

Why Customers Should Deploy Switches In Their SAN and FICON Environments

- David Lytle, BCAF
- Brocade Communications Inc.
- Monday February 4, 2013 – 1:30pm to 2:30pm

- Session Number - 14275



QR Code



David Lytle, BCAF
Principal Engineer
Global Solutions Specialist
Brocade Communications, Inc.
dlytle@brocade.com

SWITCHED FIBRE CHANNEL



Switched Fibre Channel Fundamentals

Legal Disclaimer



- All or some of the products detailed in this presentation may still be under development and certain specifications, including but not limited to, release dates, prices, and product features, may change. The products may not function as intended and a production version of the products may never be released. Even if a production version is released, it may be materially different from the pre-release version discussed in this presentation.
- NOTHING IN THIS PRESENTATION SHALL BE DEEMED TO CREATE A WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, STATUTORY OR OTHERWISE, INCLUDING BUT NOT LIMITED TO, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT OF THIRD-PARTY RIGHTS WITH RESPECT TO ANY PRODUCTS AND SERVICES REFERENCED HEREIN.
- Brocade, Fabric OS, File Lifecycle Manager, MyView, and StorageX are registered trademarks and the Brocade B-wing symbol, DCX, and SAN Health are trademarks of Brocade Communications Systems, Inc. or its subsidiaries, in the United States and/or in 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.
- There are slides in this presentation that use IBM graphics.





Notes as part of the online handouts

I have saved the PDF files for my presentations in such a way that all of the audience notes are available as you read the PDF file that you download.

If there is a little balloon icon in the upper left hand corner of the slide then take your cursor and put it over the balloon and you will see the notes that I have made concerning the slide that you are viewing.

This will usually give you more information than just what the slide contains.

I hope this helps in your educational efforts!

DIRECT-ATTACHED STORAGE OVERVIEW

DIRECT ATTACHED STORAGE

Simplest method for connecting storage to a computer.

Using this method, a hard drive, or group of hard drives, is directly connected to a laptop, workstation, or server using a cable.


The storage can be internal or external to the server.

Direct Attached Storage (DAS)

Simplest storage connectivity

Benefits

- Simplest implementation
- Dedicated storage
- Easier to troubleshoot




Challenges

- Cannot share storage resources
- Results in lower storage utilization
- Results in lower bandwidth utilization

Direct Attached Storage (DAS)

Attributes of DAS Configurations

- Host/storage access
 - Host I/O and Storage I/O severely limits scalability of connectivity
- Channel (I/O) utilization
 - One-to-One relationships do not typically maximize channel efficiency
 - Mainframe NPIV cannot be deployed for Linux storage connectivity
- Point to Point I/O Paths
 - Any failure of the path reduces access and will cause application performance to suffer
- Mainframe CHIPID and Storage Interface Optics
 - Longwave SFP's used for host, switch and storage interfaces which equal some additional cost \$\$



SWITCH-ATTACHED STORAGE OVERVIEW

SWITCH ATTACHED STORAGE

Consists of one or more arrays of hard drives that are shared between multiple servers.

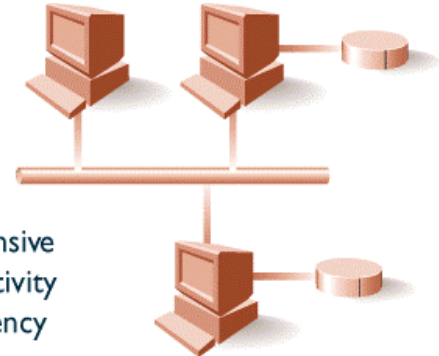

SANs removes the storage from behind the servers and centralizes it over a high-speed, scalable storage network.

Fibre Channel — Hybrid Connectivity




Channels

- Hardware intensive
- High speed
- Low latency
- Short distance
- Master/slave
- Robust data integrity
- Low-level error detection



Networks

- Software intensive
- Large connectivity
- Moderate latency
- Long distance
- Peer-to-peer
- Fragile
- ULP error detection



Fibre Channel

- Hardware intensive
- High speed
- Large connectivity
- Low latency
- Long distance
- Master/slave and peer-to-peer
- Robust data integrity
- No station management

Combines desirable features of both worlds, each with its own paradigm

Switch-Attached Storage

Storage Area Networking (SAN)





Benefits

- Isolation of storage traffic
- Flexibility to present storage everywhere at any time -- scalability
- Low latency/high performance
- Reliable delivery
- Disaster tolerance

Challenges

- Dedicated storage network takes deep planning and good architectures
- More complex to manage




Complete your sessions evaluation online at SHARE.org/BostonEval

Switch-Attached Storage (SAN)

Attributes of SAN Configurations

- Virtually Unlimited Storage access
 - Scales through the fabric with few limitations
- Maximizing Channel Utilization
 - Storage access can be consolidated to drive more efficient use of the host I/O
- Improves Failure Recovery
 - Multiple storage definitions can reside on same I/O for failover and recovery purposes
- Enables I/O Virtualization for Linux on System z
 - NPIV allows multiple address access to storage over a single interface
- Improves remote DR integration
 - Distance applications including Disk-to-Disk replication can be easily adapted through the fabric



Switched-Storage is a Best Practice



- Brocade switching devices do not cause performance problems within a local data center
- Architected and deployed correctly, Brocade switching devices do not cause performance problems even across very long distances
- In fact, use of Brocade switched-FC and Brocade FCIP long distance connectivity solutions can even enhance DASD replication performance and long distance tape operations effectiveness and performance
- Mainframe Switched-FICON is the only way to efficiently and effectively support Linux on System z connectivity (NPIV)
- Switched-FC is the only way to really take advantage of the full value of the open systems or Mainframe I/O subsystems

Complete your sessions evaluation online at SHARE.org/BostonEval

How Are Directors and Switches Different?




B-Series can run at up to 1600MBps on a port-by-port basis

FC Switches

- Good Availability up to **99.99%**
- Based upon *motherboard* design
- Some redundant components like power supplies and fans
- **24-80** Fiber Channel ports
- Decent fabric Scalability (*100's of ports*)
- Motherboard problems will require the switch to be replaced!

It is not when it is working, but rather when a problem occurs, that truly differentiates a Director from a Switch!



FC Directors

- Superb Availability up to **99.999%**
- Based on **discrete, redundant** parts
- Complete Redundancy and hot swap FRUs throughout the architecture
- Highest port counts – up to **384** ports
- Superior fabric Scalability (*1,000s of ports*)
- + Online Error Recovery (*non-disruptive failover*)
- + *Online Repair of the error (hot swap)*

90% or More of Customers Should Deploy Director-based, Switched-FC Fabrics!

How Are Directors and Switches Different

Physical Differences



- Since switches are **motherboard-based**, they are engineered to run at the then current line rate – cannot be upgraded -- and
 - Each port of an 8Gbps switch can be run using 4Gbps or 8Gbps SFPs
 - Each port of a 16Gbps switch can be run using 8Gbps, 10Gbps or 16Gbps SFPs
 - Failing SFPs can be hot-swapped but physical ports cannot be replaced
 - A switch must be completely replaced to repair a failed physical port(s) or ASIC

- Directors have **discrete, redundant components** that are engineered to run at current line rate – but can be upgraded – and non-disruptive firmware loads
 - Today each port of a 16Gbps Director can run using 8Gb, 10Gb or 16Gbps SFPs
 - Failing SFPs can be hot-swap replaced (along with fans and power supplies...)
 - New blades can replace blades that have failing or failed physical ports


- There are **16Gbps HBAs now and in the future, IBM will have 16Gbps mainframe CHPIDs**
 - Brocade's older 8Gbps Directors are upgradable to 16Gbps now
 - Once again providing our customers with investment protection!
 - But existing 8G switches will have to be completely swapped out and replaced with newer 16G capable switches in order to achieve 16G fabrics

How Are Directors and Switches Different

Availability



- **Complete, non-disruptive, firmware Hot Code Load is supported on Director class switches**
 - Since 2000
- **Comprehensive, non-disruptive, firmware Hot Code Load is not currently supported for FCIP blades and extension switches:**
 - On extension switches and blades, the FCIP tunnels will go down for 10-15 seconds and all traffic in the tunnels will be disrupted.
- **Brocade switches do try to support non-disruptive firmware upgrades but when upgrading firmware on a fixed port, motherboard-based switch, customers may experience recoverable errors (IFCCs):**
 - Switches have only a single CP on the ASIC and can drop frames if very busy
 - The errors for dropped frames are part of normal fibre channel recovery so this recovery happens with FCP and FICON channels and devices as well.
 - The only difference is that mainframes report absolutely everything. For most of our mainframe customers, they just need a warning that they may see a few IFCCs during firmware upgrades
 - IFCCs, regardless of the impact to traffic flow, then their best option is to always deploy a Director-class platform.



Switched-FC for Mainframe FICON and FCP



This Section Is For The Mainframers In The Audience

There are special reasons why a user should consider deploying Switched-FICON



- A smaller or older System z can start at about US\$50,000 while an IBM System zEC12 can be a little more 😊 but zBC12 helps makes even the newest technology more affordable by starting at around US\$75,000
- zEC12 provides a max of 320 FICON Express8S CHPIDs
- zBC12 provides a max of 128 FICON Express8S CHPIDs
- z196 provides a max of 320 FICON Express8S CHPIDs
- z114 provides a max of 128 FICON Express8S CHPIDs
 - CHPID ports to storage ports/mainframe are limited
 - Large Sequential throughput per CHPID/System z is limited
 - zEC12: $320 \times 620 \text{ MBps} = 198,400 \text{ MBps}$ – or 39% of full duplex 8G

Complete your sessions evaluation online at SHARE.org/BostonEval

System z: Why Deploy Switched-FICON




- With direct-attached FICON you must consume one CHPID to access one storage port
 - Very wasteful since neither CHPIDs nor DASD storage ports can make use of the full bandwidth of any of the channel paths
 - So CHPIDs and Storage Ports are always under-utilized resources that you have paid full price to deploy
- The CHPIDs on most mainframe channel cards cannot really perform at their listed line rate -- FICON Express8S running zHPF is the exception
 - In real use cases the data rate is less than ½ the performance these channel CHPIDs are rated to achieve

....**BUT**....

Customer's can use switched-FICON, making use of Fan In – Fan Out, to mitigate these limitations!



Mainframe Channel Cards

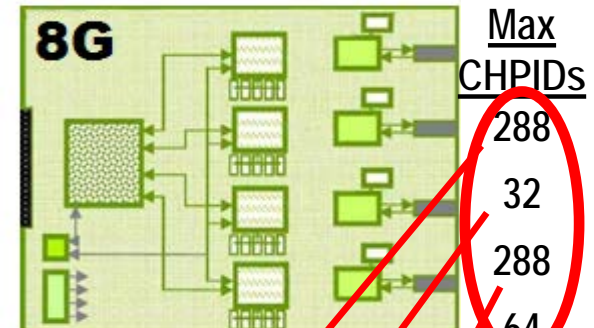


FICON Express8 – 4 ports
800MBps+800MBps=1600MBps

FICON Express8

- z10
- 2, 4 or 8 GBps link rate
- **Cannot Perform at 8Gbps!**
- Standard FICON Mode:
 32% ≤ 620 MBps Full Duplex
 out of 1600 MBps
- zHPF FICON Mode:
 46% ≤ 770 MBps Full Duplex
 out of 1600 MBps
- 40 Buffer Credits per port
 - Out to 5km
 assuming 1K frames

FICON switching devices will provide BCs for long distances



FICON Express8 – 4 ports
800MBps+800MBps=1600MBps

FICON Express8


- zEC12, zBC12, z196, z114
- 2, 4 or 8 GBps link rate
- **Cannot Perform at 8Gbps!**
- Standard FICON Mode:
 32% ≤ 620 MBps Full Duplex
 out of 1600 MBps
- zHPF FICON Mode:
 46% ≤ 770 MBps Full Duplex
 out of 1600 MBps
- 40 Buffer Credits per port
 - Out to 5km
 assuming 1K frames

Faster Processors but fewer total CHPIDs available

One or more IBM graphics are used above

Mainframe Channel Cards

Standard, long PCIe card

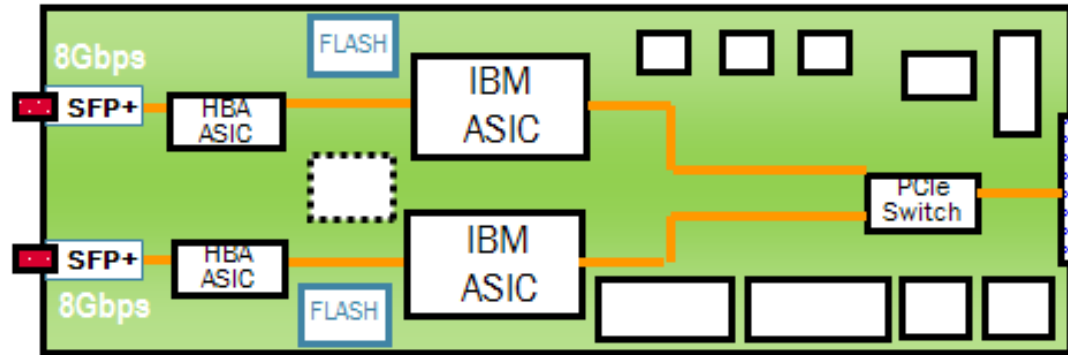


FICON Express8S – 2 ports
800MBps+800MBps=1600MBps

FICON Express8S

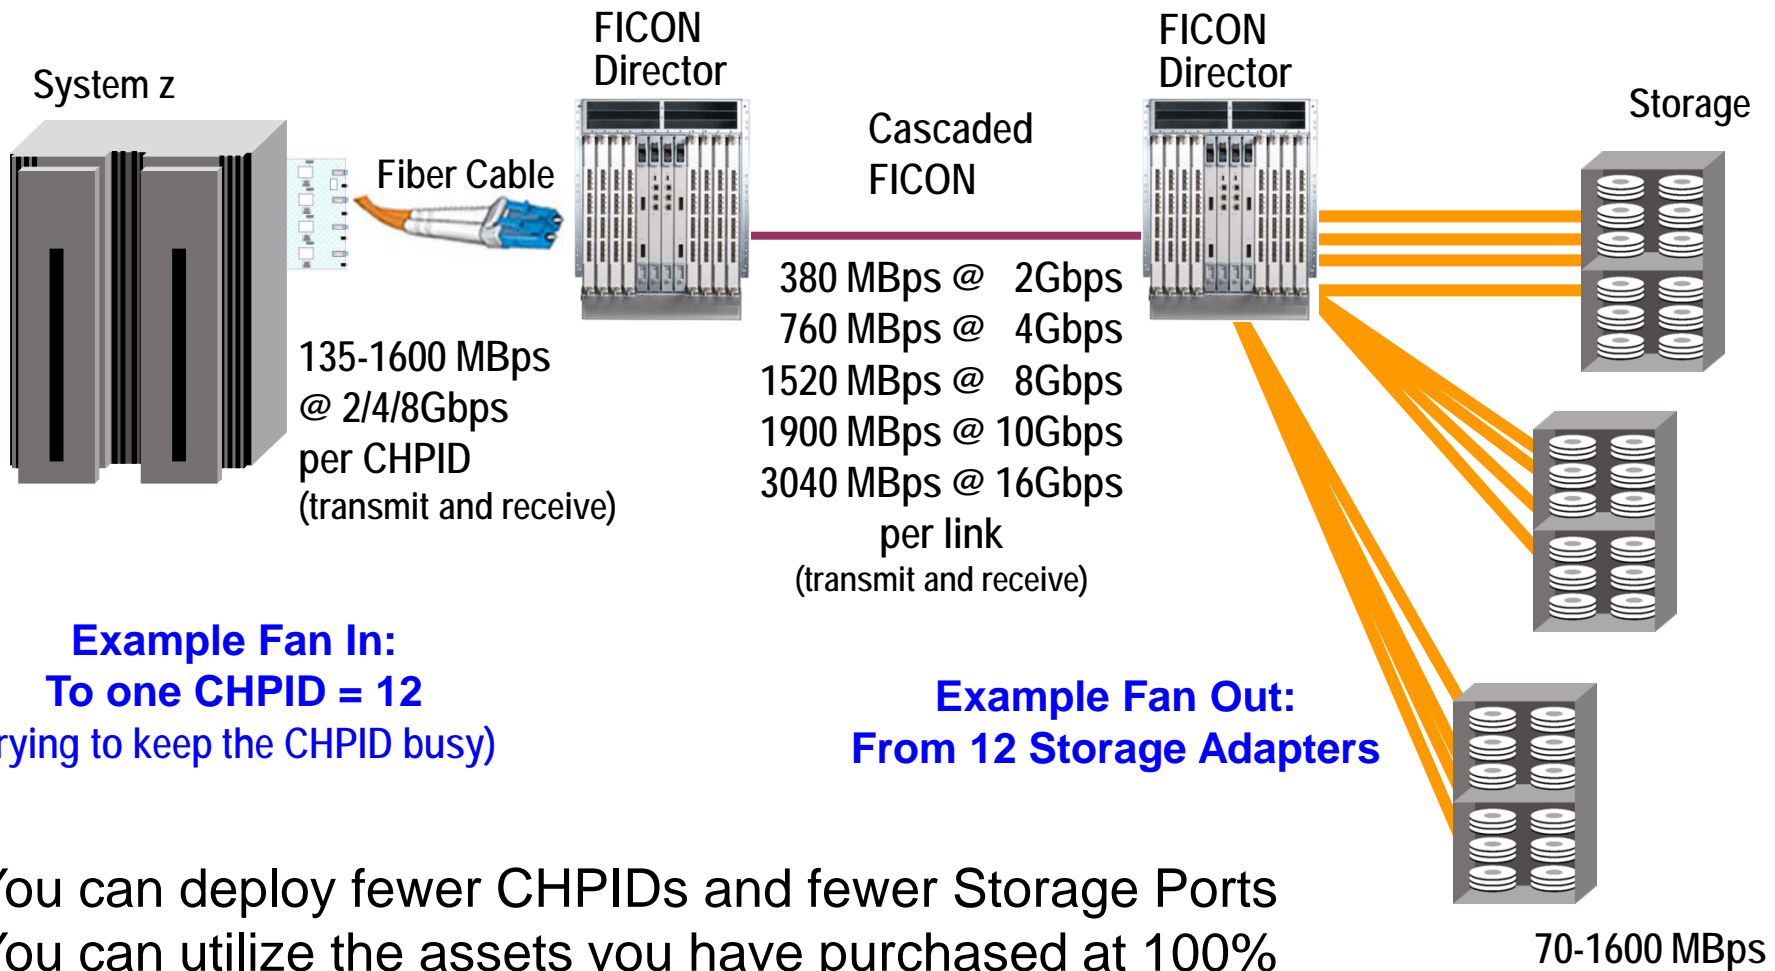
- zEC12, zBC12, z196, z114
- 2, 4 or 8 GBps link rate
- zHPF Performs at 8Gbps!
- Standard FICON Mode:
 - ≤ 620MBps Full Duplex out of 1600 MBps
- zHPF FICON Mode:
 - ≤ 1600 MBps Full Duplex out of 1600 MBps
- 40 Buffer Credits per port
 - Out to 5km assuming 1K frames

32%
100%



- For FICON, zHPF, and FCP environments
 - CHPID types: FC and FCP
 - 2 PCHIDs/CHPIDs
- Auto-negotiates to 2, 4, or 8Gbps
- Increased performance versus FICON Express8
- 10KM LX - 9 micron SM fiber
 - Unrepeated distance - 10 kilometers which 6.2 miles
 - Receiving device must also be LX
- SX - 50 or 62.5 micron multimode fiber
 - Distance variable with link data rate and fiber type
 - Receiving device must also be SX
- 2 channels of LX or SX (no mix)
- Small form factor pluggable (SFP) optics
 - Concurrent repair/replace action for each SFP

Fan In-Fan Out Reduces System Bottlenecks



Example Fan In:
To one CHPID = 12
(trying to keep the CHPID busy)

Example Fan Out:
From 12 Storage Adapters

- You can deploy fewer CHPIDs and fewer Storage Ports
- You can utilize the assets you have purchased at 100%
- You can scale up very easily without purchasing a lot of hardware
- You actually achieve a higher level of system availability

New z/OS and System z Functionality



System z functionality that **REQUIRES** customers to deploy switched-FICON:


- **FICON Express8 CHPID buffer credits:** Only 40 BCs per FICON Express8 and FICON Express8S CHPID limits long distance direct connectivity to ~5km. So customers can use up to 1,300 BCs at 8G and about 7,000 BCs at 16G, on a port on FICON switching devices, for longer distances.
- **FICON Dynamic Channel Management:** Ability to dynamically add and remove channel resources at Workload Manager discretion can be accomplished only in switched-FICON environments.
- **zDAC:** Simplified configuration of FICON connected disk and tape through z/OS FICON Discovery and Auto Configuration (zDAC) capability of switched-FICON fabrics. Required for z/OS v1.x users.
- **NPIV:** Excellent for Linux on the Mainframe, Node_Port ID Virtualization allows many FCP I/O users to interleave their I/O across a single physical channel path

Complete your sessions evaluation online at SHARE.org/BostonEval

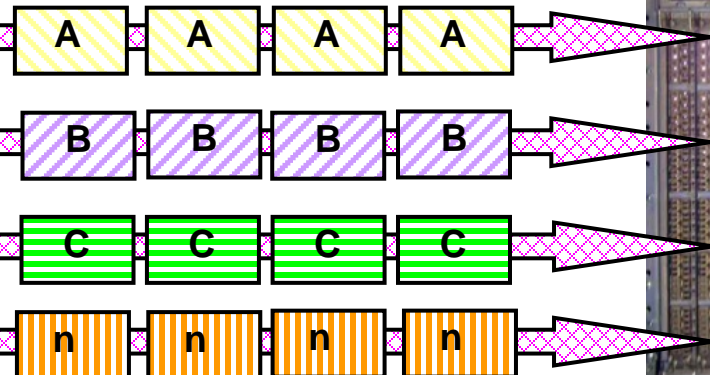


A Simplified Schematic - Linux without NPIV


An Example of
Linux on System z without NPIV



One FCP CHPID
per Linux guest



FICON
Director



For 300-3,000
guests,
no parallelism so it
is very difficult to
drive I/O for lots of
Linux images with
only 256 CHPIDs

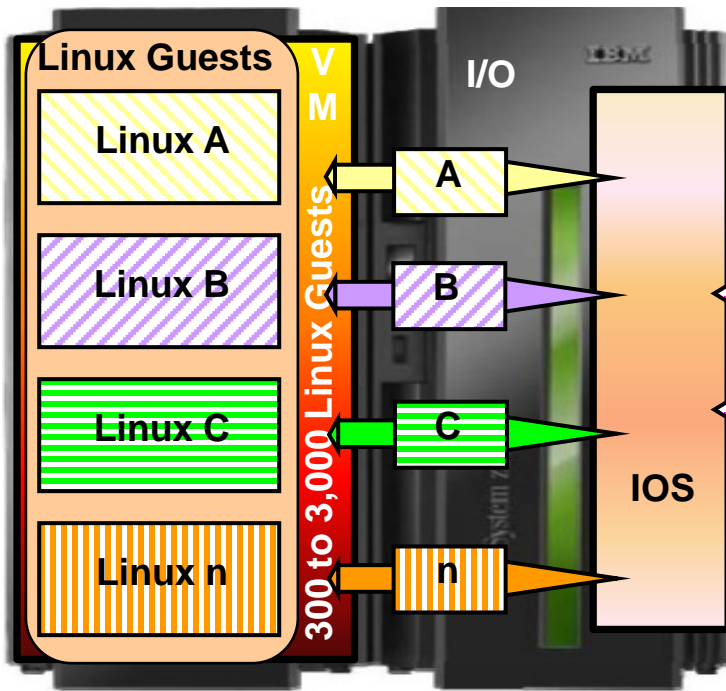
Probably very little
I/O bandwidth
utilization per
CHPID and
switch port

Linux on System z can run in its
own LPAR(s) but usually it is
deployed as guests under VM

A Simplified Schematic - Linux with NPIV


NPIV is ONLY available in a switched-FICON fabric!

An example of System z when using NPIV



One FCP channel for many Linux guests

FICON Director NPIV enabled



Lots of Parallelism

Fewer switch ports required!

Much better I/O bandwidth utilization per path

8Gbps Is Great For NPIV!

Using FICON Dynamic Channel Management



FICON Dynamic Channel Path Management (DCM) provides the ability for the z/OS system to manage FICON channel path assignment dynamically based on current workload conditions and availability characteristics.

- z/OS allows pools of FICON CHPIDs to be unassigned so that workload manager can use them when it is operating in GOAL mode.
- But mainframe channels and control units must be Switch Attached in order to make use of Dynamic Channel management (DCM)

Using zDAC with zEnterprise



Simplified configuration for FICON connected DASD and tape through z/OS Discovery and Auto Configuration (zDAC)


- zDAC is only useful for switched-FICON storage/host connections
 - zDAC must make use of the FICON fabric name server
 - zDAC is/will be allowed for non-switched FICON beginning with z/OS v2
- Uses intelligent analysis to help validate that server and storage definitions are compatible with each other
- Transparent to existing configurations and settings
- Invoked through and integrated with z/OS Hardware Configuration Definition (HCD) and z/OS Hardware Configuration Manager (HCM)
- Use on single systems or across your Sysplex's

Complete your sessions evaluation online at SHARE.org/BostonEval



Using zDAC with zEnterprise

z196,/z114/zEC12/zBC12



- The Fabric Name Server makes it possible to automatically discover what is new or has been changed in the fabric
- When a change is discovered, zDAC proposes a channel configuration based on:
 - High availability best practices
 - Customer configuration policies
 - Existing configurations

At z/OS V1.12+ (must have at least 1 LPAR for Dynamic I/O capability)

Other issues are described in the notes that are a part of this slide.

- zDAC attempts to make a symmetric configuration:
 - And this is repeated for each server in a sysplex





More Reasons For Deploying Switched-FICON

SHARE
Technology - Connections - Results

- Switched-FICON for RMF reports and z/OS Systems Automation control

Complete your sessions evaluation online at SHARE.org/BostonEval

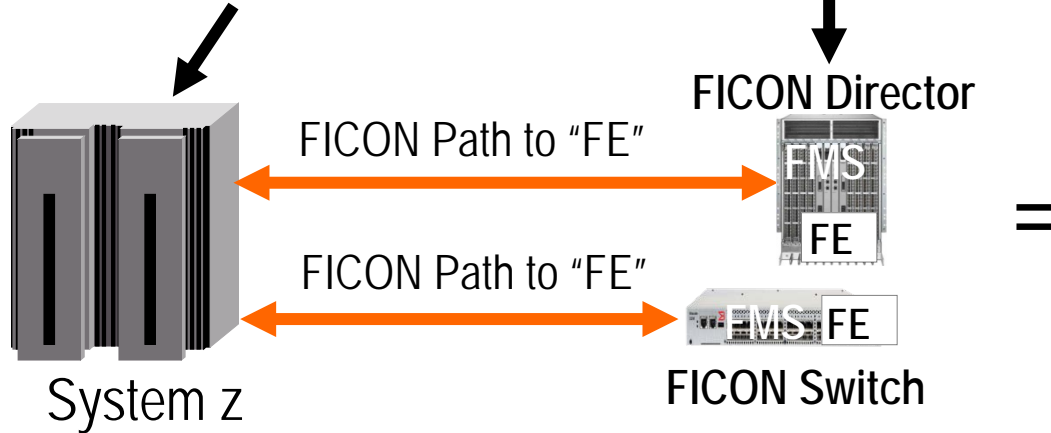


CUP on a FICON Switch For Use By RMF

Sys1.Parmlib options allow RMF to produce the RMF FICON Director Activity Report

Control Unit Port (CUP) on a FICON switching device allows z/OS to access switch and fabric information and provide that into RMF

FICON Director Activity Report per FICON Domain ID per Interval




FICON DIRECTOR ACTIVITY												PAGE 1
z/OS VARS		SYSTEM ID PRD1		START 04/12/2009-04:30:00		INTERVAL 000.15.00		END 04/12/2009-04:45:00		CYCLE 1.000 SECONDS		
IOCF = A2		CR-DATE: 03/27/2009		CR-TIME: 16.43.51		ACT: ACTIVATE						
SWITCH	DEVICE: 032B	SWITCH ID: 2B	TYPE: 006140	MODE: 001	MAN: MCD	PIANT: 01	SERIAL: 00000131					
PORT	-CONNECTION-	AVG FRAME	AVG FRAME SIZE	PORT BANDWIDTH (MB/SEC)	ERROR							
ADDR	UNIT	ID	PACING	READ	WRITE	-- READ --	-- WRITE --	COUNT				
05	CHP	05	0	849	1436	8.63	17.34	0				
07	CHP-H	6B	0	1661	1385	0.87	0.52	0				
09	CHP	15	7	833	1429	11.96	20.49	0				
0C	CHP-H	64	0	939	1099	0.39	0.50	0				
0D	CHP	6B	1	1328	1823	3.56	12.73	0				
0F	CHP-H	66	0	1456	1075	1.05	2.61	0				
10	CHP	64	0	644	1380	0.03	0.13	0				
13	CHP-H	19	0	907	885	0.58	0.45	0				
16	CHP	12	0	1241	1738	0.27	1.72	0				
17	CHP	0B	0	655	1650	0.10	0.82	0				
1A	CHP	15	0	1144	1664	0.65	1.18	0				
1B	CHP	0D	0	510	1759	0.12	1.72	0				
1E	CHP-H	05	0	918	894	0.59	0.45	0				
1F	CHP	21	0	1243	1736	0.27	1.70	0				
20	CU	E900	0	1429	849	17.66	8.85	0				
	CU	E800										
	CU	E700										
22	CHP	10	0	923	1753	0.55	2.78	0				
23	CHP	54	0	1805	69	0.80	0.00	0				
24	CHP	64	0	89	1345	0.00	0.00	0				
27	CHP	6B	0	1619	82	0.01	0.00	0				
28	CHP	95	27	918	1889	10.32	30.56	0				
28	CHP	70	0	69	2022	0.00	0.71	0				

- FICON Management Server (FMS) is a license to enable Control Unit Port (CUP) on a FICON switching device – always uses the “embedded” port x”FE”
- FICON Director Activity Reports are very useful to customers who would like to understand their average frame sizes traversing their fabrics as well as information about how buffer credits are being utilized

Using Buffer Credits is how FC does Flow Control, also called "Frame Pacing"

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

IODF = A2		CR-DATE: 03/27/2009	CR-TIME: 18.43.51	ACT: ACTIVATE	START 04/12/2009-04.30.00		INTERVAL 000.15.00
SWITCH DEVICE: 032B		SWITCH ID: 2B	TYPE: 006140	MODEL: 001	MAN: MCD	PLANT: 01	SERIAL: 0000HIJKLMN
PORT ADDR	-CONNECTION- UNIT ID	AVG FRAME PACING	AVG FRAME SIZE READ WRITE	PORT BANDWIDTH (MB/SEC) -- READ -- -- WRITE --		ERROR COUNT	
05	CHP-H 05	0	849 1436	8.63	17.34	0	
07	CHP 6B	1	1681 1395	50.87	10.32	0	
09	CHP 15	0	833 1429	11.96	20.49	0	
0C	CHP-H 64	0	939 1099	0.39	0.50	0	
0D	CHP 6B	0	1328 1823	3.56	12.73	0	
0F	CHP-H 66	0	1496 1675	1.85	2.61	0	
10	CHP 64	0	644 1380	0.03	0.13	0	
13	CHP-H 19	0	907 885	0.58	0.45	0	
16	CU C800	0	1241 738	20.97	5.72	0	
	CU CA00			70.10	3.82	0	
1A	CHP 15	0	1144 1664	0.65	1.18	0	
1B	CHP 0D	0	510 1759	0.12	1.72	0	
1E	CHP-H 05	0	918 894	0.59	0.45	0	
1F	CHP 21	0	1243 1736	0.97	1.70	0	
20	CU E900	0	1429 849	17.66	8.85	0	
	CU E800						
	CU E700						
22	CHP 10	0	923 1753	0.55	2.78	0	
23	CHP 54	0	1805 69	20.80	7.30	0	
24	CHP 64	0	89 1345	0.00	0.00	0	
27	CHP 6B	0	1619 82	0.01	0.00	0	
28	SWITCH 95	270	598 789	50.32	10.56	0	
2B	CHP 70	0	69 2022	0.00	0.71	0	

Interval 000.15.00
 In the last 15 minutes

↓
 This port had a frame to send but did not have any Buffer Credits left to use to send them.

And this happened 270 times during the interval.

And this is an ISL Link!
 Indicators of Potential Buffer Credit Starvation

Fabric with zHPF Enabled



More Reasons For Deploying Switched-FICON

SHARE
Technology - Connections - Results

- Point-to-Point versus switched-FICON Reliability and Availability

Complete your sessions evaluation online at SHARE.org/BostonEval



Reliability versus Availability



- **Reliability is NOT the same as Availability!**
- **Reliability is a measurement of the dependability of the customer's system, fabric and/or devices**
- Often characterized as Mean Time Before Failure (MTBF)
- **Availability is the ability of a system, fabric and/or device to continue to provide services when they are needed, without delay, even if reliability has failed**
- In the data center this is typically discussed as a percentage such as two-9s, three-9s, four-9s or five-9s (99.999%) of availability
- The calculation for this percentage is the length of time that a given system, fabric and/or device will be online and functioning during the course of a years worth of time – it is really an up-time ratio
- So Availability is NOT the same as Reliability and customers do realize that it is very difficult to achieve High Availability unless a system, fabric and/or device does have High Reliability!

Complete your sessions evaluation online at SHARE.org/BostonEval



Availability

- **High availability** refers to a system or component that is continuously operational for a desirably long length of time.
- Availability can be measured relative to "100% operational" or "never failing with a difficult-to-achieve standard of availability for a system or product being "five-9s" (99.999 percent) availability.

	Calculating Percent of Availability Downtime	Seconds	Minutes	Hours	Days
Availability	Downtime per year (31,536,000 seconds in a year) Calculation: (Seconds in year * %) Outages would cause the following max downtime:	31536000			
99.9999%	32 seconds	31.536			
99.999%	5 minutes, 15 seconds	315.36	5.256		
99.99%	52 minutes, 36 seconds	3153.6	52.56		
99.95%	4 hours, 23 minutes	15768	262.8	4.38	
99.9%	8 hours, 46 minutes	31536	525.6	8.76	
99.5%	1 day, 19 hours, 48 minutes	157680	2628	43.8	1.825
99%	3 days, 15 hours, 40 minutes	315360	5256	87.6	3.65

Five-9s is really the ultimate goal for data center deployments.

These figures are just an **"annual risk of outage"** not a guarantee that an outage will actually occur

Customers might go for years before having a problem but that does not matter. Availability is calculated on an annual basis.

Loss of Availability Is Very Costly!

- Few customers know what an hour of downtime costs them!
- Some industry statistics may help to put a context to potential losses from downtime. The numbers differ, depending on the source, but they give some idea of possible impact. Check these yourself on the internet!
- *Even if a vendor **GAVE YOU the hardware**, or if they even **GAVE YOU the Hardware and ALL Maintenance costs** – if it is not deployed as five-9s of high availability, and then their system causes you to take an outage, the hardware cost [maintenance cost] will not offset the cost of the outage!*

Example:

Vendors provide Virtual Fabrics hosted on a single chassis.


If you use two VFs on one chassis for your HA redundancy, but then the chassis firmware fails, both of the VFs fail as well.

All of the I/O paths on these VFs fail which would be poor availability.


INDUSTRY SECTOR	AVG. COST/HOUR*
Energy	\$2.8 Million
Financial Institutions	\$1.4 Million
Information Technology	\$1.3 Million
Insurance	\$1.2 Million
Health Care	\$686,000
Transportation	\$668,586
Utilities	\$643,000
Avg. Cost Per Hour Per Employee	\$205

Source: : META Group, Inc., "Quantifying Performance Loss: IT Performance Engineering and Measurement Strategies"

Availability After A Component Failure




Point-to-Point Deployment of FICON



**...BUT...
Storage Port
Remains
Available!**

- A failure of a FICON CHPID or cable or storage port means that you lose two valuable resources:
 - Channel port will become unavailable AND
 - Storage port becomes unavailable for everyone!
- A failure **anywhere** affects both the mainframe connection and the storage connection
 - The WORST possible reliability and availability is provided by a direct-attached FICON and/or SAN storage topology!



- In a switched-FICON environment, only a connection segment is rendered unavailable:
 - The non-failing side remains available
 - If the storage port has not failed, its port is still available to be used by other CHPIDs
 - If the CHPID has not failed, its port is still available to be used by other storage ports

And There Are Many More Reasons For Deploying Switched-FICON in Mainframe Shops



- Can host both SAN and FICON on the same I/O infrastructure
- Balancing workload across all the ports in a Port Group
- Intermixing Long wave and Short wave FICON Connections As You Desire
- FICON switching devices provide lots of scalability not possible with direct attached FICON
- Consolidation of Channel Cards, CHPIDs and Storage

Summary



- Using a switched-FC infrastructure provides strategic value to operations
- Cost of host I/O adapters and storage I/O adapters can be avoided through consolidation and multi-device access through single HBA, CHIPID and/or storage port
- Channel efficiency can be achieved through Fan In Fan Out topologies
- Storage access and expansion can be scaled with few limitations
- Data protection schemes can be easily incorporated
- For the mainframe, NPIV technology can reduce the overhead of I/O created by Linux when running on System z or zEnterprise

.....Newly Scheduled Presentation.....



Brocade SAN and FICON Update

Please consider attending to discover the innovation of Brocade's Gen 5 Fibre Channel Architecture

This will be my next presentation at SHARE!



Wednesday August 14, 2013 -- 11:00am to 12:00pm -- Session 14482

Please Fill Out Your Evaluation Forms!!

Thank You For Attending Today!



QR Code

This was session:

14275

And Please Indicate on those forms if there are other presentations you would like to see in this track at SHARE!

Your Eval



My Reaction!

- 5: “Aw shucks. Thanks!”**
- 4: “Mighty kind of you!”**
- 3: “Glad you enjoyed this!”**
- 2: “A Few Golden Nuggets!”**
- 1: “You Got a nice nap!”**

Monday August 12, 2013 -- 1:30pm to 2:30pm -- Session 14275