidths up to 4x. This means that if one circuit is an OC-3, the most the other circuit(s) can be is OC-12, because the bandwidth delta is 4x.

FCIP Trunking

FCIP trunking is a method for managing the use of WAN bandwidth and providing redundant paths over the WAN that can protect against transmission loss due to WAN failure. Trunking is enabled by creating logical circuits within an FCIP tunnel. A tunnel can have multiple circuits. You can configure up to 6 circuits on tunnels between 7800 switches and up to 10 on tunnels between FX8-24 blades. Each circuit is a connection between a pair of IP addresses that are associated with source and destination endpoints of an FCIP tunnel, as shown in Figure 7.




FIGURE 7 FCIP Tunnel and Circuits

Adaptive Rate Limiting

Adaptive Rate Limiting (ARL) is performed on FCIP circuits to change the rate in which the FCIP tunnel transmits data through the IP network. ARL uses information from the TCP connections to determine and adjust the rate limit for the FCIP circuit dynamically. This allows FCIP connections to utilize the maximum available bandwidth while providing a minimum bandwidth guarantee. ARL is configured on a per-circuit basis because each circuit may have available differing amounts of bandwidth.

ARL applies a minimum and maximum traffic rate, and allows the traffic demand and WAN connection quality to dynamically determine the rate. If traffic is flowing error-free over the WAN, the rate grows towards the maximum rate. If TCP reports an increase in retransmissions, the rate reduces towards the minimum. ARL never attempts to exceed the maximum configured value and reserves at least the minimum configured value. The aggregate of the minimum configured values cannot exceed the speed of the Ethernet interface, which is 1 Gbps for GbE ports or 10 Gbps for 10 GbE ports.

Compression

Compression is recommended in every type of architecture, including those built for RDR/S. There are three modes of compression besides off:

- Mode 1: Brocade optimized Lempel–Ziv (LZ), is a hardware-implemented compression algorithm that is suitable for synchronous applications because it adds a mere 10 μ s of added latency. In addition, Brocade LZ can accommodate the maximum ingress rate for which the 7800/FX has been built, so it is line rate and poses no bottleneck for ingress traffic. LZ typically gets about a 2:1 compression ratio.

- Mode 2: Brocade optimized Dynamic Huffman Coding, is a software with hardware assist compression algorithm. Software-based algorithms are not suitable for synchronous applications, because they add too much processing latency. Brocade Dynamic Huffman Coding can accommodate up to 8-Gbps ingress from the FC side. For the Brocade 7800, that means 8 Gbps for the entire box. For the FX blade, that means 8 Gbps for each FCIP complex, of which there are two, one for each 10GE interface. The 10GE interfaces belong to the complex for 10GE interface 1 (XGE1). Mode 2 has been designed to work efficiently with an OC-48 WAN connection. Mode 2 typically gets about a 2.5:1 compression ratio.
- Mode 3: Deflate, also known as GZIP, is entirely a software-based algorithm and not suitable for synchronous applications. Deflate takes the tradeoff between compression ratio and compression rate further. The maximum rate per FCIP complex is 2.5-Gbps ingress from the FC side. Mode 3 has been designed to work efficiently with an OC-12 WAN connection. Mode 3 typically gets about a 4:1 compression ratio.

6 – Performance and Tuning Optimization

When you finish this section you should be able to perform the following tasks:

- Given a performance scenario, determine an appropriate solution
- Describe strategies for maximizing throughput in a fabric

Trunking Overview

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 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.

Masterless Trunking

Masterless trunking means that if the master port goes offline, one of the slave ports automatically becomes the new master port, thus avoiding traffic disruption. The new master port uses the old master port area and the old master port is assigned a new, unused area. In this way, the PID of the trunk does not change if the master port goes offline. If trunking is not masterless, and if the master port goes offline, traffic disruption can occur because the slave ports in the trunk group go offline to select the new master port and then come back online.

Masterless trunking is supported for most platforms and trunking types:

- All F_Port trunking is masterless.
- ISL and ICL trunking is masterless.
- EX_Port trunking is masterless, except for the following:
 - Enterprise-class platforms with VF disabled.

Requirements for Trunk Groups

The following requirements apply to all types of trunking:

- The Trunking license must be installed on every switch that participates in trunking.
- All of the ports in a trunk group must belong to the same port group.
- All of the ports in a trunk group must be running at the same speed.
- All of the ports in a trunk group must be configured for the same distance.
- All of the ports in a trunk group must have the same encryption, compression, QoS, and FEC settings.
- Trunk groups must be between Brocade switches (or Brocade adapters, in the case of F_Port trunking). Brocade trunking is proprietary and not supported on M-EOS or third-party switches.
- There must be a direct connection between participating switches.

- Trunking cannot be done if ports are in ISL R_RDY mode. (You can disable this mode using the `portCfgIslMode` command.)
- Trunking is supported only on FC ports. Virtual FC ports (VE_ or VEX_Ports) do not support trunking.

F_Port Trunking for Access Gateway

You can configure trunking between the F_Ports on an edge switch and the N_Ports on an Access Gateway module.

F_Port trunking keeps F_Ports from becoming disabled when they are mapped to an N_Port on a switch in Access Gateway mode. With F_Port trunking, any link within a trunk can go offline or become disabled, but the trunk remains fully functional and there are no reconfiguration requirements.

EX_Port Trunking

You can configure EX_Ports to use trunking just as you do regular E_Ports. EX_Port trunking support is designed to provide the best utilization and balance of frames transmitted on each link between the FC router and the edge fabric. You should trunk all ports connected to the same edge fabrics.

After initiation, the first port from the trunk group that comes online is designated as the master port. The other ports that come online on the trunk group are considered the slave ports. Adding or removing a slave port does not cause frame drop; however, removing a slave port causes the loss of frames in transit.

The restrictions for EX_Port frame trunking are the same as for E_Ports—all the ports must be adjacent to each other using the clearly marked groups on the front of the product.

Configuring Encryption and Compression

On a given ISL between two 16 Gbps E_Ports, you can configure each port for encryption, compression, or both. Your encryption and compression settings must match at either end of the ISL. Port segmentation will occur during port initialization if these configurations do not match. Before configuring a port for encryption, you must configure the port for authentication using the `authUtil` and `secAuthSecret` commands:

- Use the `authutil` command to enable switch authentication, enable the DH-CHAP authentication protocol for ports that support encryption, and select the appropriate DH (Diffie-Hellman) group (4 or “*”).
- Use the `secauthsecret` command to configure a pre-shared secret on both sides of the ISL for all ports configured for in-flight encryption. A secret of at least 32 characters is recommended. Maximum is 40 characters.

Port segmentation will occur during port initialization if authentication fails.

If you need to disable authentication on a port that has encryption or compression configured, you must first disable encryption or compression on the port, and then disable authentication.

Fibre Channel Link Lengths and Loss Budget

The fiber optic link loss budget, also known as the “channel insertion loss,” “link margin,” or “power budget” of the link, is a measure of signal power loss expressed in decibels (dB). Link loss is a combination of fiber attenuation related to the distance of the link and the connectors or splices in the link. The length of Fibre Channel links is explicitly defined in FC-PI-4 and FC-PI-5. Multimode links are defined with a maximum link length and 1.5 dB of connector loss. Single-mode links have been defined with a maximum length and 2.0 dB of connector loss. [Table 2](#) illustrates the standard link lengths for various types of fibers and speeds.

TABLE 2 Standard Link Lengths

FC Link Speed	OM2 Link Length (m)	OM3 Link Length (m)	OM4 Link Length (m)	Single-mode Link Length (m)
1 Gbps	500	860	N/A ¹	10,000+
2 Gbps	300	500	N/A ¹	10,000+
4 Gbps	150	380	400	10,000+
8 Gbps	50	150	190	10,000+
10 Gbps	82	300	N/A ¹	10,000+
16 Gbps	35	100	125	10,000+

1. The OM4 was standardized after 1, 2, and 10 Gbps FC were released.

The different lengths cause different link loss budgets. For a given fiber type and link speed, the standards define these link loss budgets as shown in [Table 3](#). OM4 fiber did not exist when 1, 2 and 10 Gbps FC were defined.

TABLE 3 Link Loss Budget

FC Link Speed	OM2 Link Length (m)	OM3 Link Length (m)	OM4 Link Length (m)	Single-mode Link Length (m)
1 Gbps	3.85	4.62	N/A	7.8
2 Gbps	2.62	3.31	N/A	7.8
4 Gbps	2.06	2.88	2.95	7.8
8 Gbps	1.68	2.04	2.19	6.4
10 Gbps	1.8	2.6	N/A	9.4
16 Gbps	1.63	1.86	1.95	6.4

Measuring Fiber Optic Link Loss

If you want to find the actual link loss, two types of devices can measure the link loss. An Optical Time Domain Reflectometer (OTDR) accurately measures the loss in a given link in dB. An OTDR can measure loss across an end-to-end link or across each cabling section. It is also used to locate cable cuts or a section of degraded cable within a link.

Sources of Congestion

Frame congestion is primarily caused by latencies somewhere in the SAN—usually storage devices and occasionally hosts. These latencies cause frames to be held in ASICs and reduce the number of buffer credits available to all flows traversing that ASIC. The congestion will back up from the source of the latency to the other side of the connection and start clogging up the fabric. This creates what is called back pressure. Back pressure can be created from the original source of the latency to the other side and all the way back (through other possible paths across the fabric, to the original source again). Once this situation arises the fabric is very vulnerable to severe performance problems.

Sources of high latencies include:

- Storage devices that are not optimized or where performance has deteriorated over time.
- Distance links where the number of allocated buffers has been miscalculated or where the average frame sizes of the flows traversing the links has changed over time.
- Hosts where the application performance has deteriorated to the point that the host can no longer respond to incoming frames in a sufficiently timely manner.

Other contributors to frame congestion include behaviors where short frames are generated in large numbers such as:

- Clustering software that verifies the integrity of attached storage.
- Clustering software that uses control techniques such as SCSI RESERVE/RELEASE to serialize access to shared file systems.
- Host-based mirroring software that routinely sends SCSI control frames for mirror integrity checks.
- Virtualizing environments, both workload and storage, that use in-band Fibre Channel for other control purposes.

Mitigating Congestion

Frame congestion cannot be corrected in the fabric. Devices exhibiting high latencies, whether servers or storage arrays, must be examined and the source of poor performance eliminated. Since these are the major sources of frame congestion, eliminating them will typically address the vast majority of cases of frame congestion in fabrics.

Topologies

A typical SAN design comprises devices on the edge of the network, switches in the core of the network, and the cabling that connects it all together. Topology is usually described in terms of how the switches are interconnected, such as ring, core-edge, and edge-core-edge or fully meshed. At this point the focus is on switch topology with ISLs—device connectivity is discussed in later sections. The recommended SAN topology to optimize performance, management, and scalability is a tiered, core-edge topology (sometimes called core-edge or tiered core edge). This approach provides good performance without unnecessary interconnections. At a high level, the tiered topology has a large number of edge switches used for device connectivity, and a smaller number of core switches used for routing traffic between the edge switches, as shown in [Figure 8](#).




FIGURE 8 Four scenarios of tiered network topologies (hops shown in heavier, orange connections)

The difference between these four scenarios is device placement (where devices are attached to the network) and the associated traffic flow.

- **Scenario A** has localized traffic, which can have small performance advantages but does not provide ease of scalability or manageability.
- **Scenario B**, also called edge-core, separates the storage and servers, thus providing ease of management and moderate scalability.

The edge-core topology ([Figure 8](#)—Scenario B) places initiators (servers) on the edge tier and storage (targets) on the core tier. Since the servers and storage are on different switches, this topology provides ease of management as well as good performance, with most traffic only traversing one hop from the edge to the core. (Storage-to-storage traffic is two hops if the second storage is on another core switch], but the two cores can be connected if fabrics are redundant.) The disadvantage to this design is that the storage and core connections are in contention for expansion. In other words, this topology allows for only minimal growth.

- **Scenario C**, also known as edge-core-edge, has both storage and servers on edge switches, which provides ease of management and is much more scalable.

The edge-core-edge topology (in [Figure 8](#)—Scenario C) places initiators on one edge tier and storage on another edge tier, leaving the core for switch interconnections or connecting devices with network-wide scope, such as Dense Wavelength Division Multiplexers (DWDMs), inter-fabric routers, storage virtualizers, tape libraries, and encryption engines.

Since servers and storage are on different switches, this design enables independent scaling of compute and storage resources, ease of management, and optimal performance—with traffic traversing only two hops from the edge through the core to the other edge. In addition, it provides an easy path for expansion as ports and switches can readily be added to the appropriate tier as needed.

- **Scenario D** is a *full-mesh* topology, and server to storage is no more than one hop. Designing with ICLs is an efficient way to save front-end ports, and users can easily build a large (for example, 1536-port) fabric with minimal SAN design considerations.

A full-mesh topology (Figure 8—Scenario D) allows you to place servers and storage anywhere, since the communication between source to destination is no more than one hop. With optical ICLs on the Brocade DCX 8510, customers can build a full-mesh topology that is scalable and cost effective compared to the previous generation of SAN products.



Note

Hop count is not a concern if the total switching latency is less than the disk I/O timeout value.

7 – Migration, Integration, and Validation

When you finish this section you should be able to perform the following tasks:

- Given an existing fabric, identify migration strategies to upgrade the fabric with new technology
- Given a set of existing fabrics and network devices, determine a consolidation or upgrade plan that minimizes disruption
- Develop a plan to validate a design

Migration Strategies

There are several main migration strategies that most projects follow:

- Simple replacement of existing switches
- Integration or consolidation of fabrics

Using this method, the existing fabrics are supplemented with new switches and Directors. Host and storage connections are migrated to the new switching elements within the same fabrics over time and the old switching elements are removed when they are no longer needed (i.e. no longer contain any active connections). This method may be more complex than the “parallel fabric” method since the old fabrics must be compatible with the new switching elements. This may require the old switches/Directors to first be upgraded (firmware) and/or modified (e.g. interoperability mode and other switch settings) prior to the new switches/Directors being added. However, in this model, since the same physical fabrics are being used, host and storage ports may not have to be moved in coordinated groups.

- Dual fabrics (rip-and-replace)

Using this method, new fabrics are built using new switches and Directors. These new fabrics are independent from the old fabrics (i.e. the fabrics are not merged) and host and storage connections are migrated to the new fabrics over time. This method tends to be more straightforward and less risky since it does not require new Fibre Channel switches and Directors to be added to the existing fabrics. However, in order for hosts to have access to the target storage devices in the new fabrics, hosts and storage must be moved.

- Dual fabrics and using FC Routing

If the parallel fabric method is used, the old and new fabrics may be interconnected via Fibre Channel routing technology. This allows selected resources within each fabric to have access to selected resources in the other fabrics. This method leverages the advantages of the parallel fabric method (separate fabrics) with those of the integrated fabric method (more flexibility in when specific host and storage connections are moved).

Use Case for FC-FC Routing

The primary reasons for using FC-FC Routing are as follows:

- A limited number of LUNs shared between fabrics
- A limited number of servers that need to share LUNs between fabrics
- Share archiving devices like tape libraries
- The migration of legacy M-EOS or Brocade FOS fabrics to current Brocade FOS-based platforms
- OEM support
- Security separation in a managed services environment

Integrating and Consolidating Fabrics

- New switches are integrated into the existing fabrics and fabrics merge
- Hosts and storage can be moved independently
- Benefits:
 - Allows existing host-to-switch cabling to remain in place
 - Allows many of the existing SAN components to remain in place
 - Extends the useful life of existing fabrics
 - No “new” fabrics to move to
- Considerations:
 - High level of effort and scheduling
 - Global impact/disruptive changes
 - More need for risk management
 - Difficult to test new architecture ahead of time

The tasks for merge planning include:

- Obtain accurate SAN configuration details for the existing environment. Storage administrators need this information for obvious connectivity reasons, but also to determine the impact of any Core PID, Fabric OS, or Domain ID changes. Domain ID and PID changes are especially relevant for the HP-UX servers, since a change could require the storage to be remapped on these servers.
- Schedule and coordinate activity windows for multiple hosts and fabrics. Because this option would likely result in maintenance windows that impact entire fabrics, storage administrators would need to coordinate activity windows that impact multiple servers possibly owned by multiple business units.

Zoning Overview

Zoning is a fabric-based service that enables you to partition your storage area network (SAN) into logical groups of devices that can access each other.

For example, you can partition your SAN into two zones, winzone and unixzone, so that your Windows servers and storage do not interact with your UNIX servers and storage. You can use zones to logically consolidate equipment for efficiency or to facilitate time-sensitive functions; for example, you can create a temporary zone to back up nonmember devices.

A device in a zone can communicate only with other devices connected to the fabric within the same zone. A device not included in the zone is not available to members of that zone. When zoning is enabled, devices that are not included in any zone configuration are inaccessible to all other devices in the fabric.

Zone Merging

When a new switch is added to the fabric, it automatically takes on the zone configuration information from the fabric.

If you are adding a switch that is already configured for zoning, clear the zone configuration on that switch before connecting it to the zoned fabric.

Adding a new fabric that has no zone configuration information to an existing fabric is very similar to adding a new switch. All switches in the new fabric inherit the zone configuration data. If the existing fabric has an effective zone configuration, then the same configuration becomes the effective configuration for the new switches.

Upgrading Firmware

Upgrading means installing a newer version of firmware. Downgrading means installing an older version of firmware. In most cases, you will be upgrading firmware; that is, installing a newer firmware version than the one you are currently running.



Note

Ensure all new features are working correctly prior to upgrading the next switch in your fabric.

Connected switches

Before you upgrade the firmware on your switch, you need to check the connected switches to ensure compatibility and that any older versions are supported.

Port Decommissioning

Fabric OS 7.0.0 provides an automated mechanism to remove an E_Port or E_Port trunk port from use. This feature identifies the target port and communicates the intention to decommission the port to those systems within the fabric affected by the action. Each affected system can agree or disagree with the action, and these responses are automatically collected before a port is decommissioned.

All members of a trunk group must have an equal link cost value in order for any of the members to be decommissioned. If any member of a trunk group does not have an equal cost, requests to decommission a trunk member will fail and an error reminding the caller of this requirement is produced.

Following are restrictions of port decommissioning:

- The local switch and the remote switch on the other end of the E_Port must both be running Fabric OS 7.0.0 or later.
- Port decommissioning is not supported on links configured for encryption or compression.
- Port decommissioning is not supported on ports with DWDM, CWDM, or TDM.
- Port decommissioning requires that the lossless feature is enabled on both the local switch and the remote switch.

Access Gateway Supported Hardware

All Fabric OS switches must be running Fabric OS v6.1.0 or later; all M-EOS switches must be running M-EOSc 9.1 or later, M-EOSn must be running 9.6.2 or later, and Cisco switches with SAN OS must be running 3.0 (1) and 3.1 (1) or later.

Fabric OS v7.0.0 supports the following Brocade hardware platforms for Access Gateway:

- Brocade 300
- Brocade 5100
- Brocade M5424
- Brocade 5450
- Brocade 5460
- Brocade 5470
- Brocade 5480
- Brocade 6510
- Brocade 8000
- Brocade VA40-FC

8 – Security

When you finish this section you should be able to perform the following tasks:

- Identify requirements for restricting which switches/devices may join a fabric
- Identify security features that restrict administrative access to a switch

Device and Switch Access

There are four ACL policies that restrict switch and device access to a fabric:

- Switch Connection Control (SCC) policy restricts which switches may join a fabric
The switch connection control (SCC) policy is used to restrict which switches can join the fabric. Switches are checked against the policy each time an E_Port-to-E_Port connection is made.
- Device Connection Control (DCC) policy restricts which FC devices can connect to which FC switch ports
Multiple DCC policies can be used to restrict which device ports can connect to which switch ports. The devices can be initiators, targets, or intermediate devices such as SCSI routers and loop hubs. By default, all device ports are allowed to connect to all switch ports; no DCC policies exist until they are created.
- Advanced Device Security (ADS) policy restricts device access to the fabric at the Access Gateway
Advanced Device Security (ADS) is a security policy that restricts access to the fabric at the AG level to a set of authorized devices. Unauthorized access is rejected and the system logs a RASLOG message. You can configure the list of allowed devices for each F_Port by specifying their Port WWN (PWWN). The ADS policy secures virtual and physical connections to the SAN.
- Fabric Element Authentication (AUTH) policy requires shared secrets and digital certificates to authenticate devices or switches
The authentication (AUTH) policy allows you to configure DH-CHAP authentication on switches. By default the policy is set to PASSIVE and you can change the policy. All changes to the AUTH policy take effect during the next authentication request. This includes starting authentication on all E_Ports on the local switch if the policy is changed to ON or ACTIVE, and clearing the authentication if the policy is changed to OFF. The authentication configurations will be effective only on subsequent E_ and F_Port initialization.

Administrative Access

There are three ACL policies that restrict administrative access to a fabric:

- Fabric Configuration Server (FCS) policy restricts which switches may change the configuration of the fabric.

The Fabric Configuration Server (FCS) policy identifies a specific switch (the Primary FCS) from which fabric configuration must be performed, as well as a set of switches (the Backup FCSes) that may replace the Primary FCS should it leave the fabric.

- IP Filter (IPFilter) policy restricts incoming administrative traffic on the IP management interfaces.

The IP Filter (IPFilter) 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.

- Password (PWD) policy distributes the user account and password database to other switches in the fabric.

The Password (PWD) policy allows the user accounts and password database from the local switch to be distributed to other switches in the fabric.

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