Session 17503

BROCADE

z13 and Brocade Resilient, Intelligent and Synergistic I/O Processing

David Lytle, BCAF
z System Technologies
Principal Engineer / Global Solutions Specialist
August 2015

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 non-infringement of third-party rights with respect to any products and services referenced herein.

Agenda

- Brocade and IBM
- New IBM z13 Resilient and Intelligent I/O
  - 9 Cool Capabilities and Features
- FICON Resources

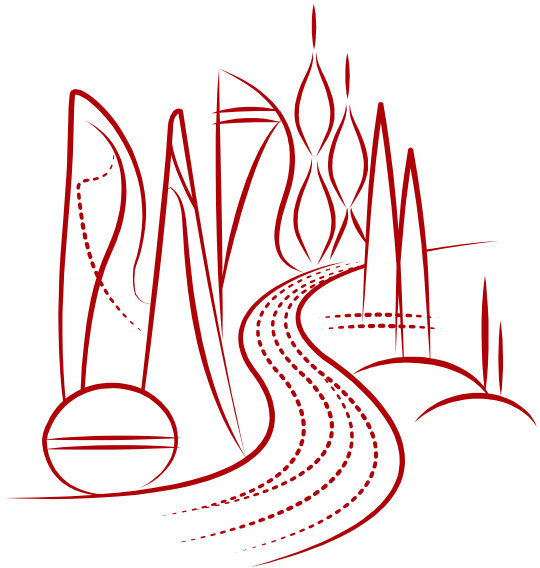

Brocade and IBM

29 Years of Partnering in the Enterprise

IBM z System

Relationship

If you lost an open systems server would your business work continue? **Probably.**

If you lost your mainframe, would your business work continue? **Probably Not.**

BROCADE

IBM z13

Announced on Jan 14, 2015

"91% of new customer-facing apps access data from z Systems" (IBM)

"The same Sys Administrator of 2015 is handling 5 times more MIPS than they were handling in 2005" (BMC)

z13 - A Superlative Processor

With Significant Synergy When Used With Brocade Gen 5 FICON Infrastructure

- 5 GHz Processors (among the world's fastest)
- 141 Configurable Processors (cores) available
- Up to 50% better response times
- Up to 40% capacity improvement over zEC12
- 30% better performance for Linux and Java
- 17x faster analytics when using the DB2-Analytics Accelerator (AA)
- Up to 8,000 virtual machines in one system
- Provides FICON Express8S to attach 2 Gbps storage
- FICON Express16S (98,000 IOPS and 2600 MBps FD with zHPF)
- Up to 160 FICON Express16S cards – 320 CHPIDs
- Many new, advanced functions (e.g. FEC, Fabric Priority)

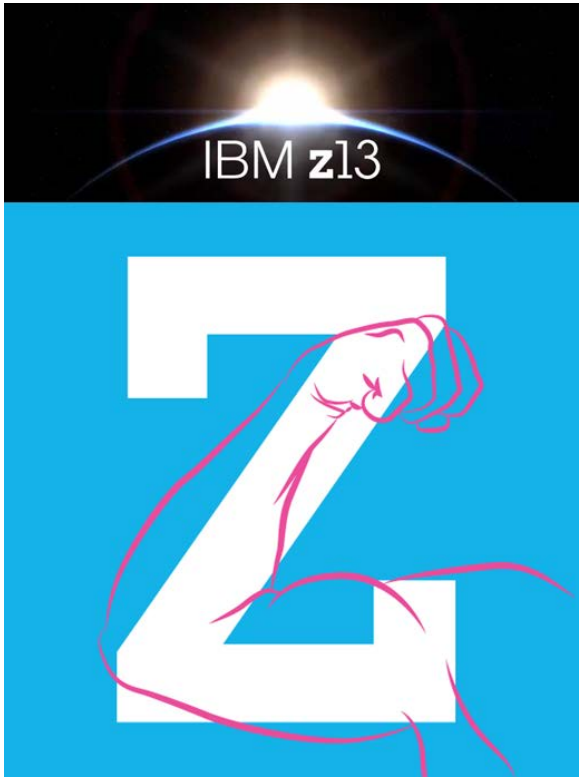

z13 Synergy with Brocade for Resilient and Intelligent I/O

Brocade has provided over 90% of all Mainframe I/O infrastructure (switching and extension) since 1994

6,000+	FICON Directors worldwide
50,000+	Brocade Directors worldwide
200,000+	Brocade SANs in production
39 million+	Brocade SAN switch ports installed

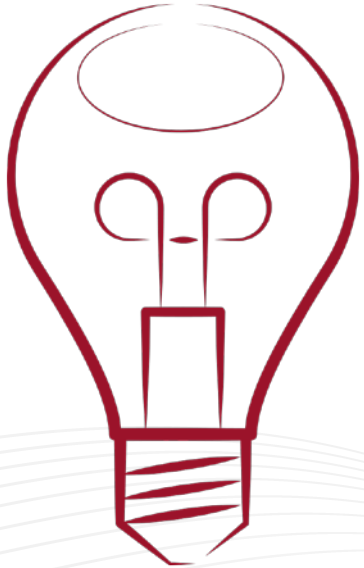

z13 Is A Superlative Processor

9 Cool Capabilities that Demonstrate z13 and Brocade Synergy

1. FICON Express16S works best when attached to Brocade
2. 16Gb Forward Error Correction for robust and efficient channel links
3. SAN Fabric I/O Priority for better control over I/O traffic
4. Brocade CUP Diagnostics for better alerting and troubleshooting
5. Read Diagnostics command to really understand a failing link situation
6. z/OS Health Checker makes use of CUP Diagnostics
7. Link Verification to be sure deployed infrastructures are robust and stable
8. Port Decommission and Recommission for easier infrastructure management
9. FICON Flexible Scalability through Inter-chassis Links (ICLs)

FICON Express16S (FX16S)

Synergy with Brocade at up to 16 Gbps Connectivity

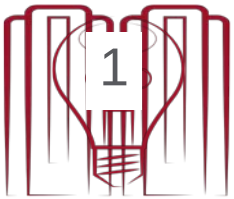

- From 90 (CHPID) to thousands of buffer credits (switch)
- Attached to Gen 5 switching, allows all FX16S features to be capitalized upon:
  - Forward Error Correction
  - Read Diagnostics Parameters
  - I/O infrastructure exercising for optimal deployment
  - NPIV handles up to 8,000 virtual guests using SCSI
- Switch Fan In Fan Out can minimize FX16S ports and channel features required by the user
  - Optimizes each link's performance
  - All CHPIDs run at 16Gbps regardless of storage attachment line rate
  - Enables Fabric I/O Priority with Workload Manager
  - Ensures highest end-to-end availability

FICON Express16S

- With IBM z13 and z/OS v2.1
- Since FOS v7.2.1d

Best Way To Deploy FICON Express16S

Achieve complete value of 16 Gbps channels

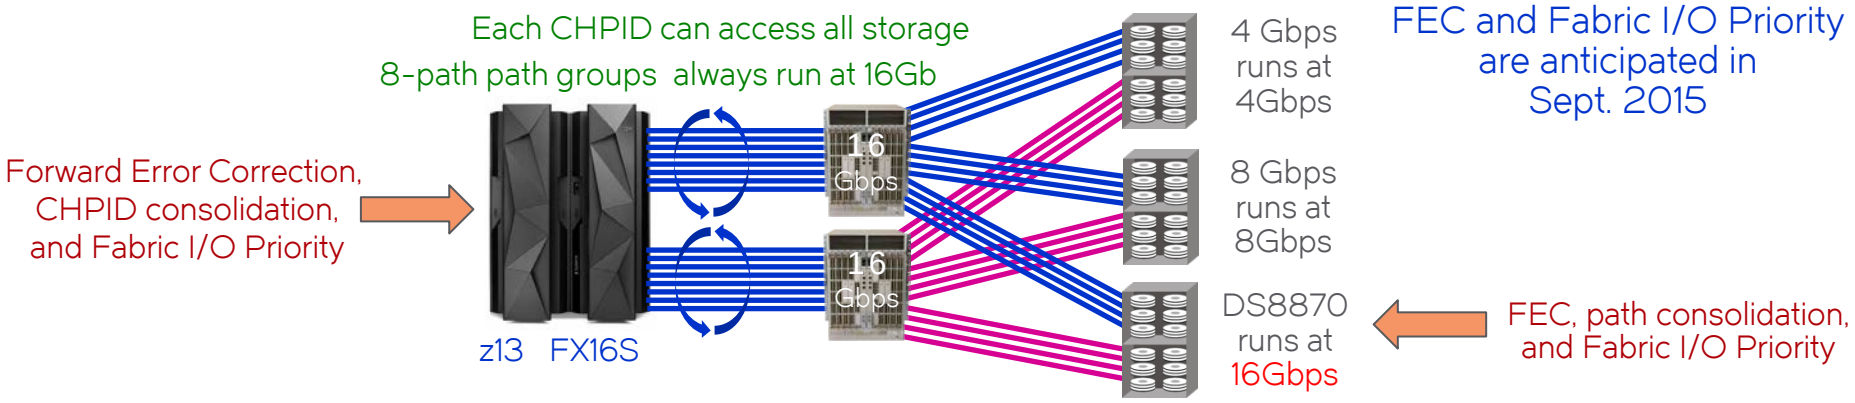

Every asset does more work and the user receives more value for his investment!

- Attaching 16Gb CHPIDs to 16Gb Director allows each CHPID to always run at 16 Gbps:
  - Storage negotiates to 4, 8 and 16Gbps as it attaches to Director – multiplexing to each array
  - 4Gb storage, naturally, will not make as much use of a 16Gb CHPID as 8Gb storage can
    - Fan In Fan Out becomes very useful
  - Switching maximizes the utilization of both host and storage ports in a FICON SAN
  - Currently, Forward Error Correction can only be used with ISL links and not F_Port links

Forward Error Correction (FEC)

Enhanced Signal Condition

- Provides error correction on top of 64b/66b encoding and improves reliability by reducing bit errors (adds equivalent of 2.5 db of signal strength)
- Ensures high data reliability
- Guarantees higher and more deterministic performance
- Provides robust data security at these extremely high data rates.
- Prepares for Gen 6 infrastructures featuring higher speeds of 32GFC and 128GFC, which must be FEC enabled

- Supported from FOS v7.0 (for z13 requires v7.3.1b+)
- Supported from z13 z/OS 2.1 GA 1.5 (FX16S only)
- Anticipated in September 2015

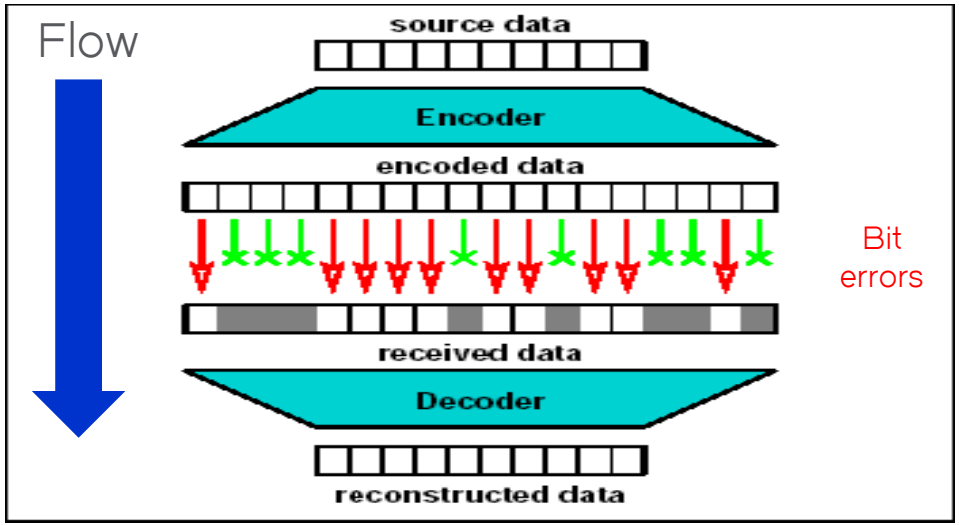

IBM Provides this on FICON Express16S and DS8870 16G Links
Brocade provides this on 10G and 16G ISL Links

Forward Error Correction (FEC)

Cascaded and non-Cascaded Infrastructures

- These diagrams show how FEC is deployed by Brocade and IBM (anticipated in Sept 2015)

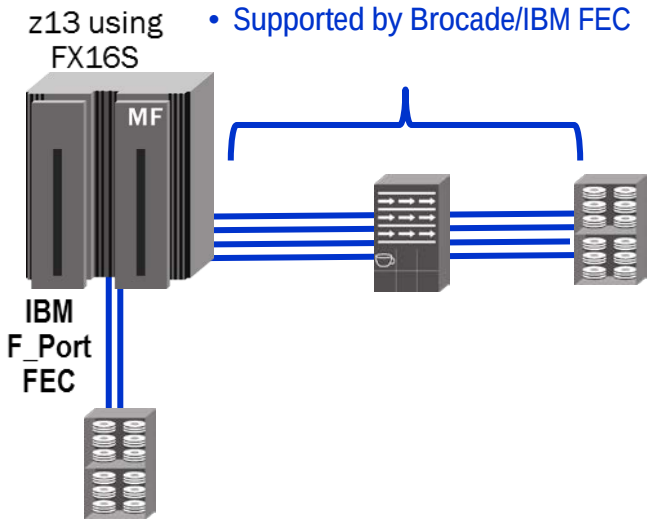

IBM claims that FEC is designed to provide the *equivalent* reliability improvement as doubling the optical signal strength.

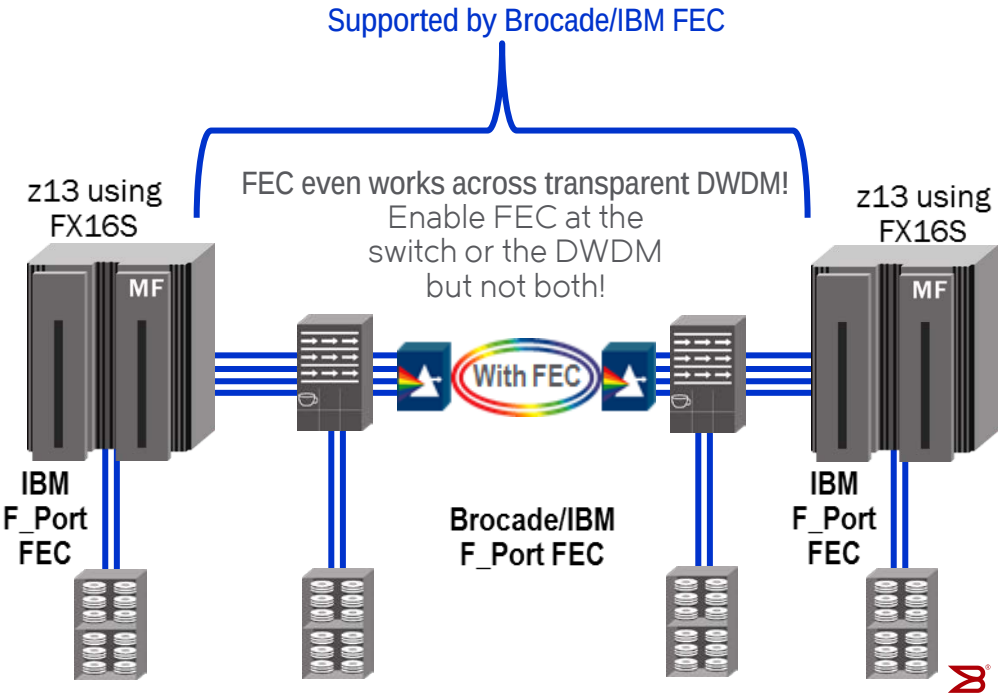

SAN Fabric I/O Priority

Application Driven Quality of Service

- Translate application importance into Fabric Priority (Quality of Service)
- Fabric I/O Priority works in conjunction with z/OS workload manager in Goal Mode
- Fabric I/O Priority (QoS) is preserved in each frame – CS_CTL bits

- Supported on z13 only
- Supported from FOS v7.3.1b
- Supported from z/OS v2.1 GA1.5
- Availability is 30 September 2015

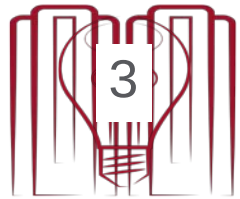

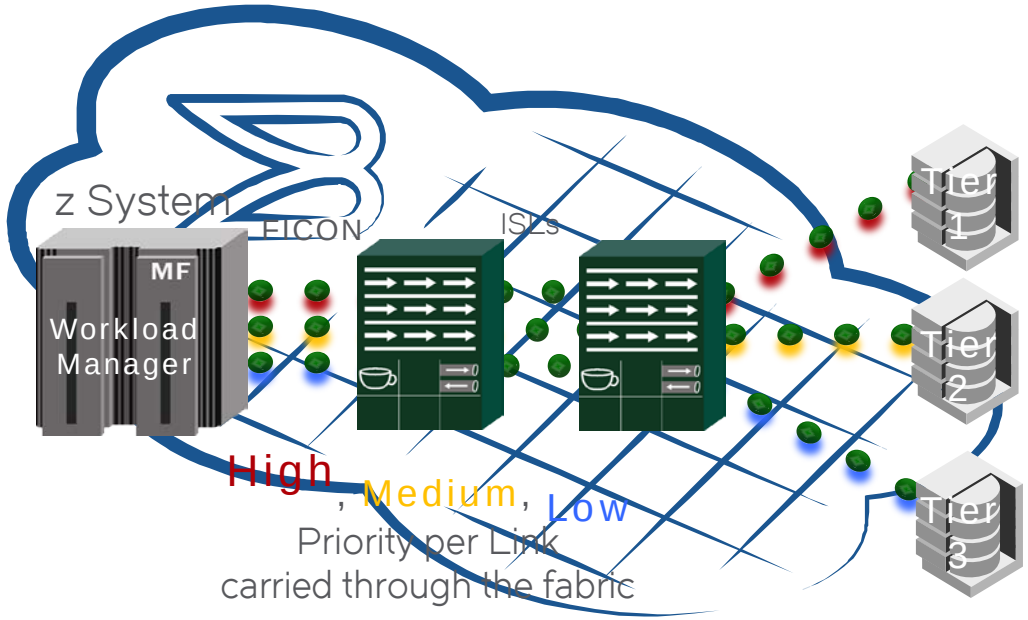

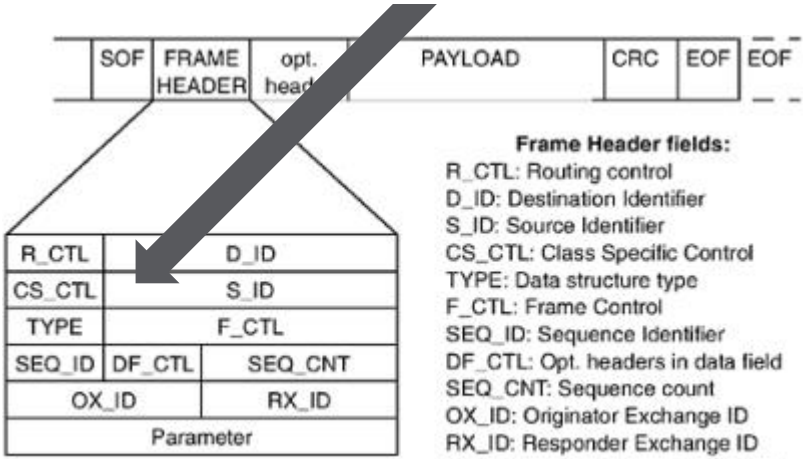

Frame Header fields:

R_CTL: Routing control
D_ID: Destination Identifier
S_ID: Source Identifier
CS_CTL: Class Specific Control
TYPE: Data structure type
F_CTL: Frame Control
SEQ_ID: Sequence Identifier
DF_CTL: Opt. headers in data field
SEQ_CNT: Sequence count
OX_ID: Originator Exchange ID
RX_ID: Responder Exchange ID
Parameter: Frame type dependent

IBM Provides this on FICON Express16S and DS8870 16G Links
Brocade provides this on 10G and 16G ISL Links

Fabric I/O Priority providing QoS via CS_CTL bits

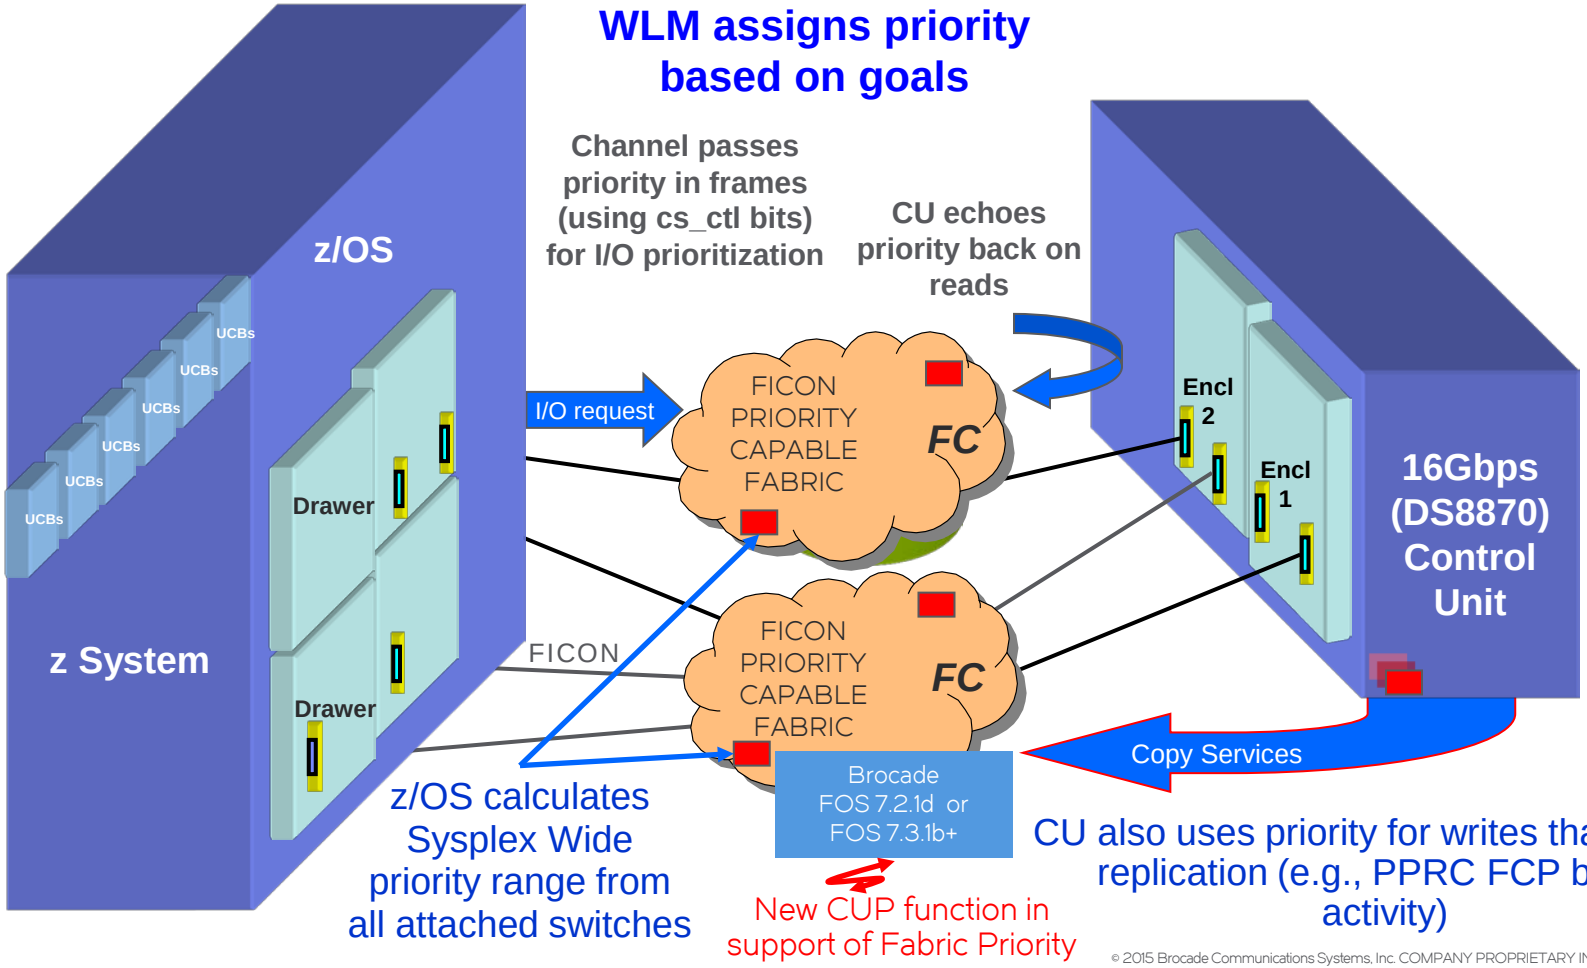

Brocade CUP Diagnostics

Route Information and SFP Health Checking

We work with IBM to provide unique, additional capabilities to the mainframe environment, such as CUP Diagnostics, which provides information about:

- Health - SFP optics information
- Single Points of Failure
- Flow Descriptions:
  - ROUTE=TODEV
    - Show the path through the fabric from the channel to the device
  - ROUTE=TODEV,HEALTH
    - Adds SFP power levels, transmit/receive utilization statistics, and error counts to the report
  - ROUTE=FROMDEV
    - Show path from the device to the channel
  - ROUTE=FROMDEV,HEALTH

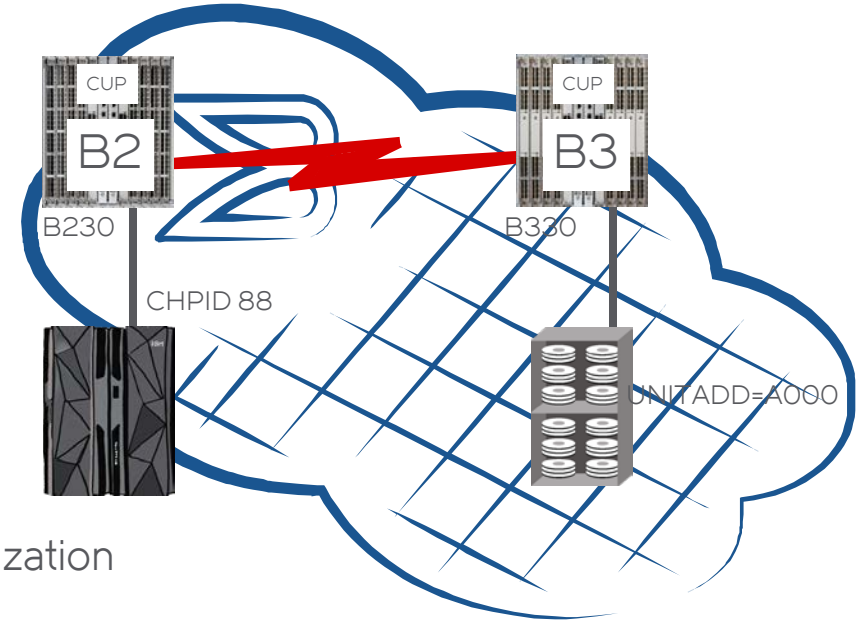

D M=DEV(A000,(88)),ROUTE=TODEV,HEALTH

CUP Diagnostics

Command Output Example - Routing Information

```
D M=DEV(A000,(88)),ROUTE=TODEV,HEALTH

DEVICE 0A000   STATUS=ONLINE

Source to destination routing information follows:

Switch Domain=B2, Type=Source Director
                  Group
Port Type  From   To     Agg Dyn Speed Misc
30  Entry Chan  ------------------------
5B  Exit  B230   B30A   ------------------

Switch Domain=B3, Type=Destination Director
                  Group
Port Type  From   To     Agg Dyn Speed Misc
0A  Entry B25B   B301   ------------------
01  Exit  --------------------------------
```

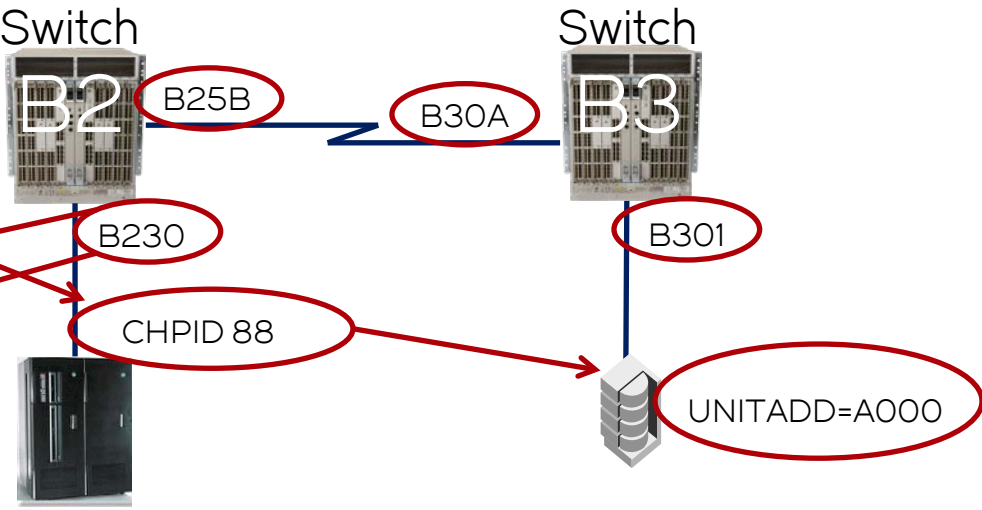

CUP Diagnostics – Find Optical Signal Specs in MyBrocade

Command Output Example – Health Information

Health information follows:

Fabric Health=No health issues

Switch Domain=B2, Health=No health issues

Port	Health	%Util Trn/Rcv	%Delay Trn/Rcv	Error Count Trn/Recv	Opt Signal Trn/Recv
5B	PORT NORMAL	0/0	0/0	0/0	-1.8/-2.2

Switch Domain=B3, Health=No health issues

Port	Health	%Util Trn/Rcv	%Delay Trn/Rcv	Error Count Trn/Recv	Opt Signal Trn/Recv
0A	PORT NORMAL	0/0	0/0	0/0	-2.0/-1.3
01	PORT NORMAL	0/0	0/0	0/0	-1.8/-2.8

Interface Verification

SFP Health through Read Diagnostics Parameter

- New z13 System Channel Subsystem Function
- Similar to CUP Diagnostics D M= (Health)
- This is z/OS Read Diagnostic Parameters (RDP)
  - Created to enhance path evaluation
- Automatically differentiate between errors caused by dirty links and those errors caused by failing optical components
- Provides the optical characteristics for the ends of the link:
  - Enriches the view of Fabric components
- z/OS Commands can display optical signal strength and other metrics without having to manually insert light meters.

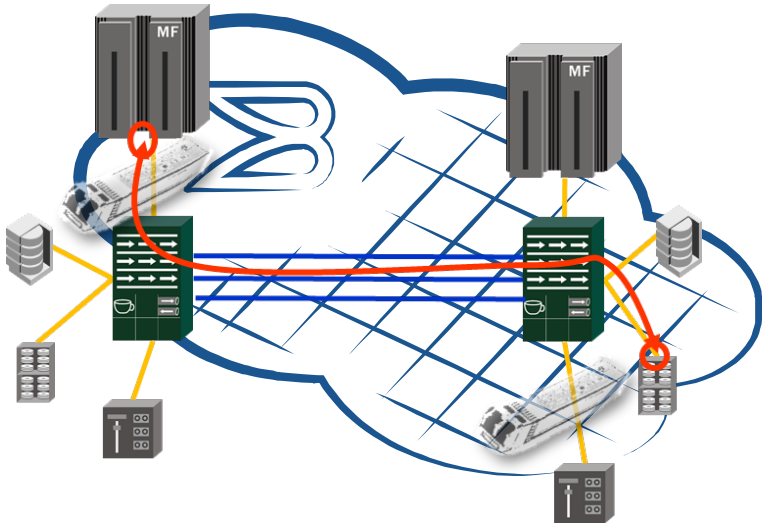

A T11 committee standard

- Supported from FOS v7.3.1b5
- Supported by z13 GA2.0

zHealth Checker

Leverages CUP Diagnostics

- FICON fabric issues have resulted in elusive and unacceptable I/O service times that negatively affect application response time:
  - RMF Device Activity reports show high average service times
  - RMF Queuing reports show high “initial command response” time
- zHealth Checker tests Single Points of Failure:
  - Analyze connections and paths
  - Identify common components
- zHealth Checker provides a Flow Description:
  - Identify fabric routes
  - Examine utilization statistics
  - Assess performance and errors

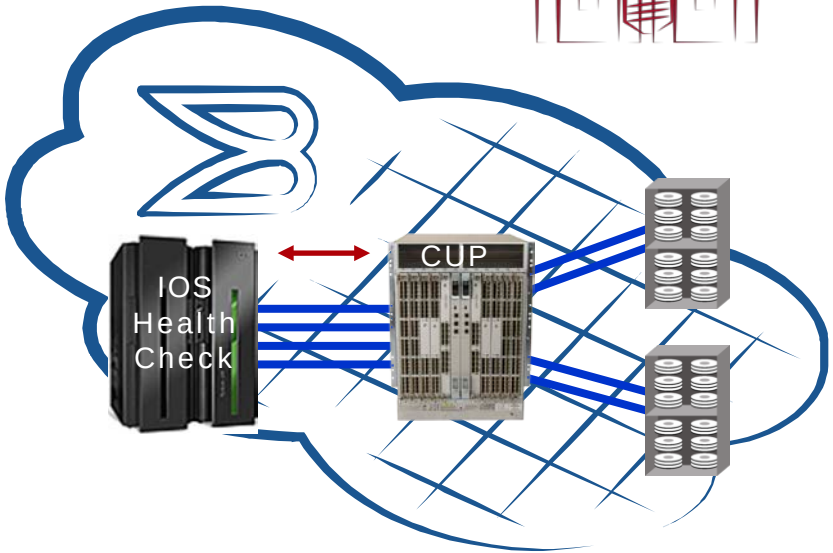

Detecting inconsistent CMR

- Since IBM z/OS v1.7
- Since FOS v7.1

ISL Link Verification

New Diagnostic Port (D_Port)

- Ensures optical and signal integrity for Gen 5 Fibre Channel optics and cables, simplifying deployment and support of high-performance fabrics.
- It leverages ClearLink Diagnostic Port (D_Port) capabilities of Gen 5 Fibre Channel platforms.
- Integrates path evaluation with Fabric technology
- Exploits advanced optical connector functions
- Supports full path evaluation prior to deployment

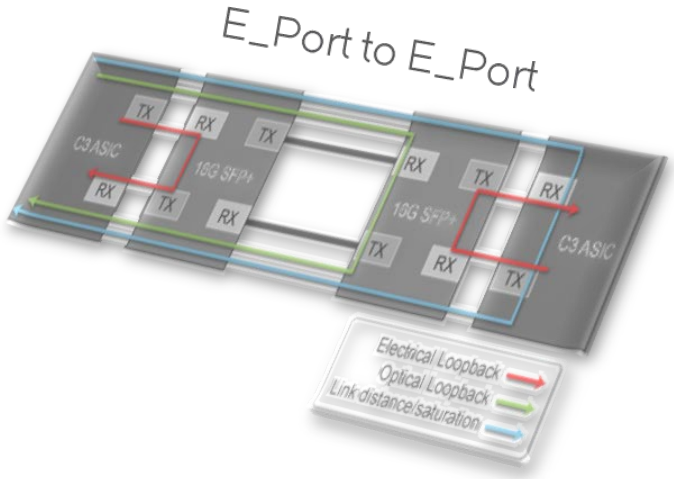

- Brocade D_Port since FOS v7.1

ClearLink Diagnostics Functional Details

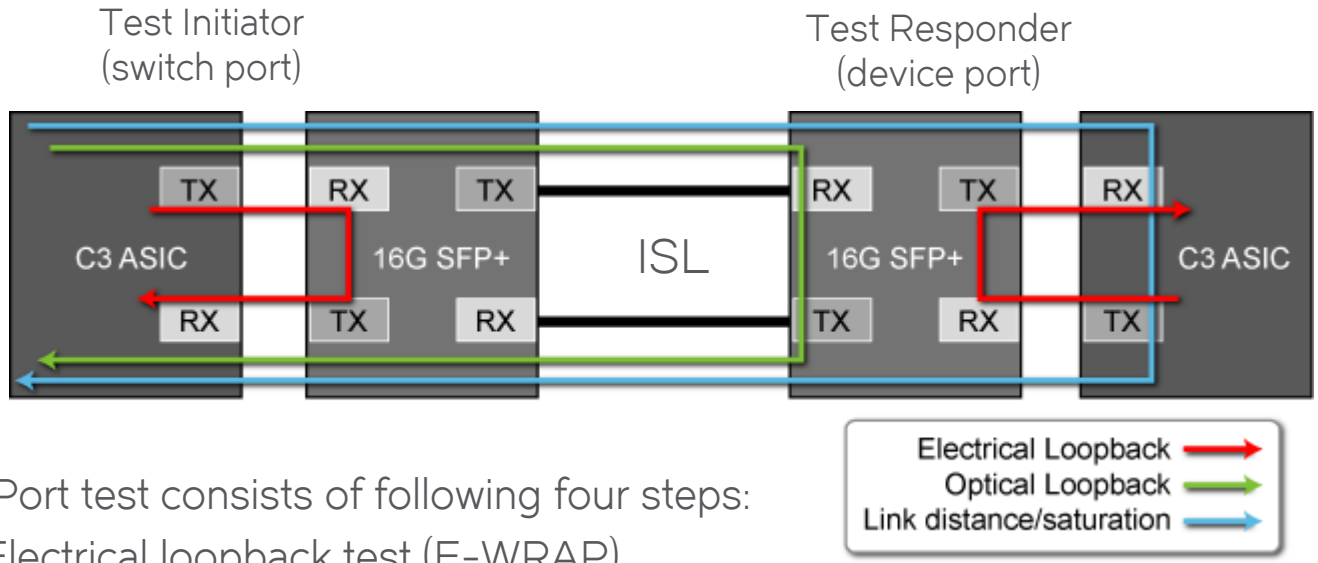

- D_Port test consists of following four steps:
  - Electrical loopback test (E-WRAP)
  - Optical loopback test (O-WRAP)
  - Link traffic test
  - Link latency and distance measurement

Complimentary Mainframe Link Verification

Use the IBM Mainframe I/O Exerciser in conjunction with D_Port Diagnostics

- Integrates channel evaluation with ISL evaluation
- Exploits advanced optical connector functions
- Supports full path evaluation prior to deployment
- IBM I/O Exerciser simplifies the chore of exercising its I/O connections before bringing up z/OS and running production work.
  - http://www14.software.ibm.com/webapp/download/preconfig.jsp?id=2014-03-04+12%3A13%3A17.385915R&S_TACT=&S_CMP=
- Complimentary capability to Brocade's ClearLink Diagnostics which checks and exercises ISL connections between switches before bringing up and running production work.

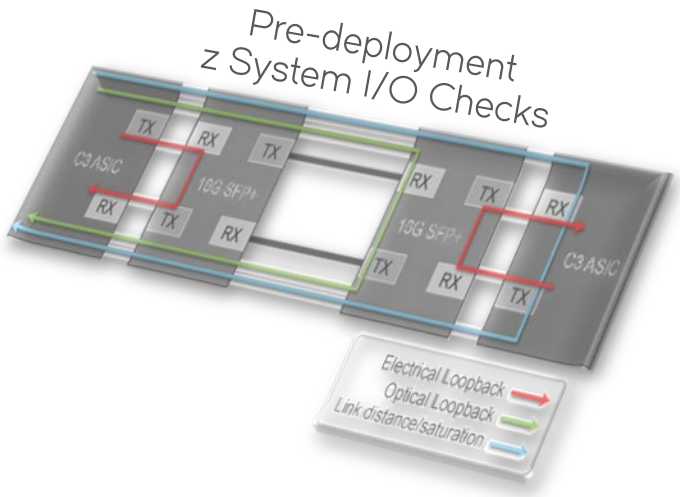

- IBM I/O tool became available March 4, 2014
- Brocade D_Port since FOS v7.1

Gentle Handling of ISL Links

Port Decommission / Recommission

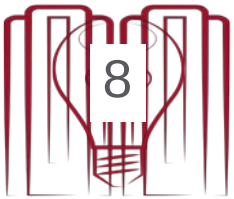

- Mechanism to decommission active ports
  - ISLs and F_Ports
  - Allow host to quiesce paths
  - Non-disruptive to active paths
- Cooperation between M/F and Brocade:
  - Interface with application, system manager, or z/OS:
    - Moves workload off of a target port before that port is disabled
- Integrate System Automation:
  - Select Port Decommission as an action
  - Invokes decommission process instead of immediate disabling of port

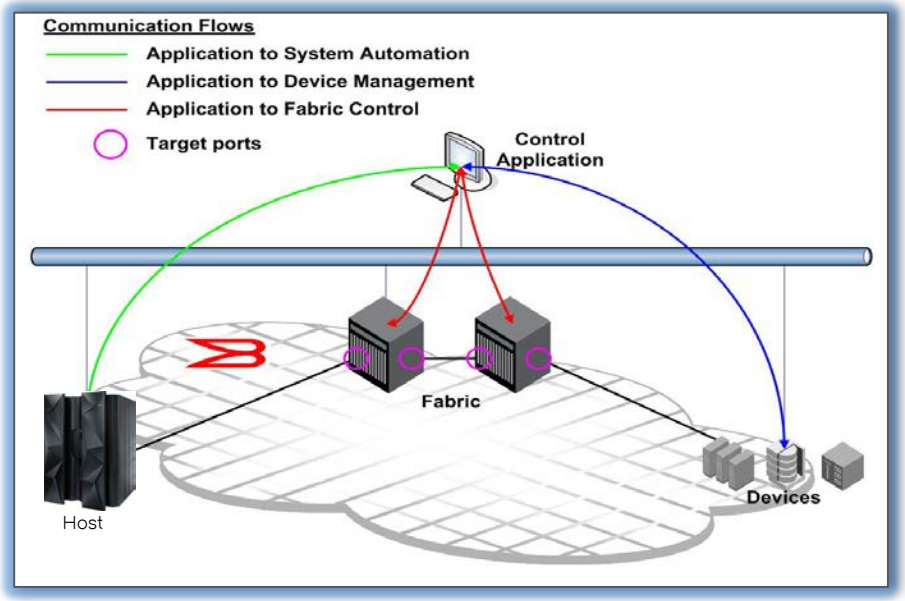

- E_Port Decommissioning since FOS v7.0
- F_Port Decommissioning since FOS 7.1

E_Port Decommissioning

FICON ISL Management for Gen 4 and Gen5

- Coordinate event with external applications:
  - Switch operating system moves routes off of the target ISL before that ISL is disabled
- Mechanism to remove an ISL non-disruptively:
  - Block/Disable an ISL port after moving the traffic flow to other routes so that removing it will be non-disruptive
- Recommissioning E_Ports:
  - Fix SFP or cable or port or path problem
  - Should run ClearLink D_Port diagnostics
  - Should run Flow Vision's Flow Generator
  - Re-enable the E_Ports on both switches

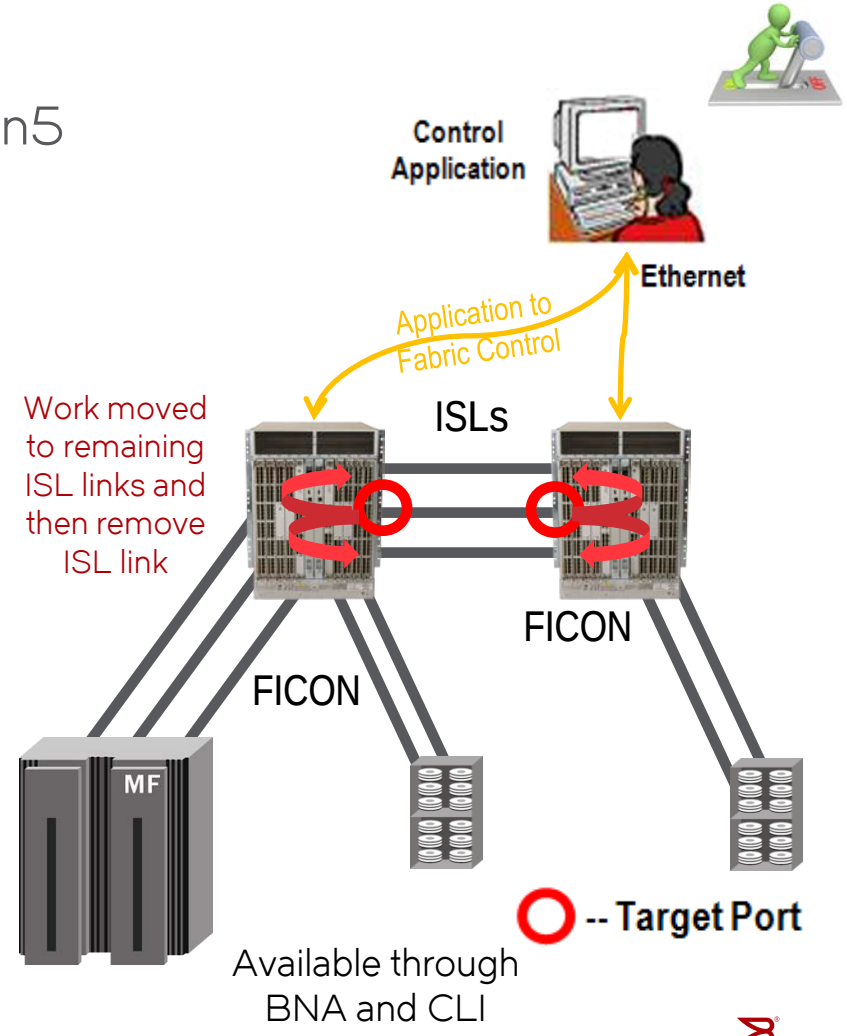

F_Port Decommissioning

FICON F_Port Management for Gen 4 and Gen5

- Mechanism to remove a node port non-disruptively:
  - Block/Disable a device port after allowing each LPAR to quiesce the path/device (CIMOM agent) so that removing it will be non-disruptive
- Coordinate event with external applications
  - Application or system manager
    - Moves workload off of a target port before that port is disabled
- Some customers have told us that now that they've implemented this they love it and use it all the time and never want to be without it again!

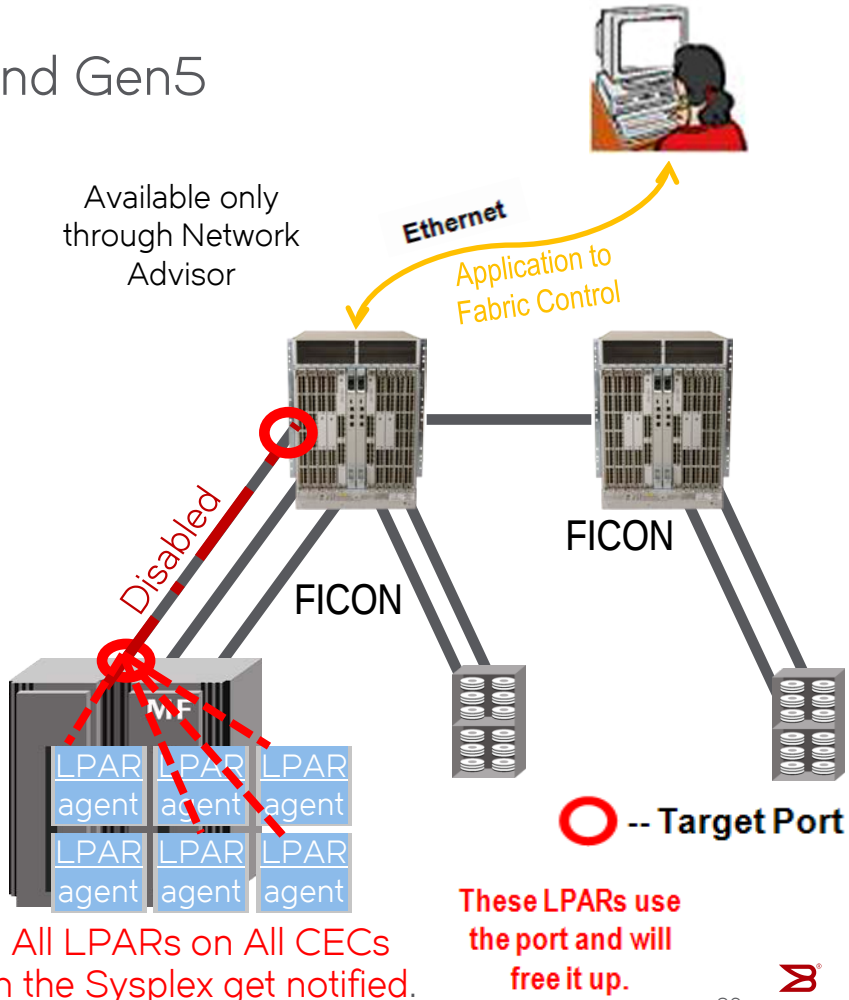

UltraScale Inter-Chassis Links (ICLs)

FICON Scalability Within a Data Center – since 2008

- Director-class switch feature only (not for 6510)
- Provides short-distance connectivity between two DCX family chassis for FICON and/or FCP:
  - 2m for 8G DCX and 50m to 2km for 16G DCX
- For customers to build a powerful core without sacrificing device ports for Inter-Switch Links (ISL).
- Minimizes latency between chassis:
  - Lowest-latency switching via a backplane vs ISLs
  - Does not count as a hop for FICON but is cascading
- Maximizes load balancing and availability

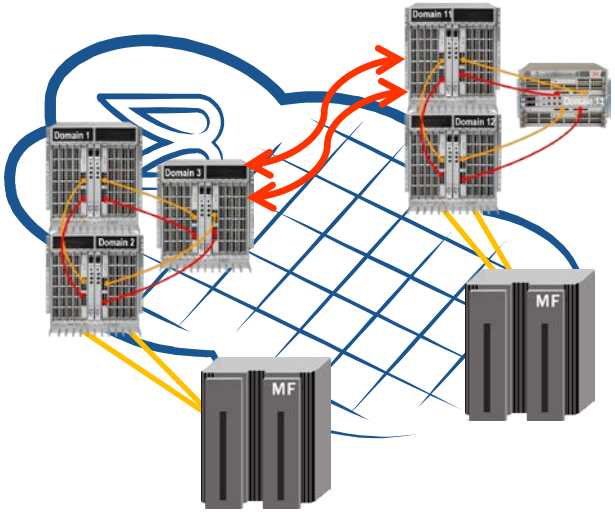

Up to 1,536 ports in a single FICON Fabric
(256p x 3c with ISLs to 3c x 256p)
All ports available for user connectivity!

Typical FICON Fabric Scalability

8 Gbps DCX introduced the innovative ICL

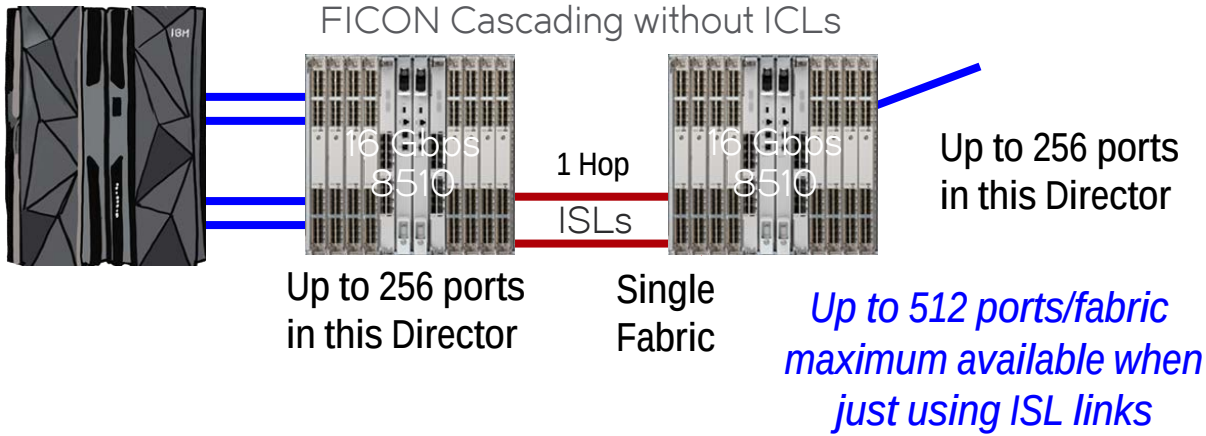

- Until the advent of the 8 Gbps DCX, FICON Cascading was limited to ISL connections between a pair of FICON switches and/or Directors -- FICON is only allowed 1-hop
- The innovative DCX developed a unique capability to connect our "core" blades together which is like an extended backplane – and the Gen 5 8510 carries on that tradition but it is much improved
- Now, with Gen 5, users can ICL chain as many as 3 Gen 5 Directors together in a pod and still have only a single fabric hop and then use ISLs to connect several pods together

Inter-chassis Link Scalability

2 Chassis Chain

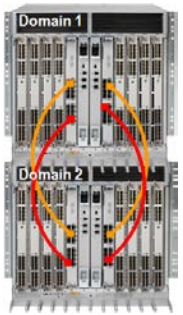

4 ICLs

4 ICLs

An ICL=64 Gbps
Four ICLs run at a max of 256 Gbps

ICLs are NOT counted as a Hop

Ingress

256 FICON ports
512p per ICL pair

No Hop
Equals 16x 16G ISLs

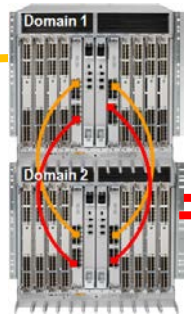

Single Fabric 1 Hop

16G ISLs

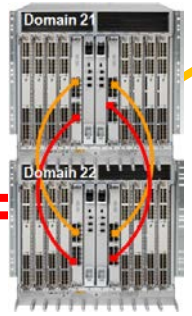

Egress

256 FICON ports
512p per ICL pair

1,024 ports/fabric at 16 Gbps!

3 Chassis Chain

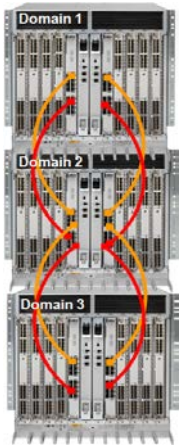

4 ICLs

8 ICLs

4 ICLs

Four ICLs at max 256 Gbps
Eight ICLs at max 512 Gbps

ICLs are NOT counted as a Hop

Ingress

256 FICON ports
512p per ICL pair
768p total

No Hop
Equals 32x 16G ISLs

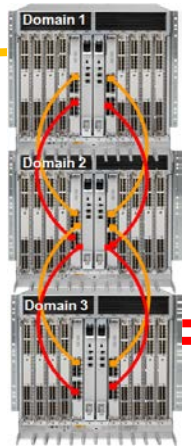

Single Fabric 1 Hop

16G ISLs

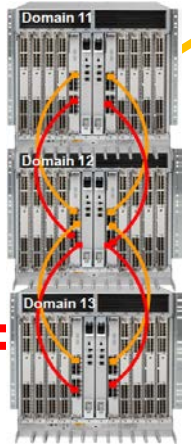

Egress

256 FICON ports
512p per ICL pair
768p total

1,536 ports/fabric at 16 Gbps!

Inter-Chassis Links (ICLs) For Extension

Exponentially expanded for Gen 5

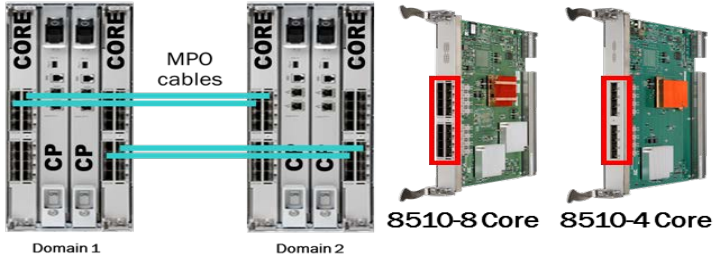

Core routing blades with optical UltraScale ICL ports

UltraScale ICL optics
QSFP (4×16 Gbps)

- 32 ICL ports available on the DCX 8510-8
- 16 ICL ports available on the DCX 8510-4
- A minimum of four ICL ports (two on each core blade) must be connected between each chassis pair.
- Uses QSFP (4 serial data paths per port)
- FOS before v7.4 uses MPO cables to break out the data paths at the receiving end
- FOS v7.4 and higher uses Parallel Single Mode cables when attaching out to 2 km

Resources

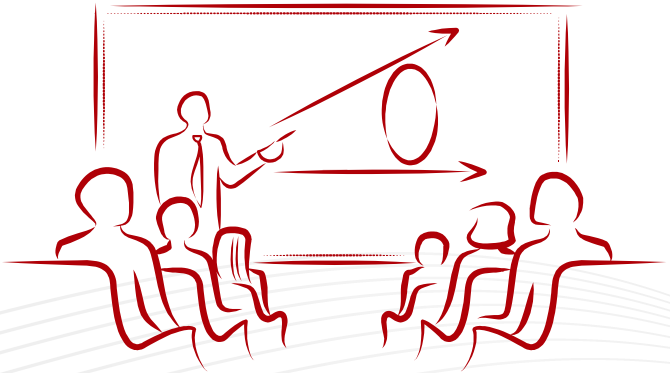

Mainframe eBook
Free at Brocade.com

- This book is the “one stop to shop” source for information on all aspects of System z connectivity.
- The book, available as an e-book, is free of charge, and can be downloaded at the below link.
- Please share the link with your peers and storage vendors instead of sharing the book via file attachment as we would like to know how popular the book really will become.
- Brocade Mainframe Connectivity Solutions.

Resources

Brocade and IBM SAN Solutions

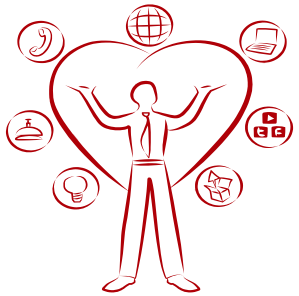

- SAN Fabric Technology Overview
  - www.brocade.com/sanfabric
- The Network Matters for Storage
  - www.brocade.com/launch/the-network-matters/index.html?intcmp=lp_networkmatters_bn_00006
- Gen 5 Fibre Channel Products
  - http://www.brocade.com/products/all/san-backbones/product-details/DCX 8510-backbone/index.page
  - http://www.brocade.com/products/all/switches/index.page?network=FIBRE_CHANNEL
- Brocade SAN Playlist on YouTube
  - http://www.youtube.com/playlist?list=PLTeNsFQA8JYc6xxqZxoKhKiZ2MMMjPU5u
- IBM- Enhancing Value to Existing and Future Workloads with IBM z13
  - http://www.redbooks.ibm.com/abstracts/redp5135.html?Open
- IBM – Get More Out of Your IT Infrastructure with IBM z13 I/O Enhancements
  - http://www.redbooks.ibm.com/abstracts/redp5134.html?Open

Additional Resources For You To Use

Visit Brocade's Mainframe Blog Page at:

http://community.brocade.com/t5/Mainframe-Solutions/bg-p/MainframeSolutions

Almost 300,000 hits

Also Visit Brocade's New Mainframe Communities Page at:

http://community.brocade.com/t5/Mainframe-FICON-Solutions/tkb-p/MainframeFICONSolutions

You can also find us on Facebook at:

https://www.facebook.com/groups/330901833600458/

- www.linkedin.com Groups

EOF

Session 17503

""Brocade presented us with a different perspective on SAN architecture. As we learned about the technical differences that Brocade offered, we realized that we could achieve our goals with less infrastructure, cabling, and complexity."

Muhammet Haydar Ertek, Storage Systems Manager at Halkbank, Istanbul, Turkey