BCAF in a Nutshell Study Guide for Exam 143-120

Exam Preparation Materials

Revision March, 2011

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

Revision: March, 2011

BCAF in a Nutshell Third Edition

Objective: The BCAF Nutshell guide is designed to help you prepare for the BCAF Certification, exam number 143-120.

Audience: The BCAF Nutshell self-study guide is intended for those who have successfully completed the CAF 200 Design and Implementation of a FICON Environment course, and who wish to undertake self-study or review activities before taking the actual BCAF exam. The BCAF guide is not intended as a substitute for classroom training or hands-on time with Brocade products.

How to make the most of the BCAF guide: The BCAF guide summarizes the key topics on the BCAF 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 - CAF 101-WBT BCAF Knowledge Assessment. To benefit from the BCAF guide, we strongly recommend you have successfully completed the CAF 200 Design and Implementation of a FICON Environment course.

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

Joe Cannata
Senior Manager of Business Development

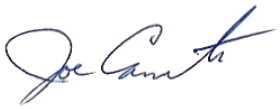

Table of Contents

List of Figures

1 - Architecture

System z Host, I/O Channel Subsystem and FICON Protocol Architecture

Mainframe Overview

- A mainframe is one machine that is a central electronics complex (CEC)
- A CEC will have many Central Processing Units (CPUs) in it
- A CEC can support logical partitions (LPARs)
- System z can support up to 60 LPARs
- CPUs are allocated to LPARs dynamically
- Special CPUs are called Service Assist Processors (SAPs)
- Each LPAR can run a different Operating System such as z/OS and Linux

Mainframe is simply an industry term for a large computer. The majority of mainframes today are manufactured by IBM. The mainframe has historically been associated with centralized rather than distributed computing but IBM refers to its mainframes as a large server. These mainframes can be used to serve distributed users and smaller servers in a computing environment. This is made possible with capabilities of these to run thousands of virtual Linux images connected via Fibre Channel to storage.

Mainframe Models: zSeries to System z

- In 2000, IBM launched the zSeries 900 line of servers which were referred to as the zSeries and included models 800, 890, 900, and 990
- In 2000, IBM renamed the System/390 to IBM eServer zSeries
- In 2005, IBM launched the z9 series which included IBM z9 Enterprise Class and Business Class
- In 2008, the z10 Enterprise Class was launched
- Since FICON has been available since the late 1990s, the 9672 G5/G6 was the first mainframe to provide support for FICON
- zSeries, System z including z9 and z10
  - 800, 890 (midrange enterprise)
  - 900, 990 (enterprise, large enterprise; business class UNIX)
  - z9, z10 (enterprise, large enterprise; business class UNIX)
  - zOS (z Operating System)
  - zVM (z Virtual Machines)
  - zTPF (z Transaction Processing Facility)
  - Linux for System z
- Elements that affect unrepeated distance
  - The type of fiber used
  - The port transceiver type

Figure 1: IBM System z10 Mainframe

What is FICON?

- FIbre CONnection (FICON) is:
  - An I/O protocol based on Fibre Channel (FC-SB-2 and FC-SB-3) used with IBM (and compatible) mainframes and storage
  - A layered protocol based on industry standards for Fibre Channel (FC) architecture using FC layer 4.
  - The successor to ESCON
- Introduced in the late 1990s, the original FICON adapter card which enabled FICON was developed for the 9672 G5 and G6 mainframes

FICON Frame Encoding

- FICON frame encoding can be either FC standard 8b/10b used for 1, 2, 4 and 8 Gbps or 64b/66b with 10 Gbps ports
- Two types of frame encoding:
  - 8b/10b encoding
  - 64b/66b encoding
- To use 64b/66b encoding, you must use a 10 Gbps ISL which cannot auto-negotiate with 1, 2, 4 and 8 Gbps ports

FICON Operating Modes

- There are two FICON operating modes: FCV and FC
  - For System zSeries and 9672 G5 and G6 servers, there were two modes supported:
    - FICON Bridge – referred to as FCV which is a FICON channel mode designed to enable access to ESCON interfaces using the FICON Bridge Adapter in the ESCON Director (9032-5)
    - FICON Native – referred to as FC which enables access to the FICON channel mode from native FICON control units in a point-to-point mode through a FICON Director (No cascaded FICON)
  - For zSeries and System z including z9 and z10, support was added for:
    - FICON Native (FC) but not FICON Bridge
    - Supported toplogies: point-to-point, switched point-to-point and cascaded FICON Directors
    - Fibre Channel Protocol (FCP) which provided support for open systems via industry standard SCSI devices

Mainframe Channels - Channel Subsystem

- The key to moving data into and out of a mainframe is the channel subsystem (CSS)
- When an I/O operation is required, the CSS is passed the request from the main processor to manage. While awaiting completion of an I/O request from the CSS, the main processor can continue processing other data
- This process supports concurrent transactions because of the interaction between the host processor and the CSS

Mainframe Channels - Logical Channel Subsystem

- The z/OS operating system is limited to a maximum of 256 channel paths (CHPIDs) for use within a logical partition (LPAR)
- To facilitate the usage of more CHPIDs, the mainframe architecture supports a logical channel subsystem (LCSS)
- The LCSS is functionally identical to the channel subsystem but up to four LCSSs can be defined within a central processor complex. CHPIDs are unique within the LCSS only; consequently, the 256 CHPID limitation can be overcome
- Usage can be maximized with NPIV and good fan-in/fan-out ratios

FICON Connection Components

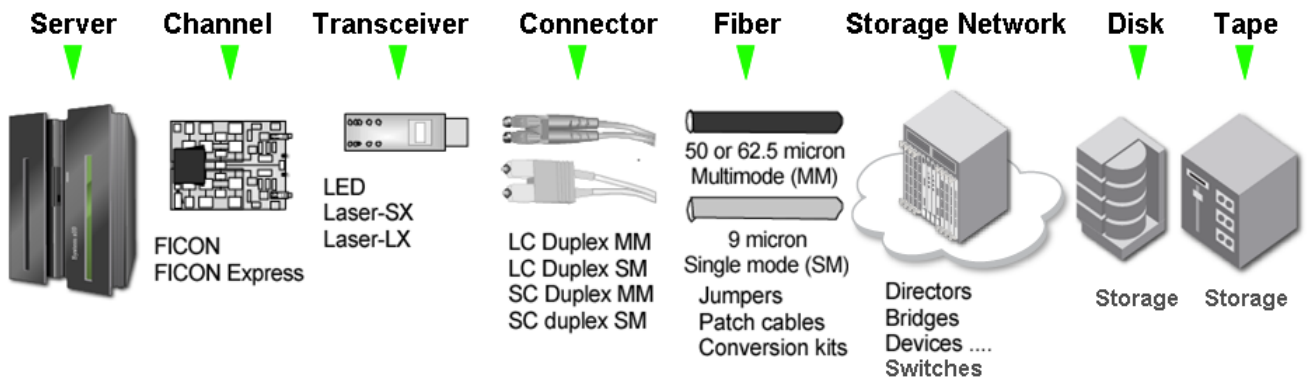

Figure 2: FICON Connection Components

FICON Express Cards

- Original FICON card:
  - 100 MB/sec full duplex
  - LX: supports 9 micron SM fiber
  - SX supports 50, 62.5 MM fiber
  - FICON Express2 (LX):
  - 200 MB/sec full duplex
  - Supports 9 micron SM fiber
- FICON Express2 (SX):
  - 200 MB/sec full duplex
  - Supports 50, 62.5 micron MM fiber
  - Supported on all z9 processors
- FICON Express4 (LX and SX):
  - 400 MB/sec full duplex (the IBM certified maximum data rate capability for FICON is 350 MB/sec)
  - Supports 9, 50, 62.5 micron SM and MM fiber
  - Supported on System z9 and z10 servers
- FICON Express4-2C adapters are supported on System z9 BC servers only, and go 550m or 4km

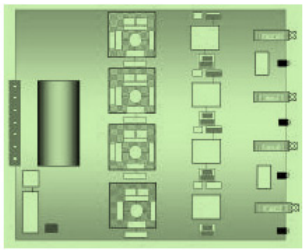

FICON Express2

- z10, z9, z990, z890
- Longwave (LX) to 10km
- Shortwave (SX)
- 1 or 2 Gbps link rate
- 270MBps FD max thru-put
- 68% of 400MBps potential
- 107 Buffer Credits per port
  - 108km @ 2G full frame / port
  - 54km @ 2G half frame / port
  - 43km @ 2G 819 byte payloads

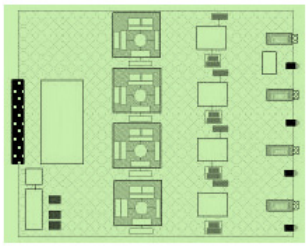

FICON Express4

- z10, z9
- 4km & 10km LX
- Shortwave (SX)
- 1, 2 or 4 Gbps link rate
- 520MBps FD max thru-put
- 65% of 800MBps potential
- 200 Buffer Credits per port
  - 101km @ 4G full frame / port
  - 51km @ 4G half frame / port
  - 40km @ 4G 819 byte payloads

FICON Express4 provides the last native 1Gbps CHPID support

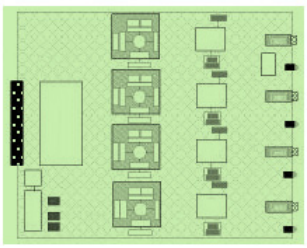

FICON Express8

- z10
- 10km LX
- Shortwave (SX)
- 2, 4 or 8 Gbps link rate
- 740MBps FD max thru-put
- 46% of 1600MBps potential
- 40 Buffer Credits per port
  - 10km @ 8G full frame / port
  - 5km @ 8G half frame / port
  - 4km @ 8G 819 byte payloads

FICON switching devices will provide BCs for long distances

Figure 3: Mainframe Channel Cards

- FICON Express 2, 4 and 8 adapters do not support CHIPID type FCV. (FCV was used with the FICON Bridge solution.)
- FICON Express8 (LX and SX):
  - 800 MB/sec full duplex.
  - Supports 9 micron SM fiber and 50, 62.5 micron MM fiber.
  - Supported only on System z10 servers.
  - Support for FCV and FCP modes.
  - Direct memory access using the QDIO architecture.
- FICON Express8 is the most current FICON card

The maximum distance for a single Director is 10km from the host mainframe. If you are trying to determine the maximum distance of a device attached to the Director to the mainframe, you need to add in that distance as well. For example, a Director could be attached to a mainframe that is 5km away and a storage device is attached to the Director which is also 5km away. The total distance then that the storage device can be from the mainframe is 10km.

Reasons to Migrate from ESCON

- FICON is an improved protocol over the ESCON protocol
  - ESCON provides a single I/O per channel whereas FICON uses multiple, simultaneous ESCON-type I/Os
  - ESCON I/O is half duplex; FICON I/O is full duplex
  - FICON provides improved support over distance decreased “data droop”
  - FICON offer channel path consilidation
  - FICON provides increased addressing
  - FICON offers a bandwidth increase
  - Buffer credit control
  - FICON adapters can support Open Systems and LINUX (FICON/FCP intermix)
  - zHPF support
- FICON uses Open Systems components
  - Fibre Channel Directors
  - Fibre Channel cables & connectors
  - Fibre Channel adapters
- Migration Considerations:
  - Buffer credits
  - SX/LX requirements as applicable

System z Channel I/O Configuration

Mainframe Physical Channel IDs (PCHIDs)

- Any CHPID can be assigned to any sub channel system location
- Because the physical channels may be associated to CHPIDs on different channel subsystems, the PCHID was created for zSeries and System z mainframes to enable the mapping to the physical channel to the CHPIDs on different channel subsystems
- PCHID is required to support LCSS
  - Used to map a CHIPD to a PCHID when using LCSS
  - PCHIDs are unique and use a three digit value for identification
  - PCHIDs must be defined in HCD (IOCP) and stored in IODF

Up to 256 CHPIDs can be assigned to a Channel Subsystem (CSS) and up to 256 CHPIDs can be assigned to a Logical Channel Subsystem (LCSS) as well. Since there can be up to four LCSSs, a total of 1024 CHPIDs can be supported. HCD is used with the CHPID mapping tool to connect a PCHID to a CHPID.

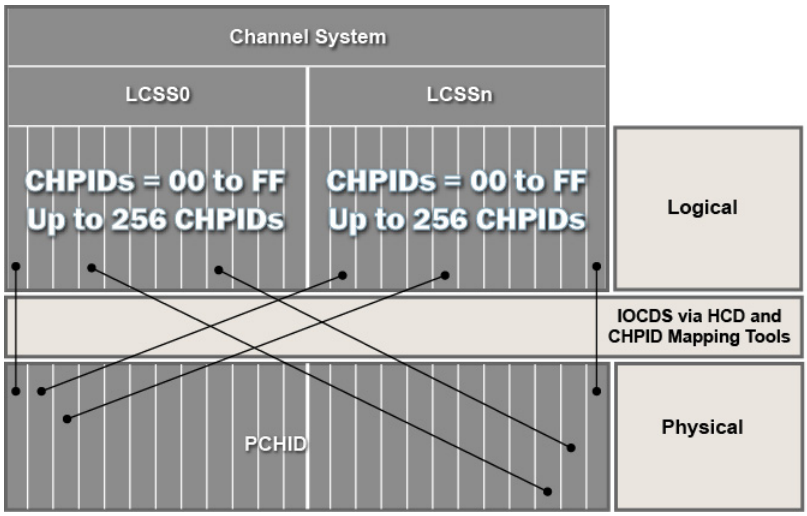

Figure 4: Mapping PCHIDs to CHPIDs

PCHIDs from LCSS can be applied to different LCSSs and thus for determining paths, you need to be able to connect the CHPID to the PCHID. (Note, LCSS0 is the first LCSS; LCSSn represents LCSS 1, 2, or 3.)

A CHPID is a logical number defined within the IOCDS; the mapping tool ties that logical CHPID number to a PCHID location in the machine, such as a FICON Interface card port. Thus if there is a problem transmitting data, you will want to know the relationship between the CHPID and the PCHID to determine the path. For more information, review the *System z CHPID Mapping Tool User's Guide* GC28-6825.

Multiple Image Facility (MIF)

MIF enables resource sharing across logical partitions within a single CSS or across CSSs. When a channel is shared across logical partitions in different CSSs, it is called "spanning."

With multiple CSSs, the IOCDS logical partition MIF Image ID is not unique. Therefore, the logical partition identifier value provides a unique value for each logical partition within the same server. The following terminology applies:

Logical partition identifier

The logical partition identifier is a number in the range of 00 to 3F and is assigned in the image profile through the Support Element (SE) or the Hardware Management Console (HMC). It is unique across the server and can also be referred to as the User Logical Partition ID (UPID). It is suggested to establish a numbering convention for the logical partition identifiers.

Logical partition name

The logical partition name is defined through HCD or IOCP. It is the name in the RESOURCE statement in the IOCP. The logical partition names must be unique across all CSSs.

MIF IDGbps

This identifier is defined through HCD or IOCP. This is the number defined in the RESOURCE statement in the IOCP. It is in the range of 1 to F and is unique within a CSS. It is not unique across multiple CSSs. Multiple CSSs can specify the same MIF ID (also known as Image ID (IID)).

- A uniprocessor system consists of a single central processor complex (CPC)
  - Consists of a single mainframe and all associated system hardware and software
  - Controlled by a single copy of the operating system
- A Sysplex has a standard communication mechanism (the cross-system coupling facility (XCF)
  - Used for z/OS system applications
  - Enables communication between application programs on one or multiple computers.

In a parallel sysplex decisions about which of these two transport mechanisms to use for routing a specific message are made dynamically. Within a single z/OS system messages are transported using cross-memory services, rather than being routed through either of the physical transport mechanisms. Applications join specific groups as individual members. On joining a group a member can send or receive messages. Individual messages are assigned to specific transport classes, based on the message's size. Each transport class owns input and output buffers. Routing decisions are made at the transport class level. Three technologies need for this are PSIFB, CFCC and a Coupling Facility

zHPF - High Performance FICON

- It is used to super-charge FICON CCW I/O performance by encapsulating multiple CCW commands into a single frame
- The FICON Express2, Express4 and Express8 features (CHPID type FC) support the native FICON protocol and the new zHPF protocol simultaneously
  - FICON Express is not supported for zHPF
- It is optimized for online transaction processing (OLTP) workloads for Media Manager such as DB2, VSAM and PDSE
  - The smaller the block size the more effective zHPF should be
- zHPF is exclusive to System z10
  - Exploitation is required by the DASD CU (IBM, HDS, EMC)
  - EMC only supports zHPF on VMAX (5874+ code release)
  - FICON Express cannot be in the path
  - There are operating system dependencies (z/OS 1.8+)
  - The code to support zHPF is in the latest driver for System z10 EC and in the "GA driver" for the z10 BC
- If an I/O is zHPF eligible, tHe Media Manager will create a TCW rather than a stream of CCWs
- zHPF eligible access methods:
  - DB2
  - PDSE
  - VSAM
  - zFS
  - HFS
  - EF SAM DS
- Supported Brocade switching devices
  - M6140
  - Mi10K
  - Brocade 48000
  - Brocade DCX-4S
  - Brocade DCX

System z I/O Channel and CUP Interdependencies

Modified Indirect Data Address Word Facility

- The Modified Indirect Data Address Word (MIDAW) facility is available in the IBM z9 and z10 processor to improve channel performance including FICON, especially when accessing IBM DB2® databases. This facility provides a method for gathering data into and scattering data from discontinuous storage locations during an I/O operation.
- The use of MIDAWs do not cause data to move quicker across the FICON link but MIDAW does reduce the number of frames and sequences flowing across the link and makes the channel more efficient by making the frames approximately 40% larger than if MIDAW was not used. For FICON, this means more information units are sent to the control unit in a single burst.
- The significant performance benefit of MIDAWs is achieved with Extended Format (EF) data sets which applications like DB2 use extensively
- No other requirements are necessary should you decide to use MIDAW in a zHPF environment

From a practical perspective, if MIDAW is installed on the host, then sustained bandwidth can be required on a FICON Director. Therefore, when configuring a FICON Director such as the M6140, you would set the port speed to sustained instead of burst negotiate. For further information, review the Implementation Guide: *Deploying Brocade Mi10K and M6140 Directors into FICON Environments* for more information. MIDAW may also improve buffer credit usage as there will be fewer frames sent.

FICON CUP Support

- CUP = Control Unit Port (an embedded not physical port)
  - Allows in-band management of Directors from the management applications on a mainframe
  - Used for configuration, monitoring, and error handling
- CUP support is provided by all Brocade certified FICON switches and Directors
  - Optional licensed feature for both B-Series and M-Series
  - FICON Management Server license required to activate CUP
  - Supported in single or cascaded FICON environments
  - Utilizes embedded port – FE
- CUP is typically used by customers with:
  - SA z/OS (System Automation software) configuration tool
  - RMF (Resource Measurement Facility) monitoring tool
- If a mainframe administrator wants to be able to do in-band management from their Hardware Management Consoles – (HMC), they need CUP
- If it is important to them to get the Service Information Messages (SIM) to the z/OS console for FICON device hardware failures
  - Interpret these messages using the IBM Mainframe Systems Messages and Codes manual
  - If there is a need to do RMF reports for reporting on FICON Director activity statistics in the Record Type 74 Subtype 7, CUP is mandatory
  - RMF option FCD/NOFCD in `ERBRMFxx` enables or disables FICON CUP Director

- Although IBM has developed a new machine type, *2027*, for FICON Directors, all Directors and switches, regardless of model or manufacture should be configured as *2032*
- Active=Saved will prevent a loss of connectivity on a Director using CUP

FICON Directors have the embedded port "FE" for the CUP port. On FICON Directors that have 256 or more ports, this has caused a slight dilemma. This logical "FE" overlaps the physical "FE" port, so the physical "FE" and "FF" ports cannot be used on these Directors for FICON connectivity. They may still be used for port swaps or for FICON/FCP intermix, however, they cannot be "genned" in HCD for FICON connectivity.

In a FICON environment, only one RMF LPAR should attempt to access the CUP port at any one time, and it is still best practice to have two or more CHPIDs with access to the CUP. However, too much activity to the FICON CUP can cause missing interrupts, leading to the potential for having a boxed device.

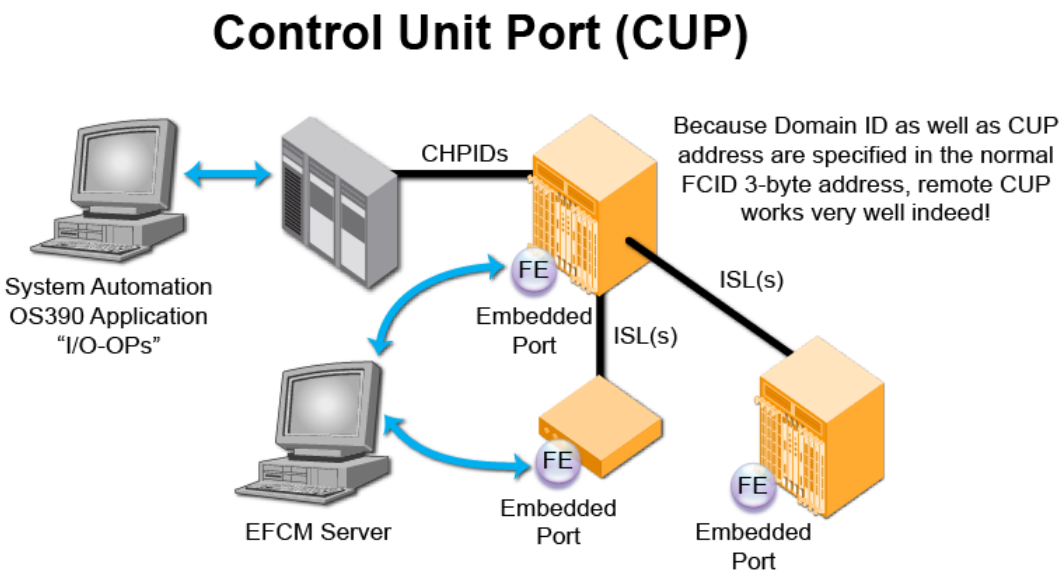

Figure 5: CUP

Footnote 1: Safe switching is the ability to manipulate ports and adjust path status non-disruptively. In order for System Automation I/O operations to provide safe switching, the System Automation I/O must have access to all switches. All the switches must be online as I/O devices on all the systems where System Automation is running in order to verify that a path can be removed before the path is actually removed.

This is a simplified representation of an FICON environment and how CUP fits into that environment. Best practice is to have two or more CHPIDs with access to the CUP and implement safe switching. Safe switching

describes the System Automation process of path removal. Safe switching is a process that involves contacting all systems in a sysplex and verifying a path can be removed before it is actually removed.

Features of CUP

- Remote switch functions via z/OS System Automation
- Reporting of failed FRUs to z/OS
- Needed to add Directors to the I/O sysgen for error reporting purposes
- May require the CUP device be taken offline before a firmware upgrade
- Prohibits communication between port pairs
- Allows the host to set the switch to an offline state

Prohibit Dynamic Connectivity Mask (PDCM)

- CUP also provides PDCM as a management tool. The PDCM is a mechanism to define port connectivity such as prohibits, allows, and blocking/unblocking of ports.
- The FICON Prohibit Dynamic Connectivity Mask (PDCM) controls whether or not communication between a pair of ports in the switch is prohibited or allowed
- Block versus Prohibit:
  - Blocking causes the firmware to send a continuous “offline” sequence to the port
    - Useful to report the link as inactive after varying a device off on the mainframe
  - Prohibit causes the firmware to “prevent” connectivity between the ports
    - Useful to force FICON traffic over specific ISLs
- PDCM controls whether or not communication between a pair of ports in the switch is prohibited or allowed
- If there are any differences in restrictions set up with Brocade Advanced Zoning and PDCM, the most restrictive rules are automatically applied

FICON I/Os

One I/O operation may be referred to as an Open Exchange, or an exchange pair.

2 - Design Considerations

Single Director vs. Cascaded Director Design Requirements

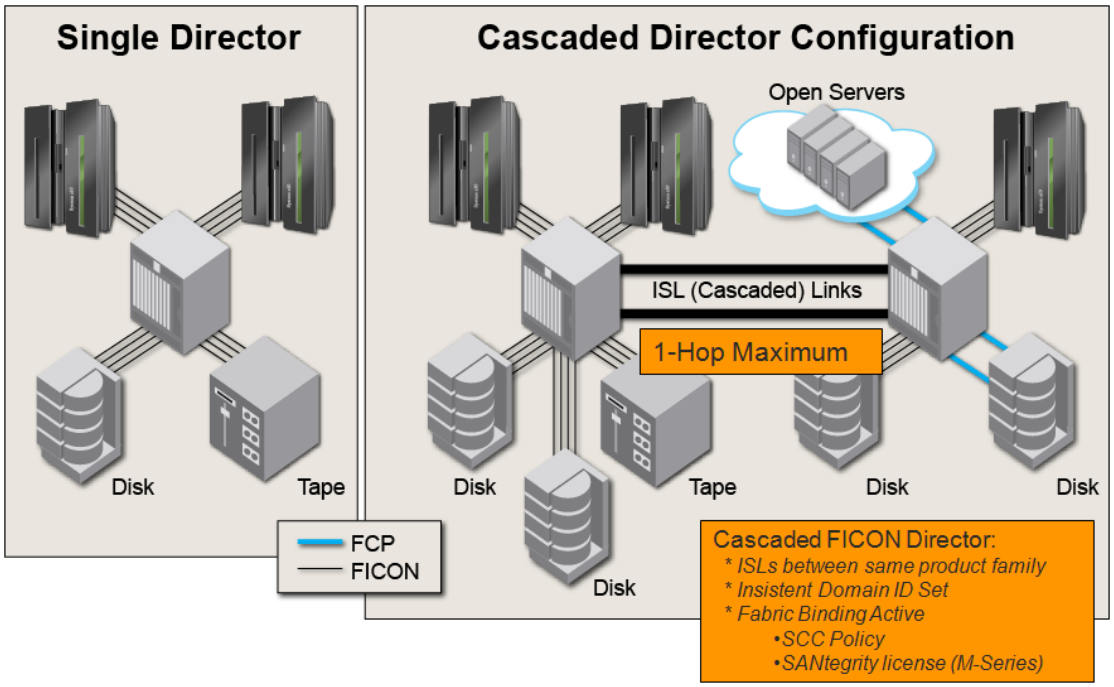

Figure 6: Single vs. Cascaded FICON Director Environments

A hop is defined as one ISL connection. More than one ISL connection is not supported between Directors. Therefore, you can have two Directors connected together through an ISL but not more than two Directors. In this example, the cascaded configuration shows how an open system server can also be included in the configuration as the open system server is attached to one Director that is attached to another Director.

- In open systems, connecting two Directors together is achieved by creating an Inter-Switch Links (ISLs) between the machines. In the FICON environment, this concept of connecting two Directors together is called cascading. M-Series Directors use preferred path.
- The standard FC routing mechanism for cascaded links is called Fabric Shortest Path First (FSPF).
- The FSPF protocol compares the cost of various paths between a source Director and a destination Director by adding the costs of the ISLs along each path. FSPF chooses the path with minimum cost. If multiple paths exist with the same minimum cost, FSPF distributes the load among these paths.
  - Database maintained of ISL and ICL link states
  - Reliable flooding is not the method of synchronization between two switches
- Required Fabric OS settings are `dlsreset` and `iodset`
  - `iodset` causes a delay to make sure frames have left the fabric
- On M-Series Directors. the rerouting delay can lead to I/O timeouts

FICON with Cascading

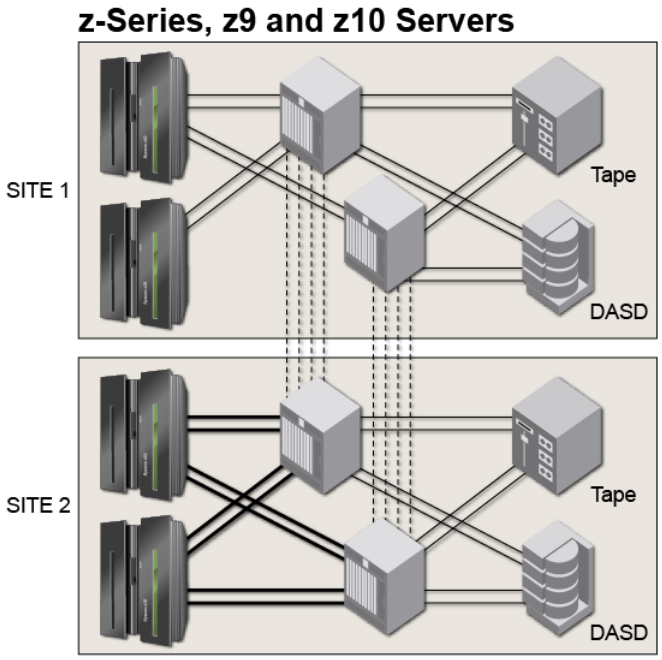

Figure 7: Multi-site FICON Environment

- High Availability FICON Redundant Fabrics as shown in the figure
- Single Director FICON Cascading
- Efficient cross-site connectivity
- Fewer CHIPID ports
- Less fiber cabling
- Easier to manage
- More scalable
- Lower cost of ownership

FCIP Support

Another way to avoid the multiple hop problemis to have a System z at a remote site and have each FICON Director attached to that mainframe. This might be done via dark fiber or potentially using FCIP blades or the Brocade 7500/7800 distance extension solutions. To support FCIP on the 7500 or 7800, the Advanced Extension license is required and provides:

- Slot-based
- FCIP Trunking
- Adaptive Rate Limiting

The 7800 offers two additional features not found in the 7500, lossless DPS and enhanced TI Zones.

2-byte Link Addressing

- ISLs are not used nor defined
- The first byte is the Domain ID (CHPID) and the second byte is the switch port (egress port)
- All CHPIDs with 2-byte link addresses require all links accessing the CHPID to be a 2-byte link
- 2-byte addressing enables cascaded switch support; the second byte enables a Domain ID to be specified in the IOCP definition
- Remember that intermixed IOCP definitions (1-byte and 2-byte link addresses) on the same channel path are not allowed
- When you create 2-byte FICON addresses in a mainframe configuration:
  - The mainframe channels will issue the Query Security Attributes CCW at initialization time (after the FLOGI) to see if the entry FICON Director supports what IBM commonly calls a "high integrity fabric“ (2-byte addressing not 1-byte addressing)
  - The channel expects the QSA Response to contain two bits indicating fabric binding and insistent Domain IDs (IDID) are configured and established throughout the entire fabric
  - If these bits come back false (either one), then the channel stops and terminates the login process
  - The QSA response will be affirmative:
    - IF and ONLY IF security is implemented on each Director in this fabric
    - Fabric Binding is enabled
    - Insistent Domain_ID is enabled
  - If any of this is NOT true, then the QSA command will indicate a "high integrity fabric" is not available, and the channel will not initialize
    - The only way to recover the channel at this point is to either use single byte addresses or turn on Fabric Binding and Insistent domain ID for every switch in the fabric

The QSA command will not be sent if 1-byte addressing is used for a FICON Director and is in the CNTLUNIT macro in the IOCDS. QSA is only used for cascaded-FICON or high integrity FICON. Anytime you see a 1-byte address for a Director, you know that Director is not part of a cascaded environment.

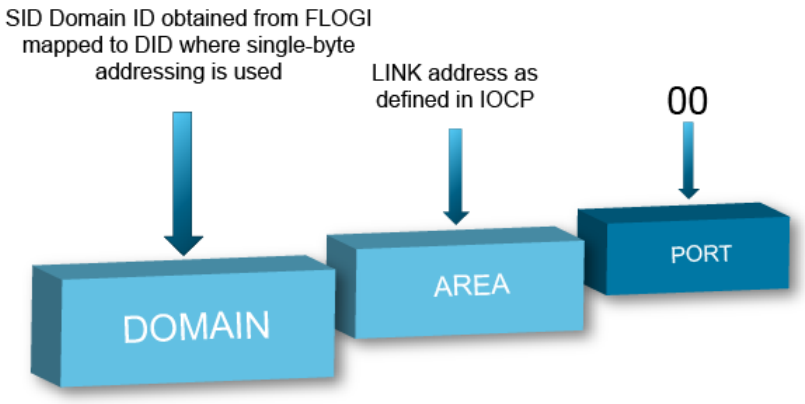

Figure 8: 2-Byte Link Addressing

Scalability Considerations

B-Series Trunking

- Trunking combines two or more physical ISLs into a single logical link, while limiting FSPF rerouting
- Trunking goals:
  - Reduce individual ISL congestion
  - Form a fault-tolerant high bandwidth logical ISL (called a trunk or trunk group) that withstands the failure of individual ISLs
- Trunk group characteristics:
  - Frames are multiplexed across ISLs in the trunk group
  - One port in the trunk group represents the link in the routing database
- ASICs preserve in-order delivery
- Trunking license required for all switches participating in trunking
- Available when the license is installed and ports are reinitialize
- Trunking is enabled by default; if previously disabled, it must be re-enabled (`portcfgtrunkport`) on the trunk ports
- Trunk ports must operate at a common port speed
- Trunk ports must originate and end in a valid port groups
- When trunking criteria is met the trunk forms automatically
- Long distance trunks require Extended Fabric licenses on both ends of the link

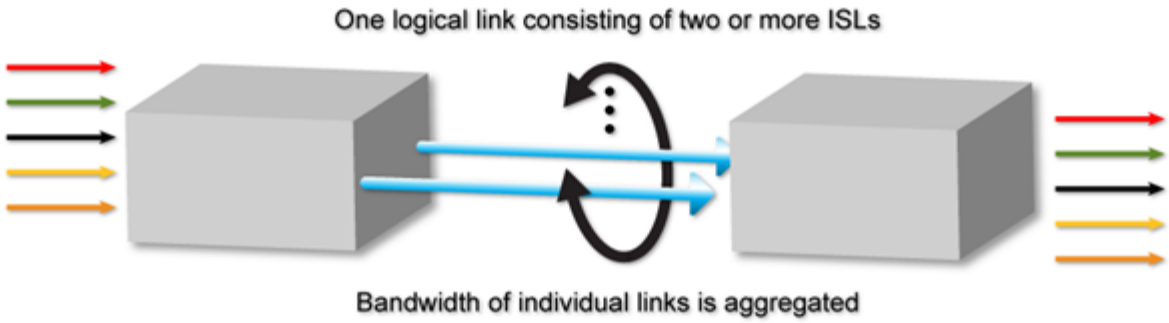

Figure 9: Sample B-Series Trunk Group

FICON Performance

Virtual Channels

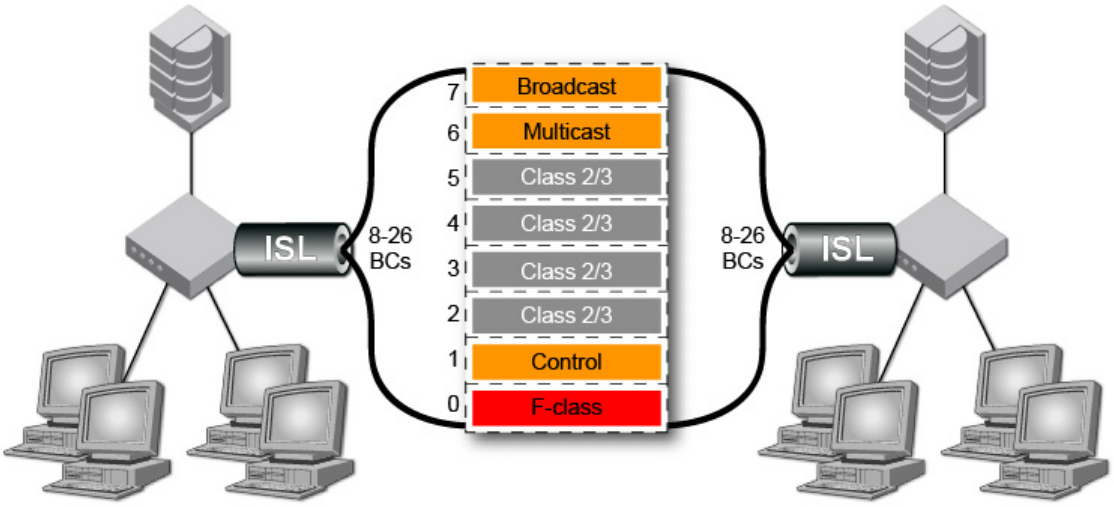

Figure 10: Virtual Channels on an ISL

- Virtual Channels split the buffer credits for an ISL into separate pools, one pool for F class traffic, some for control, and four pools for normal data traffic - effectively like four small-ISLs. Only 0-5 get buffer credits.
- Each VC has fewer buffer credits than total buffer credits for physical link; used for ISLs and ICLs
- You may need to turn this feature off if you go an extended distance as you could run out of buffer credits on a VC to drive a single device at full speed over long distances
- Reduces head of line (HOL) issues

HOL arises when data packets arrive at different input ports and are destined for the same output port. If the HOL packet of a cannot be switched to an output port because of contention, the rest of the packets in that buffer are blocked by that packet, the head of the line packet, even if there is no contention at the destination output ports for those packets.

High Availability

Buffer-to-Buffer Credits

- Define the maximum amount of data that can be sent prior to an acknowledgement
- Buffer credits are physical ASIC port or card memory resources and are finite in number as a function of cost
- Within a fabric, each port may have a different number of buffer credits
- The number of available buffer credits is communicated at fabric logon (FLOGI)
- One buffer credit allows a device to send a fixed 2112 byte frame of data (2K usable for z/OS data)
- Assuming that each credit is completely full, you need one credit for every 1km of link length over a 2 Gbps fiber

- Unfortunately, z/OS disk workload rarely produces full credits. For a 4K transfer, the average frame size for a 4K transfer is 819 bytes.
- Five credits would be required per km over a 4 Gbps fiber
- The "optimal" amount of buffer credits is determined by the distance (frame delivery time), the processing time at the receiving port, link signaling rate, and the size of the frames being transmitted:
- Credit = (round-trip-time + receiving-port-processing time) / frame-transmission-time
- Information required for calculation:
  - Link speed (1, 2, 4, 10 Gbps)
  - Actual fiber distance in km that the frame must traverse
  - The size of the frame – can be difficult to obtain
- Formula for assigning buffer credits using 2148 frame size
  - 1 Gbps where distance (in km) / 2 + 20%
  - 2 Gbps where distance (in km) + 20%
  - 4 Gbps where distance (in km) x 2 + 20%
  - 10 Gbps where distance (in km) x 6 + 20%
- If you always assume a frame size of 2148 bytes, in the vast majority of cases, it will produce too few buffer credits
- For extended links, RTT and the average frame payload size must also be considered

Mainframe Services for Distance

- IBM Peer-to-Peer Remote Copy (PPRC) also known as Metro Mirror, is a hardware-based disaster recovery and workload migration solution that maintains a synchronous copy of data in a relatively close-by remote location
- IBM Global Mirror (eXtended Remote Copy – XRC) provides data replication over extended distances between two sites for business continuity and disaster recovery. It replicates the data asynchronously and then forms a Consistency Group at a regular interval allowing a clean recovery of the application.

There are additional products uses for distance from vendors such as EMC, HP, HDS, and others. The above two are only examples which are commonly used with FICON extension products.

- GDPS (Geographically Dispersed Parallel Sysplex) is an IBM mainframe feature for continuous availability and disaster recovery across multiple, geographically dispersed sites
- GDPS can be used with FICON channels and FICON Directors between multiple sites to provide a disaster recovery solution
- Global Mirror can be combined with GDPS to provide for automated failover between sites. This combined solution provides the lowest possible RTO (Recovery Time Objective).

FICON/FCP Intermix

Why Mix FICON and FCP?

- Protocol Intermix Mode (PIM), which is also referred to as FICON/FCP intermix, is supported by IBM on mainframes to enable FICON and open systems Fibre Channel Protocol (FCP) traffic to co-exist on the same physical storage network
- FICON/FCP intermix can provide cost reductions such as the consolidation and management of both switching infrastructure and cabling plants as well as access to the latest advances in technology, such as the N_Port ID Virtualization (NPIV)
- Examples for intermix:
  - Small FICON and open systems environments with a common storage network
  - z/OS or System z servers accessing remote FICON storage via FICON cascading
  - Linux on a mainframe running z/OS using FICON to access local storage
  - Hardware-based remote DASD mirroring between two locations using FCP as a transport, for example, PPRC
  - A mainframe processor that accesses storage via FICON and storage devices via FCP to perform remote mirroring between devices (instead of ESCON)
  - Open systems servers accessing storage on the same storage network using FCP
  - Linux on the zSeries located at either of two sites using FCP to access storage

Note: Open systems use discovery; FICON uses IOCP

FICON and FCP

- A FICON channel in Fibre Channel Protocol mode (which is CHPID type FCP) can access FCP devices as follows:
  - From a FICON channel in FCP mode through a single Fibre Channel switch or multiple switches to a SCSI device
  - From a FICON channel in FCP mode through a single Fibre Channel switch or multiple switches to a Fibre Channel-to-SCSI bridge
- FCP support enables Linux on System z to access industry-standard SCSI devices
- FICON Express4, FICON Express2, and FICON Express channels in FCP mode provide full fabric attachment of SCSI devices to the operating system images, using the Fibre Channel Protocol, and provide point-to-point attachment of SCSI devices
- For disk applications, these FCP storage devices use Fixed Block (512-byte) sectors instead of Extended Count Key Data (ECKD) format used with FICON storage devices

If you are considering creating a disaster recovery site by synchronously replicating data between DASD at the two sites, you would need to use FCP protocols between the DASD but still use cascaded FICON Directors to manage the data flow. In this case, you need an intermix environment of FICON and FCP.

- The FCP channel full fabric support enables switches and Directors to be supported between the System z server and SCSI device, which means many “hops” through a storage area network

- FICON channels in FCP mode can use the Queued Direct Input/Output (QDIO) architecture for communication with the operating system

FCP channels do not use control devices but instead, data devices that represent QDIO queue pairs are defined. These queue pairs consist of a request queue and a response queue. Each queue pair forms a communication path between an operating system and the FCP channel. The enables an operating system to send FCP requests to the FCP channel via the request queue. The FCP channel then uses the response queue to pass a completion message and unsolicited status messages to the operating system.

HCD/IOCP is used to define the FCP channel type and QDIO data devices. There is no definition requirement for the Fibre Channel storage controllers and devices in IOCP, nor the Fibre Channel devices such as switches and Directors because of QDIO.

The FCP industry-standard architecture specifies that the Fibre Channel devices (end nodes) in a Fibre Channel network use using World Wide Names (WWNs), Fibre Channel Identifiers (IDs), and Logical Unit Numbers (LUNs) for addressing. These addresses are configured at the operating system level, and passed to the FCP channel together with the corresponding Fibre Channel I/O via a logical QDIO device (queue).

Virtual Fabrics and FICON

Supported Features

- Multiple Logical Switches (up to 8 on a DCX)
- Can utilize 384 ports on a DCX chassis
- Utilizes 48-port blades
- Two instances of CUP
- Not on base switch
- Only supported in Native and Interopmode 2

Cloud Environments

Inter-chassis Links (ICLs)

If thought is being given to using some sort of cloud infrastructure, there are two technolgies that could be of assistance in a FICON environment. The first is the use of ICLs on a Brocade DCX.

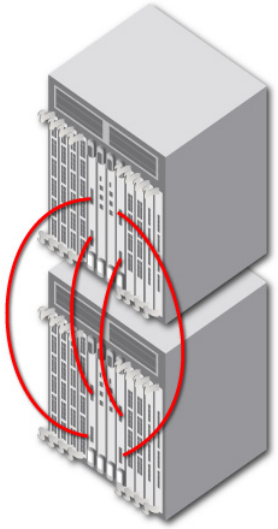

Figure 11: Brocade DCX with ICLs

- Each ICL is equivalent to 16 standard ISL ports (at 8 Gbps each) – but take up no user ports
- ICL cables can be a max of 2m in length
- ICLs are “silent” hops; IBM does not “count” them as an “official” hop for FICON cascading. By using ICLs, you can use four DCXs in a cascaded environment that does not “violate” the one hop rule.
- Cascading of Directors and switches is limited to one hop for a FICON environment with the following exceptions:
  - The DCX Backbone with ICLs consists of 2 or 3 domains, but the hop across the ICLs can be disregarded
  - Because of this DCX<->ICL<->DCX<->ISL<->DCX<->ICL<->DCX is considered one hop
  - FICON is supported for 16, 32 and 48 port blades
  - 48-port blades on a DCX require the use of Virtual Fabrics

N_Port ID Virtualization (NPIV)

NPIV is designed to allow the sharing of a single physical FCP channel among operating system images, whether in logical partitions or as z/VM guests in virtual machines. This is achieved by assigning a unique World Wide Port Name (WWPN) to each operating system connected to the FCP channel. In turn, each operating system appears to have its own distinct WWPN in the SAN environment, thus enabling separation of the associated FCP traffic on the channel.

Access controls based on the assigned WWPN can be applied in the SAN environment using standard mechanisms such as zoning in FC switches and Logical Unit Number (LUN) masking in the storage controllers.

NPIV is exclusive to the System z servers and is the preferred method for sharing SCSI devices among multiple operating system images when using a single physical FCP channel.

Benefits of Switched FICON

- Minimize the number of CHPID ports required for connection to storage.
- Switched FICON allows for simpler infrastructure management, lowered infrastructure cost of ownership, and higher data availability.
- Keep port, cable, or HBA failures to a single side of the system rather than allowing a complete connection to be disrupted.
- Can utilize Fan In – Fan Out ratios to continue to scale new or existing mainframe applications even if all the storage ports are in use.
- Distribute workload more evenly and predictably across all of the available storage ports through a well thought out Fan In – Fan Out ratio
- Ability to use NPIV with Linux
- Can provide very robust, dynamic channel path connectivity providing ease of growth and scalability throughout the FICON environment.
- Ability to implement Control Unit Port (CUP) to obtain information such as buffer credit usage and other performance statistics.
- Allow DASD arrays to contain maximum capacity that can actually be utilized due to the Fan In – Fan Out architecture and balanced use of the storage ports.
- Use FICON as a distance extension to support disaster recovery between two sites.

Minimize Mainframe Channel Card Costs with Switched FICON

- Each storage port in a point-to-point connection requires its very own physical port connection on the mainframe
- In a cascaded FICON environment, the fan-in/fan-out ratios solve this problem just like it solves other connectivity and scalability problems
  - Fan-in: Attaching servers to storage via connected server/switch ports to a single storage subsystem port
  - Fan-out: Attaching multiple storage control units to a single server/switch
- Every storage box has a maximum number of FICON connections:

  - Tape might be 1 or 2
  - Disk might be 2 to 48
- In a cascaded FICON environment, you can use fan-in/fan-out ratios
  - Use all storage ports
  - Add more capacity to the DASD
- Use CHPID fan-in to scale mainframe connectivity to the existing storage ports
- Director/switch ports are less expensive than mainframe channel cards
- NPIV can also expand the capabilities
- Can get buffer credit usage into RMF

Brocade Accelerator for FICON

FICON Read and Write Performance Over Distance

- Enables improved disaster recovery and data protection, supporting:
  - FICON Read and Write Tape Pipelining
  - FICON Emulation for IBM z/OS Global Mirror (formerly XRC)
- Faster backups, faster recoveries over distance
- Flexibility to place FICON disk and tape where needed, regardless of location
- Conserves WAN bandwidth, reducing costs
- Use port-based routing
- In-order delivery must be set
- Set TI Zones properly

Other Distance Considerations

- Channel extenders
- Asynchronous disk replication

Enhanced TI Zones

- Prior to FOS v6.4.0, a port could be in only one TI zone at a time
- Starting in FOS v6.4.0, ports can be in multiple TI zones at the same time
- Zones with overlapping port members are called enhanced TI zones (ETIZ)
  - Device ports can exist in multipleTI zones
  - Devices in a failover disabled TI zone can communicate with local devices that are not part of the same TI zone
- Enhanced TI Zoning is available only on the following platforms running Fabric OS v6.4.0:
  - Brocade 5100, 5300, 7800, DCX and DCX-4S
- For FICON, it allows a source device to reach two separate destinations over exclusive, one-hop paths, while preventing the source device from taking the two-hop path in the event of a failure
  - Can be used to isolate distance problems
- If failover is enabled:
  - Non-TI zone traffic will use the dedicated path if the non-dedicated paths are not the shortest paths
  - Non-TI zone traffic will use the dedicated path if no other E or VE paths through the fabric exist
- If failover is disabled:
  - Paths will be isolated to specific ports

Replacing Edge M3000s

If you are thinking about replacing an Edge M3000 in a cascaded tape environment, there are many factors to consider:

- Fabric managment techniques
- Domain IDs
- Whether or not to use TI Zones
- Interop modes
- Security: SCC_Policy
- Principal switch selection

Distance Settings in **`portcfglongdistance`**

Fibre Channel gigabit values reference definition Before you can calculate the buffer requirement, note the following Fibre Channel gigabit values reference definition:

- 1.0625 for 1 Gbps
- 2.125 for 2 Gbps
- 4.25 for 4 Gbps
- 8.5 for 8 Gbps

In cases where the frame size is average, for example 1024 bytes, you must allocate twice the buffer credits or give twice the distance in the long-distance LS configuration mode.

1. Use the following formula to calculate value for the desired_distance needed for Fabric OS to determine the number of BB credits to allocate:

*desired_distance = roundup [(real_estimated_distance * 2112) / average_payload_size]*

Where average_payload_size = 1024 bytes

This example uses 100 km for the real estimated distance.

*desired_distance = roundup [(100 * 2112) / 1024] = 207*

When configuring the LS mode with the `portcfglongdistance` command, enter a desired_distance value of 207 for an actual 100 km link connected to an 8 Gbps E_Port. This causes Fabric OS to allocate the correct number of BB credits.

2. Determine the speed that you will use for the long-distance connection. This example uses 8 Gbps.
3. Look up the data_rate value for the speed of the connection. Reference the Fibre Channel gigabit values reference definition to determine the data_rate value.
4. Use the following formula to calculate the number of buffer-to-buffer credits to allocate:

*BB credits = roundup [desired_distance * (data_rate / 2.125)]*

Using the values for desired_distance and data_rate from step 1 and step 3, the value for BB credits is calculated as follows:

*BB credits = roundup [(207 * 8.5) / 2.125] = 828*

3 - Implementation

Basic FICON SAN Installation Procedures

QPM Cards

				1 Gbit/sec		2 Gbit/sec		4 Gbit/sec	
Configuration	Port Type	Port Count	Limitation	F	E	F	E	F	E
1 Gbit/sec	G_Port	4	None	X	X				
2 Gbit/sec	G_Port	4	None			X	X		
Negotiate 2 Gbit/sec Max	G_Port	4	Negotiate 1Gbit/sec and 2Gbit/sec Only	X	X	X	X		
4 Gbit/sec Burst	F_Port	4	F_Port only					X	
4 Gbit/sec Sustained	G_Port	2	Only 2 ports (0 and 2) per card					X	X
4 Gbit/sec Burst Negotiate	F_Port	4	F_Port only	X		X		X	
4 Gbit/sec Sustained Negotiate	G_Port	2	Only 2 ports (0 and 2) per card	X	X	X	X	X	X

Figure 12: 4 Gbps QPM Configurations M6140

Verifying Device Registration

- The `ficonshow rnid` command is used to show which FICON devices are registered with the Director:

```
admin> ficonshow rnid
{
 {Fmt  Type PID    Registered Port WWN     Registered  Node WWN    flag Parm
  0x18 N    040d00 50:05:07:64:01:00:a9:08 50:05:07:64:00:ce:ca:bc 0x00 0x300000
  Type number:           003868
  Model number:          003
  Manufacturer:          IBM
  Plant of Manufacture: PK
  Sequence Number:       0ECABC05A0DC
  tag:                   0111
```

Domain Offsets

The M6140 has a unique offset of 0x04 that is used in calculating the port address. On B-Series products, in Fabric OS, the `configure` command is used to set any Domain ID attributes.

FICON Zoning

In general, the best practice is to put FICON devices in one zone and FCP devices in another zone. You may want to enbale default zoning as well for M-Series Directors. Activating a zone will generate RSCNs, which are ignored by FICON devices, but reconized by open systems devices.

The FICON Prohibit Dynamic Connectivity Mask (PDCM) controls whether or not communication between a pair of ports in the switch is prohibited or allowed. If there are any differences or conflicts in restrictions set up with Brocade Advanced Zoning and PDCM, the most restrictive rules are automatically applied. All FICON devices should be configured in a single zone using the "Domain, Index" notation. PDCM can then be used to "Allow" or "Prohibit" access between specific port pairs. PDCM persists across a failover because it is replicated at all times to the standby CP blade. The active PDCM configuration is saved to the IPL if the ASM bit is set on.

Preferred Path

Open trunking is recommended for most applications. Preferred pathing allows you to control which ISLs certain applications use based on the ports where the channels and devices are connected while allowing fail over to alternate paths should the preferred path fail. It should be noted the configuration should be both ways. Only one method should be used.

Use the Preferred Path method to steer application traffic down a particular path. In the case where a Preferred Path goes away, only the ports assigned to that path are re-routed to other paths. When the Preferred Path is re-established, the ports assigned to that path are re-routed back. Use the PDCM method for specific forced network control. PDCM is an M-Series equivalent feature that lets you define which ingress port takes which egress port. You must perform this configuration at the Director or enterprise class platform level.

Display Current Fill Word Setting

The easiest way to display the fill word setting is to use the CLI. To display the current settings, use the command `portcfgshow`. A fill word of 0 ("zero") represents Idles for fill words and 1 represents ARB(FF) for fill words. The portCfgshow command is limited to the Logical Switch for which the command is executed in. To display the fill word setting of ports in other Logical Switches, it is necessary to change the Logical Switch context. In the sample output below, ports 1/0, 1/1, and 1/2 are configured for a fill word of ARB(FF), while the remaining ports are configured for Idle fill words.

```
ibm_dcx_210:FID128:admin> portcfgshow
Ports of Slot 1  0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
-----------------+--+--+--+--+----+--+--+--+----+--+--+--+----+--+--+--
Speed AN AN AN AN AN AN AN AN AN AN AN AN AN AN AN AN
Fill Word 1  1  1  0  0  0  0  0  0  0  0  0  0  0  0  0
AL_PA Offset 13 .. .. .. .. .. .. .. .. .. .. .. .. .. .. .. ..
Trunk Port ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON ON
```

Cascaded FICON

Enterprise Fabric Mode

Enterprise fabric mode was intended to provide users with a common set of security features for convenience. Using it is not recommended since it turns on rerouting delay. Rerouting must be turned off. Furthermore, it turns on switch binding which may not be desirable. When using this mode, IDIDs and fabric binding are in place. Fabric Binding cannot be disabled while Enterprise Fabric Mode is active even if the switch is offline. SCC policies are supported. For fabric-wide consistency, a strict SCC policy should be used however. Use the minimum required level of binding. The minimum level of binding required for cascaded FICON is fabric binding. Try eliminate any hops of concern, a place where a frame can go an alternate way.

High-integrity Fabrics

To configure a high-integrity fabric (cascaded configuration):

1. Disable each switch in the fabric
2. For each switch:
   - Enable the IDID flag
   - Set the domain ID
3. Enable the switches; this builds the fabric
4. Configure the Switch Connection Control policies on all switches to limit connectivity to only the switches in the selected fabric
5. Save or activate the new policy
6. Enable ACL Fabric Wide Consistency Policy and enforce a strict SCC policy
7. Connect and enable channel and control unit (CU) devices. The Query for Security Attributes (QSA) response to the channel indicates that the fabric binding and IDID are enabled.

B-Series DCC Policies

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. Empty DCC policies will prevent any devices from joining the fabric.

Interop

- Operating Mode:
  - For M-Series only environment, set to McDATA Fabric Mode
  - For B-Series only environment, use interop mode 0
- If using both B-series and M-Series, use interop mode 2 on B-Series and McDATA Fabric Mode on the M-Series
- Backup the B-Series switch configuration with `configupload` after interop mode 2 is in place

McDATA Fabric Mode Restrictions

- A maximum of 2048 devices can be defined
- Extended Edge PID format cannot be use
- A maximum of 31 switches (Domain IDs)
- Domain IDs must be in the 1 to 31 value range on Fabric OS switches for successful connection to M-EOS switches. The firmware automatically assigns a valid domain ID, if necessary.
- The DCC policy or port based security is not supported in McDATA Fabric mode

Configuring a Cascaded M-Series Director with a B-Series

- Verify that the FICON Management Server and any necessary SANtegrity (for M-Series) licenses are installed
- Configure operating parameters to ensure that Periodic Date/Time Synchronization is not enabled
- Set the Product Management Style to FICON
- Configure port names to where channel systems are attached
- Change the device to McDATA Fabric Mode
- Use DCFM or command line to set McDATA Fabric Mode on the B-Series device

Top Talker

Top Talker monitors determine the flows (SID/DID pairs) 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. You can use Top Talkers to identify the SID/DID pairs that consume the most bandwidth and can then configure them with certain Quality of Service (QoS) attributes so they get proper priority. The Top Talker monitor is based on SID/DID and not WWNs. Once Top Talker is installed on a switch

or port, it remains installed across power cycles. Top Talkers supports two modes, port mode and fabric mode:

- Port mode Top Talker
  - A Top Talker monitor can be installed on an F_Port to measure the traffic originating from the F_Port and flowing to different destinations.
- Fabric mode Top Talker
  - In fabric mode, Top Talker monitors are installed on E_Ports in the fabric and measure the data rate of all the possible flows in the fabric (ingress E_Port traffic only). In fabric mode, Top Talker monitors can determine the top n bandwidth users on a given switch. You can install Top Talker monitors either in port mode or fabric mode, but not both. The Advanced Performance Monitoring license is required to implement Top Talker.

Port Swapping

Port swapping refers to the capability of a FICON Director to redirect traffic on a failed F_Port to a working F_Port without requiring a change in System z I/O configuration. In a Fibre Channel fabric, a failing port typically results in a cable being reconnected to another available port, followed by automatic discovery of the device through the name server. However, in a FICON environment, the control unit link address is defined in the channel configuration file (IOCP) of the System z server. Therefore, the FICON Director must ensure that the N_Port address for the control unit remains the same even after the cable is reconnected to a different switch port. After port swapping, the address assigned to the N_Port performing fabric login (FLOGI) by the alternate port should be the same as the one that would have been assigned by the original port. The original port will assume the port address associated with the alternate port. Ports must be contained in the same logical switch. GbE ports may not be swapped.

M-Series and B-Series Interoperability

All cascading requirements must be met, plus the following:

- Switch wide BB Credit setting for all B-Series switches participating in a fabric with M-EOS based switches must be 16. Individual port settings may be different.
- The principal switch must be an M-Series switch.
- Set the AL_PA on the B-Series switches to 13. This must be done for all ports that will carry FICON traffic. Note that for open systems connections, including NPIV channels from the mainframe, the AL_PA must remain at 0. It doesn't matter how E_Ports are configured; 13 is recommended for cable swapping, etc.
- The M-Series switches must be configured with a domain ID offset of 96 (0x60). Note that this mode limits the domain ID range to 1-31.
- B-Series switches must be configured for in interopmode=2 with a domain ID in the range of 1-31 to match the domain ID range of the M-Series switches. In interopmode=2, an offset is automatically applied, however, the resulting switch address may not always be displayed. The link address for port 0x77 on a switch configured for interopmode=2 and domain ID=6 is 6677 but in response to the `switchshow` command with FOS v6.1.0c is displayed as 067713 instead of 667713. Note that the 48-port blade is not supported
- The SCC policy must be set to tolerant on the FOS switches
- Must be running FOS v6.1.0c or higher. (Note: Web & Zoning license no longer needed with 6.1.0c.)
- If migrating to B-Series switches from M-Series, the AL_PA may be 00 and should be the same for all ports.

4 - Management

Performance Management Tools

CUP for use with RMF

- RMF is the *key* motivator for using CUP:
  - RMF is a mainframe-centric view of the channel and device activity
  - Without CUP, RMF has no idea how fabric connectivity is affecting the I/O activity that it sees on any given CHPID
  - With CUP, an RMF FICON switch shares the gathered statistics about each of its ports enabling RMF to accurately report I/O channel and device activity as well as network timings
  - Activity and summary reports can help troubleshoot issues
- CUP is turned on by implementing FICON Management Server (FMS)
- Only one LPAR in one of the systems of the sysplex should gather statistics from FICON switching devices
- Gathering FICON switch statistics is a processor heavy workload on the FICON switch so you will want to minimize both fiber cable utilization (it is an in-band gathering of statistics) and FICON device processor contention by using only one LPAR to gather RMF CUP statistics
- FICON switch statistics are gathered individually from each FICON switching device that has CUP enabled
  - Requires that SMF 74 subtype 7 records are being collected
- Missing interrupt handler processor timeout (MIHPTO) should be checked to be sure that FICON Directors or switches are set to 180 seconds. An increase may be needed when CUP ports are “boxed”. You may also need to increase the RMF collection interval as well. This is set through the Fabric OS CLI.
- In addition to FICON requirements, for B-Series and M-Series:
  - CUP license installed
  - Director port FE and FF reserved
  - `fmsmode` enabled
  - The switch is defined to the mainframe by CUP as a 256-port switch (due to CUP protocol limitation)
- CUP as an in-band management tool provides a great deal of useful information to the end-user. The data contains the following:
  - Port statistics - the number of words/frames transmitted and received
  - Switch node identifier - serial number, manufacturer, etc
  - Configuration file information - the list of configuration files residing on the Director
  - History summary (Director history buffer) - each change in status or configuration of the ports is logged in a history buffer
  - Switch configuration data - timeout values, ports, etc

RMF 74-7 Records

- Enabling RMF 74 subtype 7 (RMF 74-7) records yields an RMF report called the "FICON Director Activity Report"
- Data is collected for each RMF interval if FCD is specified in the `ERBRMFnn parmlib` member. The report captures information based on an interval which is set when you create the report.
  - The report captures a snapshot of data and the counters based on an time interval, such as 60 seconds. Often, you need to run these reports more than once and change the interval periods for troubleshooting to determine if there is a trend.
  - FMS must be enabled on the switch
  - IECIOSxx dataset `SYS1.PARMLIB` must specify `STATS=YES`
- This RMF report is often overlooked but contains very meaningful data concerning FICON I/O performance - in particular, frame pacing delay
- Frame pacing delay is the only available method to indicate a BB_Credit starvation issue on a given port

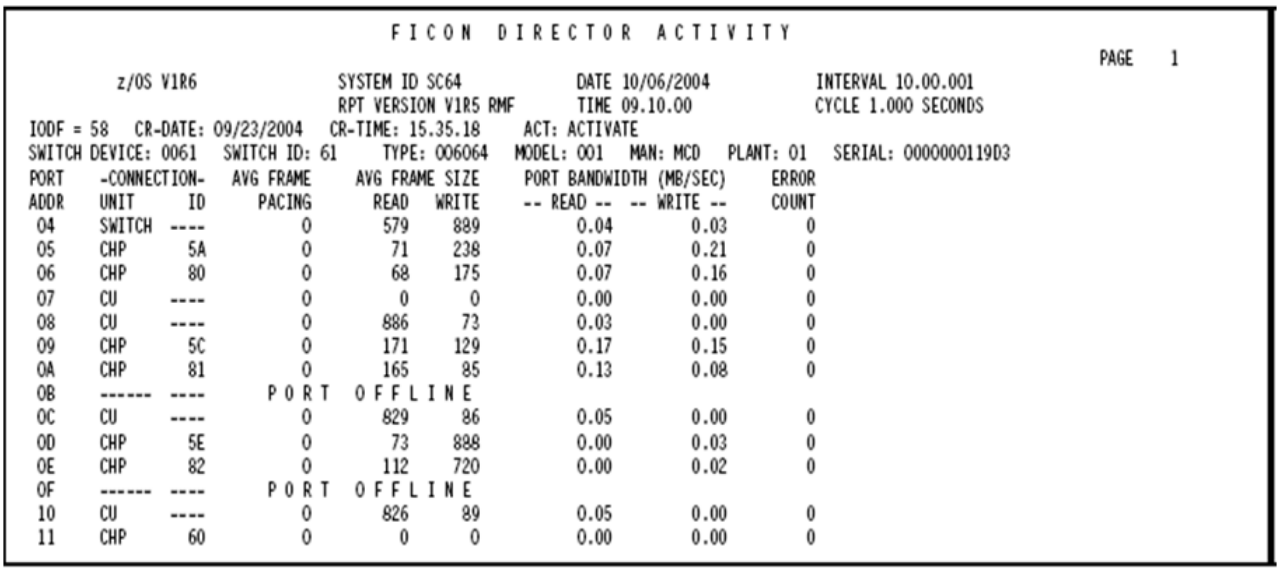

F I C O N D I R E C T O R A C T I V I T Y

PAGE 1

z/OS V1R6 SYSTEM ID SC64 DATE 10/06/2004 INTERVAL 10.00.001
RPT VERSION V1R5 RMF TIME 09.10.00 CYCLE 1.000 SECONDS

IODF = 58 CR-DATE: 09/23/2004 CR-TIME: 15.35.18 ACT: ACTIVATE
SWITCH DEVICE: 0061 SWITCH ID: 61 TYPE: 006064 MODEL: 001 MAN: MCD PLANT: 01 SERIAL: 000000119D3

PORT ADDR	-CONNECTION- UNIT	ID	AVG FRAME PACING	AVG FRAME SIZE READ	WRITE	PORT BANDWIDTH (MB/SEC) -- READ --	-- WRITE --	ERROR COUNT
04	SWITCH	----	0	579	889	0.04	0.03	0
05	CHP	5A	0	71	238	0.07	0.21	0
06	CHP	80	0	68	175	0.07	0.16	0
07	CU	----	0	0	0	0.00	0.00	0
08	CU	----	0	886	73	0.03	0.00	0
09	CHP	5C	0	171	129	0.17	0.15	0
0A	CHP	81	0	165	85	0.13	0.08	0
0B	------	----	P O R T O F F L I N E					
0C	CU	----	0	829	86	0.05	0.00	0
0D	CHP	5E	0	73	888	0.00	0.03	0
0E	CHP	82	0	112	720	0.00	0.02	0
0F	------	----	P O R T O F F L I N E					
10	CU	----	0	826	89	0.05	0.00	0
11	CHP	60	0	0	0	0.00	0.00	0

Figure 13: FICON Director Activity Report

DASD Activity Report

The DASD Activity Report is the first starting point for any I/O analysis

```
                                D I R E C T   A C C E S S   D E V I C E   A C T I V I T Y

            z/OS V1R8                SYSTEM ID ZOSA           DATE 07/22/2008          INTERVAL 10.45.340
                                     RPT VERSION V1R8 RMF     TIME 16.19.14            CYCLE 1.000 SECONDS
-
 TOTAL SAMPLES =    645   IODF = 00   NO CREATION INFORMATION AVAILABLE
-                                        DEVICE   AVG  AVG    AVG  AVG    AVG  AVG  AVG     %     %     %     AVG    %     %
 STORAGE  DEV  DEVICE   VOLUME PAV  LCU  ACTIVITY RESP IOSQ   CMR  DB     PEND DISC CONN    DEV   DEV   DEV   NUMBER ANY   MT
  GROUP   NUM  TYPE     SERIAL          RATE     TIME TIME   DLY  DLY    TIME TIME TIME    CONN  UTIL  RESV  ALLOC  ALLOC PEND
-        0430 33903    ZOSYSB     001F   0.028   17.5  0.0    0.0  0.0    0.4  0.0 17.1    0.05  0.05  0.0   3.0    100.0  0.0
         0431 33903    ZOSYSC     001F   1.581    1.2  0.0    0.0  0.0    0.3  0.1  0.8    0.13 16.88 16.9   9.0    100.0  0.0
         0432 33903    ZOST05     001F   0.000    0.0  0.0    0.0  0.0    0.0  0.0  0.0    0.00  0.00  0.0   5.0    100.0  0.0
         0433 33903    ZOST06     001F   0.102    5.8  0.0    0.0  0.0    0.2  1.0  4.6    0.05  0.06  0.0   1.0    100.0  0.0
         0434 33903    ZOST07     001F   0.031    1.5  0.0    0.0  0.0    0.2  1.1  0.2    0.00  0.00  0.0   0.0    100.0  0.0
```

Figure 14: DASD Activity Report

RMF Channel Activity Report

CHANNEL ID	PATH TYPE	G	SHR	UTILIZATION(%) PART	UTILIZATION(%) TOTAL	UTILIZATION(%) BUS	READ(MB/SEC) PART	READ(MB/SEC) TOTAL	WRITE(MB/SEC) PART	WRITE(MB/SEC) TOTAL	FICON OPERATIONS RATE	FICON OPERATIONS ACTIVE	FICON OPERATIONS DEFER	zHPF OPERATIONS RATE	zHPF OPERATIONS ACTIVE	zHPF OPERATIONS DEFER
10	FC_S	5	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.0	0.0	0.0	0.0	0.0	0.0
11	FC_S	5	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.0	0.0	0.0	0.0	0.0	0.0
12	FC_S	5	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.0	0.0	0.0	0.0	0.0	0.0
13	FC_S	5	Y	63.97	63.97	12.98	74.11	74.11	0.00	0.00	0.4	1.0	0.0	**18092**	**8.8**	**0.0**

Figure 15: RMF Channel Activity Report

Use this report to troubleshoot any performance issues.

RMF FICON Director Activity Report

```
                                        F I C O N   D I R E C T O R   A C T I V I T Y
                                                                                                                    PAGE    1
            z/OS V1R8              SYSTEM ID ANONYMOUS       START 04/12/2009-04.45.00  INTERVAL 000.15.00
                                   RPT VERSION V1R8 RMF      END   04/12/2009-05.00.00  CYCLE 1.000 SECONDS
   IODF = A2   CR-DATE: 03/27/2009   CR-TIME: 18.43.51     ACT: ACTIVATE
   SWITCH DEVICE: 032B   SWITCH ID: 2B     TYPE: 006140   MODEL: 001   MAN: MCD   PLANT: 01   SERIAL: 00000131656G
   PORT     CONNECTION     AVG FRAME     AVG FRAME SIZE     PORT BANDWIDTH (MB/SEC)     ERROR
   ADDR     UNIT      ID      PACING      READ   WRITE      READ      WRITE      COUNT
     05     CHP       05           2      1009    1420     13.53      20.69          0
     07     CHP-H     6B         151      1113     120      0.00       0.00          0
     09     CHP       15           8      1102    1559     12.26      20.41          0
     0C     CHP-H     64         842       778     937      0.40       0.49          0
     0D     CHP       6B           1      1568    1842      7.30      17.68          0
     0F     CHP-H     66           0      1905      70      0.33       0.00          0
     10     CHP       64           2       693    1382      0.03       0.11          0
     13     CHP-H     19         119       702    1121      0.36       0.72          0
     16     CHP       12          29      1443    1818      0.95       1.35          0
     17     CHP       0B          16      1176    1627      1.57       1.77          0
     1A     CHP       15           0      1463    1810      1.56       1.63          0
     1B     CHP       0D           9      1237    1621      2.00       2.11          0
     1E     SWITCH    05         134       706    1130      0.36       0.71          0
     1F     CHP       21          27      1444    1821      0.94       1.36          0
     20     CU      E900          12      1414    1004     20.98      13.68          0
            CU      E800
            CU      E700
```

Figure 16: FICON Director Activity Report for an M6140

- The two values contained in the report that are most often used for troubleshooting are:
  - AVG FRAME PACING refers to frame pacing delay. If it is a non-zero number then it reflects the number of times that I/O was delayed for 2.5 microseconds or longer due to buffer credits falling to zero.
  - ERROR COUNT which is the number of errors which were encountered during the interval.
- The AVG FRAME PACING column is where frame pacing delay is reported.
- Even if the number is zero there could be a buffer credit shortage problem:
  - Switching device is not fast enough to catch all errors (M6140, Mi10K, etc.)
  - Buffer credits went to zero but delays to send the next frame were less than 2.5 microseconds
- Sampling times can be an issue

Advanced Performance Monitoring

- Based on Brocade Frame Filtering technology and a unique performance counter engine
- Comprehensive tool for monitoring the performance of networked storage resources
- Advanced Performance Monitoring provides the following monitors
- End-to-end monitors measure the traffic between a host/target pair
- Filter-based monitors measure the traffic transmitted through a port with specific values in the first 64 bytes of the frame
- ISL monitors measure the traffic transmitted through an InterSwitch Link (ISL) to different destination domains
- Top Talker monitors measure the flows that are major consumers of bandwidth on a port
- The Top Talker (TT) feature is an enhancement to APM end-to-end monitors
  - First available in Fabric OS v6.0
  - When enabled, these monitors determine which SID-DID pairs are the major users of switch F_Port bandwidth
  - Can be enabled on specific switch F_Ports on any switch in the fabric

Buffer-to-Buffer Credits

Buffer Credit Concepts

- One of the goals of Fibre Channel Protocol (and hence FICON) is to provide reliable delivery of information from sender to receiver
  - Providing a data link with a low bit-error rate is a good start, but simply minimizing the quantity of bit-level transmission errors is not enough
  - The goal is to guarantee/ensure consistent and reliable frame delivery. Flow control is one of the primary mechanisms for providing this reliability
- There are two types of flow control mechanisms in FICON and Fibre Channel that are used:
  - End-to-End Flow Control
  - Buffer-to-Buffer Flow Control
- Define the maximum amount of data that can be sent prior to an acknowledgement
- Buffer credits are physical ASIC port or card memory resources and are finite in number as a function of cost
- Within a fabric, each port may have a different number of buffer credits
- The number of available buffer credits is communicated at fabric logon (FLOGI)
- One buffer credit allows a device to send a frame of any size up to a maximum size of 2148 bytes (max of 2112 bytes of usable data within that frame)
- Assuming that each credit is completely full (2148 bytes), you need:
  - 1 BC for every 1 KM of link length on 2 Gbps links
  - 2 BCs for every 1 KM of link length on 4 Gbps links
  - 4 BCs for every 1 KM of link length on 8 Gbps links
- Unfortunately, z/OS disk workloads rarely produce full credits. On DASD the average frame size, worldwide, is about 819 bytes
  - Hence, five BCs would be required per km when using a 4 Gbps FICON link

How Many Credits are Needed?

- Information required for the calculation of buffer credits when the goal is to keep the link 100% busy:
  - Link speed (1, 2, 4, 8, 10 Gbps)
  - Actual fiber distance in km that the frame must traverse
  - The size of the frame – this can be difficult to obtain

DCFM

Overview

- Comprehensive fabric management application
- Enables unified management of data center fabrics
- Manages and secures the flow of data across multiple fabrics
- Available in three editions that differ in supported features, hardware platforms, and supported scalability limits
  - DCFM Enterprise for Enterprise customers
  - DCFM Professional Plus for mid-sized organizations
  - DCFM Professional for small SANs

EGM

- Enhanced Group Manager (EGM) is a Fabric OS-enabled license to ensure that Brocade devices can be managed out of the box using DCFM.
- DCFM Professional Plus and DCFM Enterprise editions require the Enhanced Group Management (EGM) license on Fabric OS switches
- Not required for M-EOS switches
- Legacy Brocade switches do not require an EGM license check
- Included with all new 8 Gbps Fabric OS-based Directors, switches, and embedded products
- 8 Gbps switches shipped with Fabric OS v6.1.1+ have EGM license
- 8 Gbps switches shipped with Fabric OS v6.0.x and upgraded to Fabric OS v6.1.1+ will require additional EGM license key
-

Assigning a Device to the Call Home Center

Discovered devices (switches, routers, and directors) are not assigned to a corresponding call home center automatically. You must manually assign each device to a call home center before you use call home.

To assign a device or multiple devices to a call home center, complete the following steps.

1. Select Monitor > Event Notification > Call Home. The Call Home dialog box displays.
2. Select the devices you want to assign to a call home center in the Products List table.
3. Select the call home center to which you want to assign the devices in the Call Home Center table. You can only assign a device to one call home center at a time. If you do not select a call home center, the selection defaults to the first call home center table. If you have made a previous selection on an assigned device or filter and you do not select a call home center, the selection defaults to the previous selection's call home center.
4. Click the right arrow button. The selected devices display beneath the selected call home center. Devices assigned to a call home center do not display in the Products List table.
5. Click OK to close the Call Home dialog box.

Blade Swapping

Blade-based port swap is mainly used for FICON and is only applicable for port blades. However, the Management application does not block blade-based port swap for other application blades, including the 8 Gbps 24-port blade.

You can swap all of the ports from one blade to another blade. During this operation all ports in the selected blades are swapped. This operation disrupts the traffic on all ports for the selected blades. If GigE ports are present on the blade, only the non-GigE ports are swapped. To swap blades, you must meet the following requirements:

- The chassis must be running Fabric OS 6.3 or later.
- The chassis must have at least two blades of same type present.
- To perform the Swap Blades function you must have Read and Write access for the Product Administration privilege.

FICON Configuration Wizard

The FICON Configuration Wizard can be used to perform these tasks:

- Setting high-integrity fabric configurations
- Setting the routing policy
- Setting the insistent Domain ID

Web Tools Through DCFM

Web Tools may be invoked through DCFM. It can be used to backup or restore a configuration file, or download firmware to the switch

Web Tools - Performance Monitoring

- Web Tools predefines basic graph types, to simplify performance monitoring. A wide range of end-to-end fabric, LUN, device, and port metrics graphs are included.
- The advanced monitoring graphs give more detailed performance information to help you manage your fabric
- You can access the basic monitoring graphs on all switches
- Advanced Monitoring graphs are available only on switches that have a Brocade Advanced Performance Monitoring license activated
- Graphs:
  - Port Throughput
    - The performance of a port, in bytes per second, for frames received and transmitted
  - Switch Aggregate Throughput
    - The aggregate performance of all ports on a switch
  - Blade Aggregate Throughput
    - The aggregate performance of all ports on a port card. This graph is available only for the Brocade 48000 and Brocade DCX Backbone.
  - Switch Throughput Utilization
    - The port throughput, in Gbps, at the time the sample is taken
  - Port Error
    - CRC errors for a given port
  - Switch Percent Utilization
    - The percentage utilization for each port in a switch
  - Port Snapshot Error
    - The CRC error count between sampling periods for all the ports on a switch

Firmware and Configuration

- Backing up a switch configuration
  - B-Series
    - CLI `configupload` command
    - Web Tools
    - DCFM
  - M-Series
    - Element Manager
    - DCFM
- Restoring Switch Configuration: active=saved
  - Restores configurations in a FICON environment
  - If the switch is operating in a FICON CUP environment, and the PDCM ASM (active=saved) bit is set on, then the switch ignores the IPL file downloaded when you restore a configuration. Use this option to save CUP configuration on the Director or switches
  - If `fmsmode` is enabled in a configuration file, but is disabled on the switch, the `configdownload` command fails and displays an error message. This prevents undesirable conditions that could result from enabling `fmsmode` on a switch that does not require it.
  - `portcfgpersistentenable` for the port will ensure it remains enabled
- Firmware notes
  - GigE ports in an FCIP extended distance configuration could go offline for 30 seconds during a firmware upgrade
  - FTP is required for firmware downloads to a Brocade switch
    - DCFM has a built-in FTP servers
    - CLI requires an FTP service be running on the host containing the firmware

Advanced Performance Monitoring

Advanced Performance Monitoring is a licensed feature that provides a comprehensive tool for monitoring the performance of networked storage resources. Additional performance monitoring features, such as CRC error reports, are provided through Web Tools and DCFM. Advanced Performance Monitoring provides the following monitors:

- End-to-End monitors measure the traffic between a host/target pair (filter-based).
- Frame monitors measure the traffic transmitted through a port with specific values in the first 64 bytes of the frame.
- ISL monitors measure the traffic transmitted through an InterSwitch Link (ISL) to different destination domains.
- Top Talkers monitors measure the flows that are major consumers of bandwidth on a switch or port.

Merging FICON Fabrics

If you want to join two cascaded FICON fabrics, they must be merged. For long distance connections, it is recommended that you create the Extended Fabrics configuration first, have an active connection, and have the E_port and cable distances values ready before you merge the fabrics. A merge conflict is possible. You can resolve the following types of switch configuration conflicts:

- Domain ID
- TOV
- Buffer To Buffer Credit
- Disable Device Probe
- Route Priority Per Frame
- Sequence Level Switching
- Suppress Class F
- Long Distance Setting
- Data Field Size
- VC Priority

SAN Health

IOCP

- SAN Health is compatible with FICON switches and can be used to review IOCP data with the SAN environment.
- The IOCP file contains the path configuration statements for a given mainframe system.
  - SAN Health will match these path statements with the physical switch ports found in the SAN Health audit
- While SAN Health is run, all interaction with the switches is displayed and can be stopped at any time.
- The diagnostics capture completes in approximately two minutes.
- SAN Health will remain open and capture performance data for the duration that you have set the performance capture

DEVICE MAP FOR FR-FICON

Table Of Contents

FR_Switch6						
Dom	Port	Speed	Description	Model	Port World Wide Name	Additional Information
6	0	2 G	IBM zSeries	IBM2064-1C1	50:05:07:64:01:40:5e:17	SN55CCD-ec00
6	1	2 G	IBM zSeries	IBM2064-1C1	50:05:07:64:01:40:74:a9	SN55CCD-fd01
6	2	2 G	Hitachi Storage	HTC2105-F20	50:06:0e:80:04:76:80:c1	SN30336 Cluster13 PortB SN30336-00a4
6	3	4 G	Hitachi Storage	HTC2105-F20	50:06:0e:80:04:5a:b7:49	SN23223 Cluster5 PortK SN23223-0024
6	4	2 G	IBM zSeries	IBM2064-2C5	50:05:07:64:01:40:69:38	SN5505D-0304
6	5	4 G	IBM zSeries	IBM2096-S07	50:05:07:64:01:00:b1:8a	SNFD56D-8017
6	6	4 G	IBM zSeries	IBM2096-S07	50:05:07:64:01:40:b2:4e	SNFD56D-801a
6	7	2 G	Hitachi Storage	HTC2105-F20	50:06:0e:80:04:76:80:81	SN30336 Cluster9 PortB SN30336-0084
6	8	2 G	Hitachi Storage	HTC2105-F20	50:06:0e:80:04:76:80:08	SN30336 Cluster1 PortJ SN30336-0000

Figure 17: SAN Health Output with IOCP Data

To the standard SAN Health Report and Visio diagram, Brocade adds FICON information including the Serial/ Sequence number. This information appears anywhere where there is port specific information in the report. AN EXcel spreadsheet is also produced with CHPIDs and CNTLUNITs.

5 - Troubleshooting

Data Collection

Management Tools for Data Collection

- DCFM
- Fabric OS CLI
- SAN Health

Troubleshooting FICON CUP

- `fmsmode` must be enabled
- Port 254 may not be used on a B-Series Director

IFCCs

- The System z FICON channel's approach to capturing error information at the time of a FICON I/O operation failure is to detect an Interface Control Check (IFCC). The FICON purge-path-extended function requests, collects, and transfers link error statistical buffer information from each Fibre Channel port in the FICON channel-to-CU path to the System z host (MVS) console.
- The FICON purge-path-extended function is used by the mainframe to issue a query to determine the end-to-end path information
- Example of an IFCC Error Message: IOS051I
  - `2027-20:17:04 (MVS 20:16:54) IOS051I INTERFACE TIMEOUT DETECTED ON D829,D0,E7,**02,PCHID=0520`
  - A possible MIHPTO issue
- The message help provides this:
  - `IOS051I INTERFACE TIMEOUT DETECTED ON` dev, chp, cmd, stat
  - Explanation: The channel subsystem detected a timeout condition during the operation of device. This message is issued as an informational message only for the interface timeouts that occur on native FICON channels
  - In the message text:
    - dev `(D829)` is the device number.
    - chp `(D0)` are the channel path identifiers (CHPID), if known, otherwise, this field is set to asterisks
    - cmd `(E7)` is the failing command code, if known; otherwise, this field is set to asterisks
    - stat `(**02)` is the device and subchannel status, if known; otherwise, this field is set to asterisks. In the example, the `**02` indicates an IFCC which indicates a problem with the device.
    - `PCHID=0520` is the physical channel path identifier

Using the supportsave File

FICON-specific information captured:

- RNIDs
- The FMS mode
- LIRRs

Host-based and Path Troubleshooting

Purge Path Extended Function

The purge path extended function provides enhanced capability for FICON problem determination. The FICON purge path error-recovery function is extended so that it transfers error-related data and statistics between the channel and entry switch and the control unit and its entry switch to the host operating system.

FSPF Monitoring

- FSPF messages are carried by Class F traffic
- They are not sent across multiple hops
- They all begin with a common header
- Communications is between fabric controllers (0xFFFFFD)
- A topology database is replicated across all switches in the fabric

GTF Trace

GTF trace is a useful tool for analyzing problems with I/O operations to devices attached to FICON channels. More than in an ESCON or parallel environment, it is necessary to be aware that the CCW chain traced on a SSCH entry may not be the CCW chain that is sent on the interface. There are several events, such as the execution of a PCI, that need to be considered.

Open Exchanges

- An Open Exchange (OE) is a logical resource that represents an I/O being processed by a channel. A FICON channel is limited to 64 concurrent OEs.
- An Information Unit (IU) is a logical resource, comprised of 1 to 4 credits, which may contain a CCW, a CCW and associated data, or just data. A FICON channel can support a maximum of 16 concurrent IUs for each OE.
- The number of OEs is very highly correlated with service time
- Service time dramatically increases as the OEs exceed 8

Using DCFM in Troubleshooting

Performance Issues

- Use the DCFM Real-Time graph for analysis

Link Issues

- Use the Element Manager and call Web Tools for a B-Series switch
  - Check LIRRs

Taking the Test

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

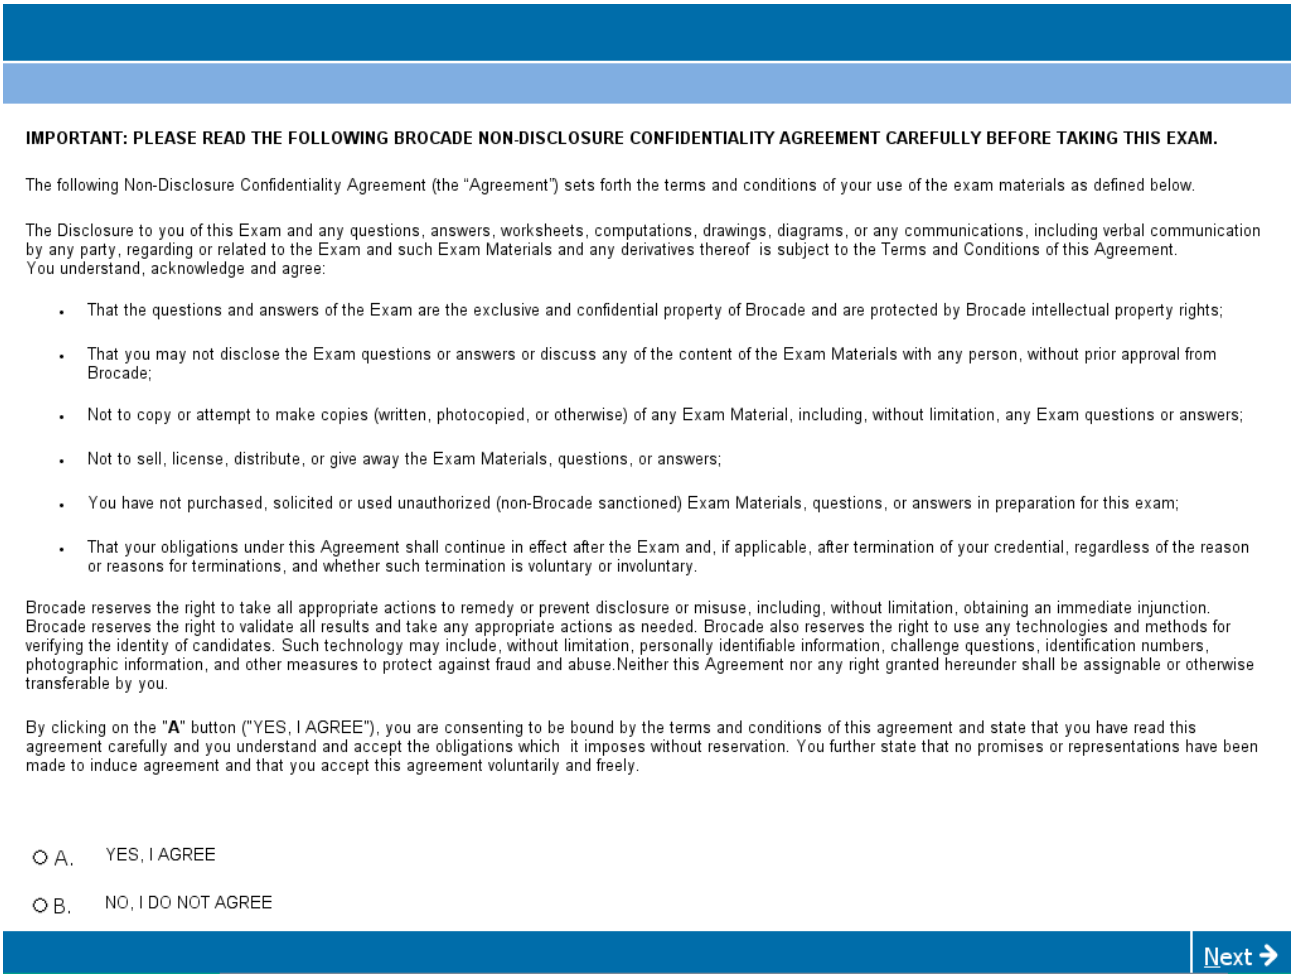

IMPORTANT: PLEASE READ THE FOLLOWING BROCADE NON-DISCLOSURE CONFIDENTIALITY AGREEMENT CAREFULLY BEFORE TAKING THIS EXAM.

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

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

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

Next →

Figure 18: Sample NDA