Session 13011

#SHAREorg

SHARE
Technology · Connections · Results

QR Code

SAN 201 Further on Storage Area Networking FICON (FIber CONnection) and a little SAN (Storage Area Networking)

David Lytle – Brocade – dlytle@brocade.com
Mike Blair – Cisco – mblair@cisco.com

SAN 201

SHARE in San Francisco 2013

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!

Agenda for Session 13014

1st Session 13014 – 11:00am – 12:00pm

- Types and Components of Storage
- Let's talk Fibre Channel
- FC Buffer Credits
- Fabric Routing / Virtual Fabrics / Partitioning
- Security / Zoning

…. then continued in the 2nd session…

2nd Session 13011 – 3:00pm – 4:00pm

- **History**
- **Mainframe Terminology, Connectors, Cables, and Wavelengths**
- **Addressing in FICON**
- **ESCON Status, zHPF and NPIV**
- **Buffer Credits, CUP, RMF, BC/DR Solutions**

Brocade, Cisco and IBM Graphics are used throughout this presentation.

Agenda for Session 13011

Session 13011

- **History (DAVE)**
- **Mainframe Terminology, Connectors, Cables, and Wavelengths**
- **Addressing in FICON**
- **ESCON Status, zHPF and NPIV**
- **Buffer Credits, CUP, RMF, BC/DR Solutions**

Mainframe Time Lines

Exciting things in SAN today has been on System z for years

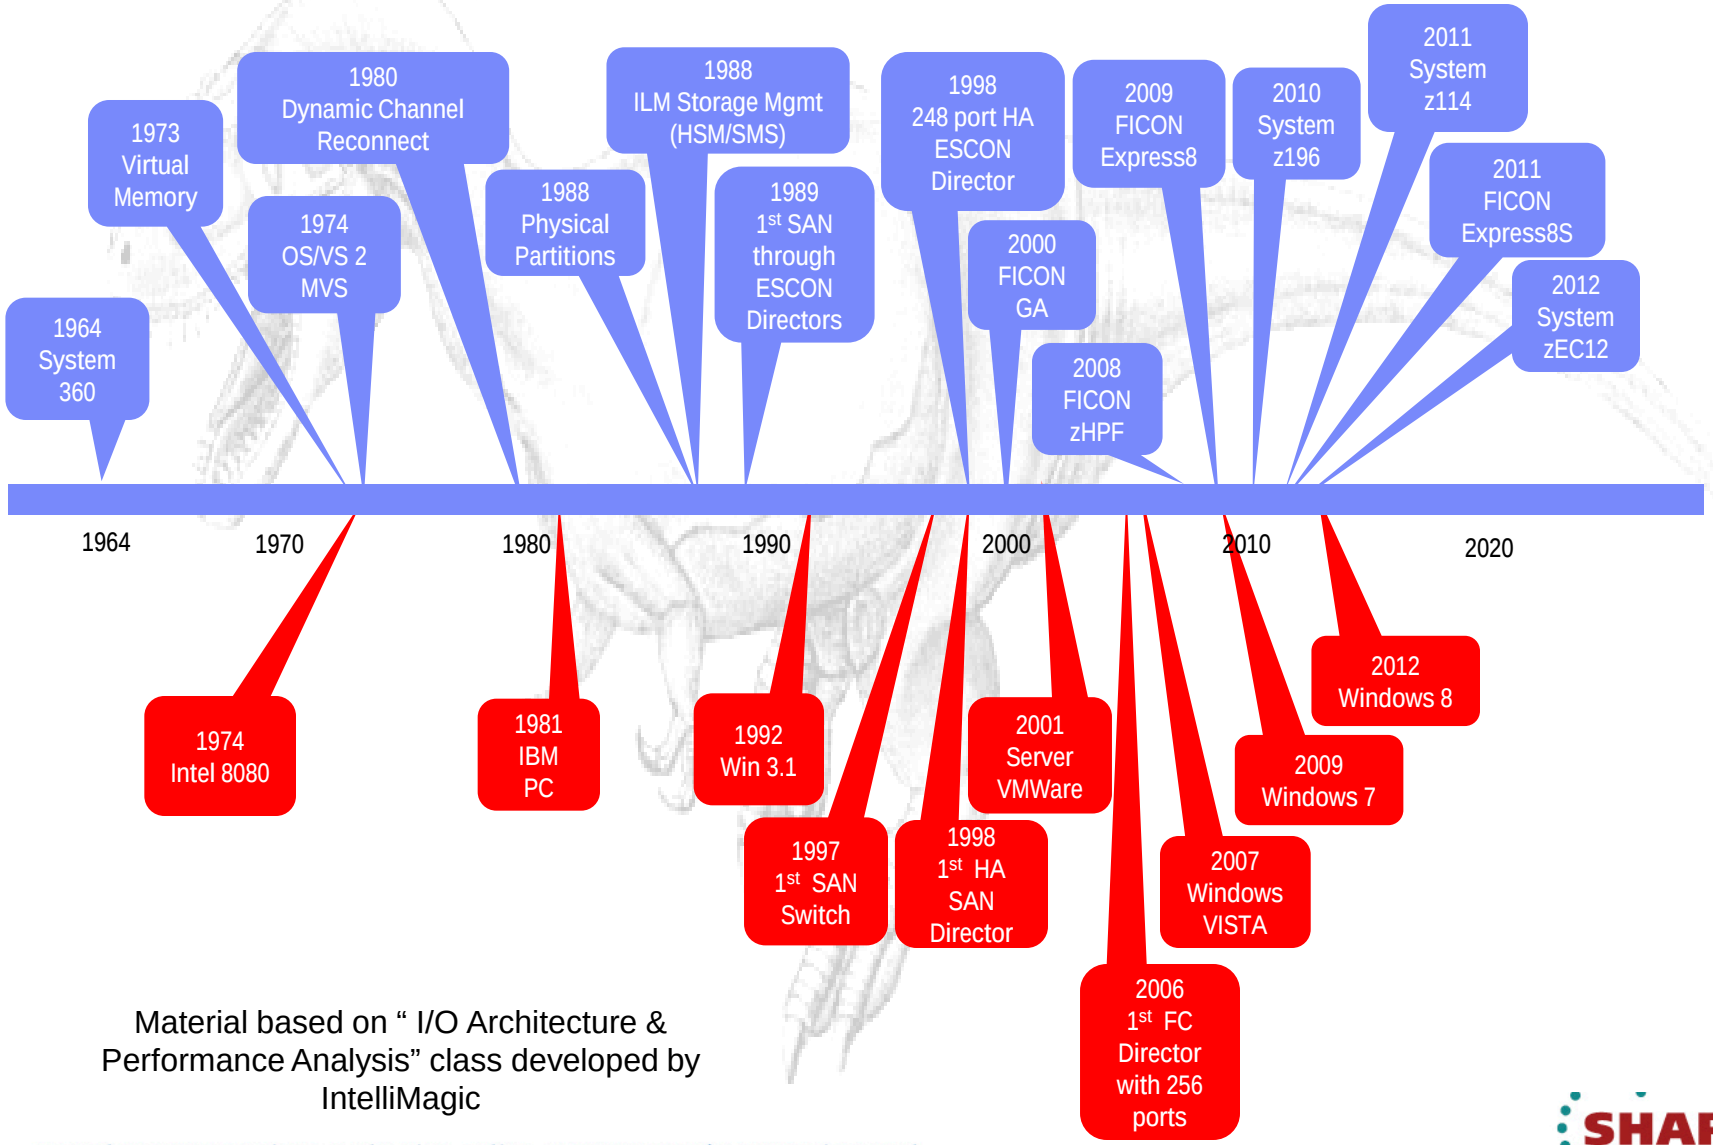

Material based on " I/O Architecture & Performance Analysis" class developed by IntelliMagic

Agenda for Session 13011

Session 13011

- History
- **Terminology, Connectors, Cables, and Wavelengths**
- Addressing in FICON
- ESCON Status, zHPF and NPIV
- Buffer Credits, CUP, RMF, BC/DR Solutions

SAN Terminology -- *Definitions*

- "Fibre" is the protocol/architecture used to transport frames
  - As in "the fibre channel protocol"
- "Fiber" is the glass cable used to attach connections
  - As in "the fiber cable"
- Allows up to 16 million nodes of connectivity
- Historically has used 8-bit/10-bit encoding/decoding to translate an 8-bit byte of information to 10-bit format for serial transmission within the payload of a fibre channel frame
  - But 10Gbps and 16Gbps utilize 64b66b for more efficiency
- Variable frame size with a maximum user payload of 2112 bytes.
- Supports transfers of files that are blocked at up to 128MB in size.
- Supports full duplex transmission
  - Simultaneous send and receive

SAN Terminology -- *Fiber Channel Link*

- Light wavelengths in fiber are expressed in nanometers
- Speed of light in fiber cable is 2/3^{rds} the speed of light in a vacuum
  - Light travels at ~5 nanoseconds per meter (3.3 ft) of distance in glass
- **Multimode fiber** is used for numerous frequencies which are all short-wave frequencies (62.5, 50 micron) of laser light.
  - Always used with short wave optics (transceivers)
  - Used for local distance connectivity (~33-1,640 feet...or...10-500 meters)
- **Single-mode fiber** has a smaller core that allows only one frequency of light (9 micron) which is long-wave laser light.
  - Always used with long wave optics (transceivers)
  - This is used for longer distance connectivity (up to 15.5 miles or 25 km)
- Optical power budgets, or link loss budgets, measured in decibels (dBs), are used to manage optical signal loss.

Why Customers Use Mainframes

- Organizations that run an IBM System z platform have unique business requirements.
- They need a **very powerful, highly available, robust and scalable environment** that can support high-volume transaction processing with demanding batch windows and large, critical application portfolios.
- There are tremendous opportunities (like private cloud computing) for leveraging the vast computing power of System z to help greatly simplify the computing environment while substantially reducing overall costs and power consumption.
- All of this makes System z a critically important platform for the future of large enterprise-scale computing!

Mainframe Terminology

System z

This is one CHPID

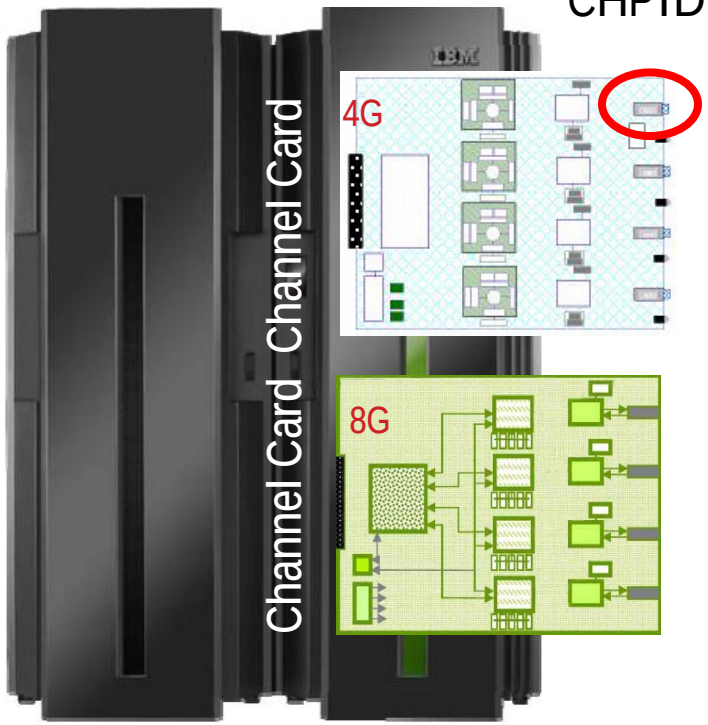

Up to 336 FICON CHPIDs on z9 / z10
Up to 320 FICON CHPIDs on z196
Up to 128 FICON CHPIDs on z114
Up to 320 FICON CHPIDs on zEC12

All CHPIDs on a channel card must be the same -- LX or SX – no mixture

Channel or Channel Path Identifier (CHPID)

A physical connectivity port which is embedded in the mainframe processor's channel system

A CHPID will be either long wave or short wave and depending upon the installed channel card it could be 1Gbps, 2Gbps, 4Gbps or 8Gbps referred to as FICON Express or FICON ExpressX (where X is 2, 4 or 8) or FICON Express8S

Mainframe Terminology

System z Mainframe (M/F)

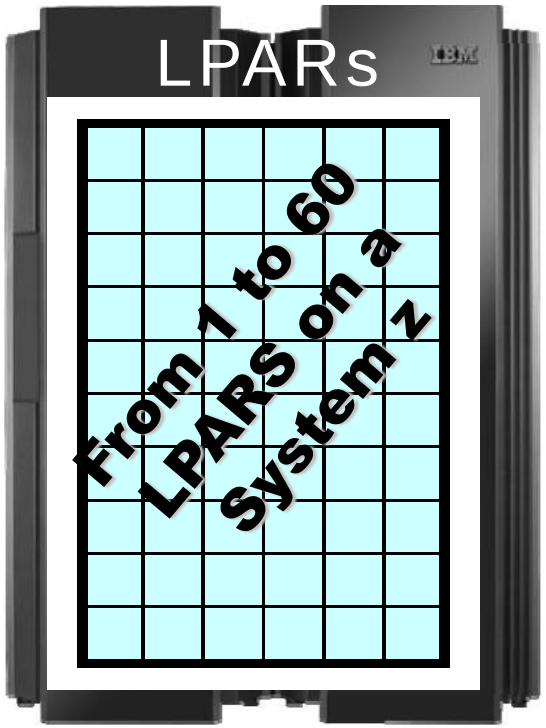

Each LPAR runs its own operating system and has some number of processors assigned to it as CPs and specialty engines

Logical Partitions (LPARs)

- Available since 1988 this is a way of dividing up a mainframe's capacity into Logical PARtitions and isolating each LPAR from every other LPAR to ensure RAS.
- System 390 and zSeries could run up to 15 partitions per physical mainframe.
- System z can run up to 60 partitions per physical mainframe.

Mainframe Terminology

Channel Path

The fiber between the channel and the storage subsystem, as well as the interface adapter on the subsystem and any interleaving directors

A channel path can be a Direct Attached (Point-to-Point) path or it can be a switched path

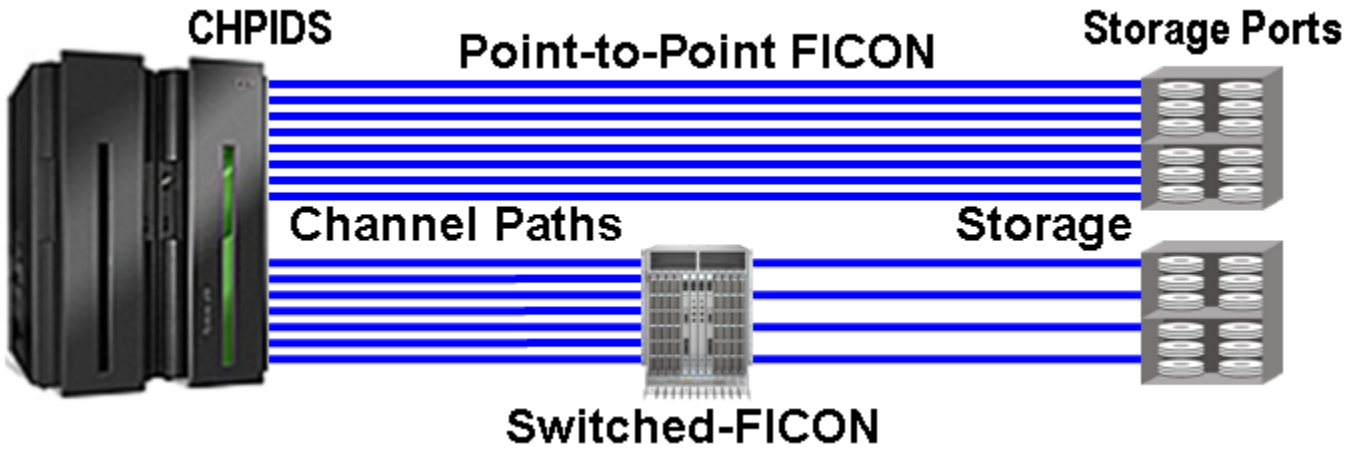

Mainframe Terminology

ESCON Channel Paths – 1990

- This old protocol provides a circuit switched, unidirectional data transfer mechanism.
- Once proprietary, it did become a FC standard – SBCON
- Once a data transfer for an I/O from channel to subsystem or subsystem to channel has begun, no other I/O operations can employ the channel until that transfer has been completed.

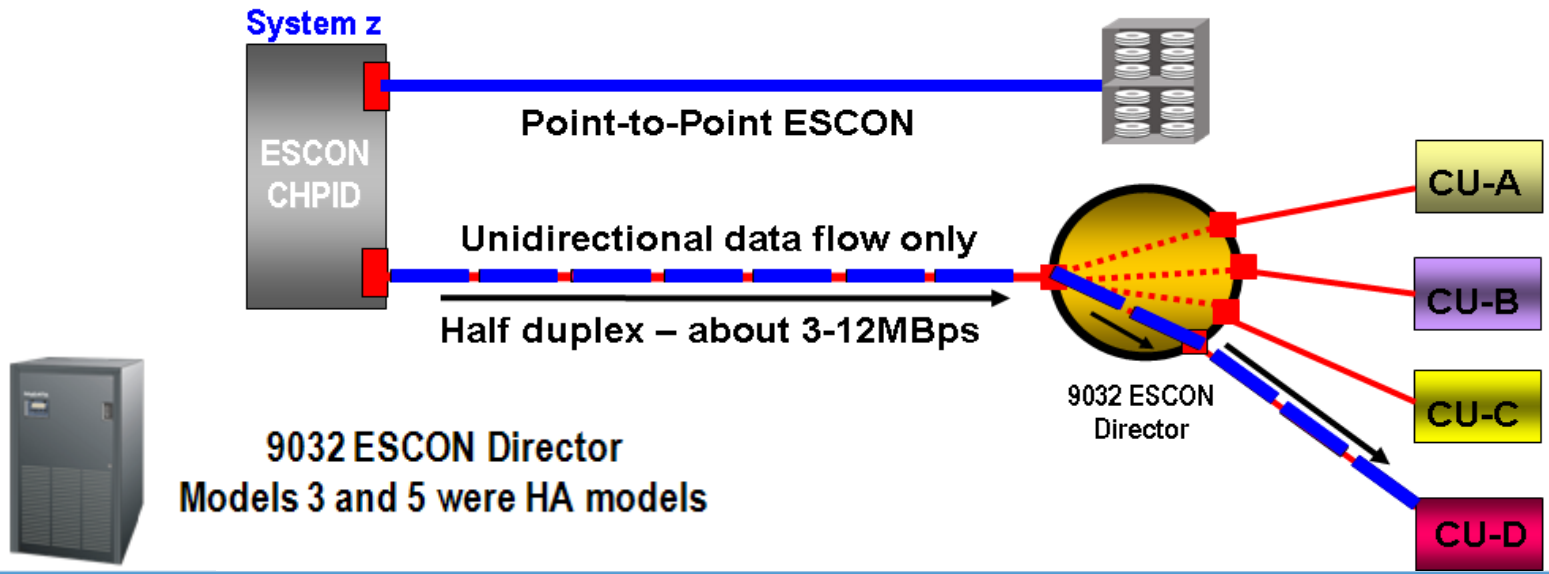

Mainframe Terminology

FICON Channel Paths - 2000

- This newer but very mature protocol provides a robust packet switched, bi-directional data transfer mechanism.
- System z10, z196, z114 and zEC12 can do up to 64 open exchanges unless zHPF is driving the I/O in which case ~600 OEs can be active.

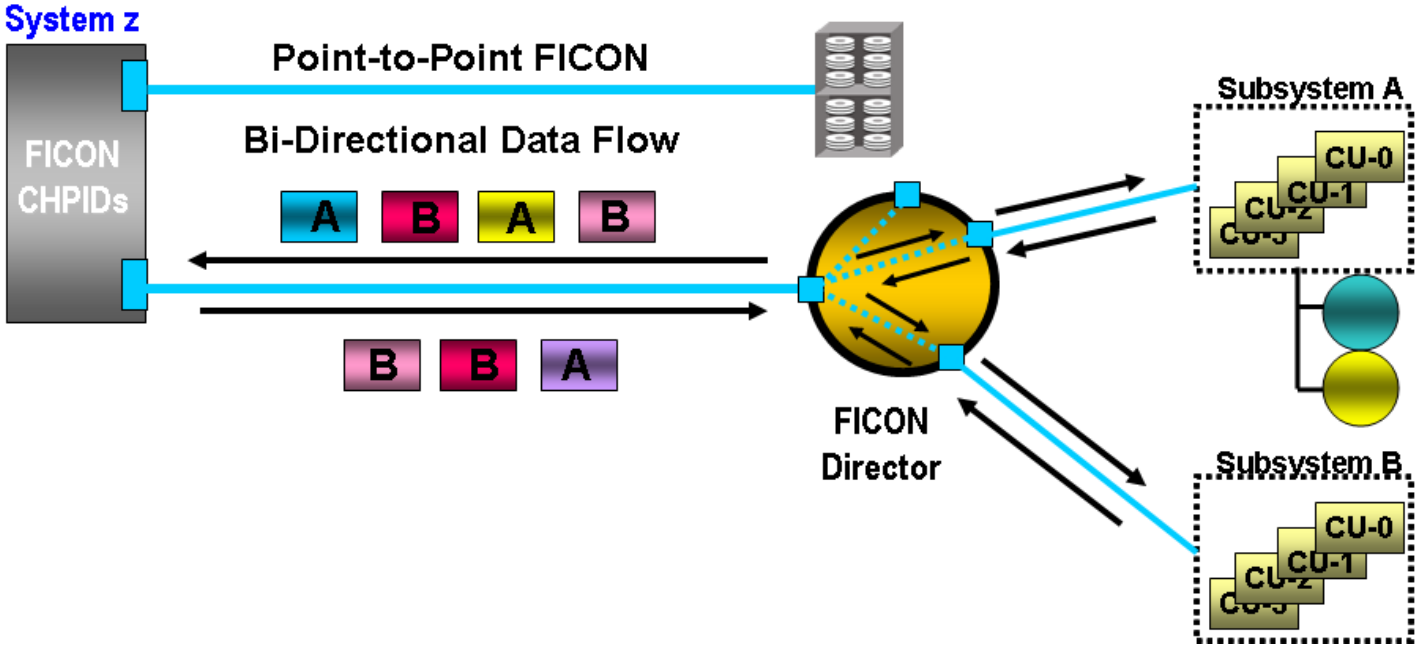

ESCON versus FICON I/O Protocol

- **ESCON among other things is:**

- Very old (September 1989)
- Proprietary protocol / now SBCON
- Half-Duplex at ~10-14 MBps
- Maximum of about 1,200 IOps
- Short distances of ~3-9 km
- Restricted number of control units per channel - 15
- Only 1,024 devices per channel allowed
- No channel consolidation going from bus-tag to ESCON
- Lots of multi-mode optical cables used for connectivity needed to be managed

- **FICON among other things is:**

- Most Current (December 2000)
- FC Standards-based protocol
- Full-Duplex at 130-1600 MBps
- 23,000 (CCW) or 92,000 (TCW) IOps
- Long distances of Local-to-25 km
- Better use of the 15 control units than ESCON
- 16,384 devices per channel are now allowed
- Channel consolidation from 2:1 up to 16:1 when going to FICON
- Fewer single-mode and/or multi-mode cables are needed makes management easier

So What Is FICON (FIber CONnection)?

- FICON has always been non-proprietary – based upon Fibre Channel Standards
- FICON was a major technical improvement over ESCON and continues to mature and get better and faster all the time
- Worldwide, there are still about ~20 to 30% of ESCON users that still need to completely migrate to FICON
  - The zEC12 mainframe does not host any ESCON CHPIDs
  - ESCON Directors will soon go to End of Support
- Many factors are pushing end users towards FICON:
  - Mainframes are getting faster and more flexible
  - DASD and Tape storage is getting faster and faster
  - Customer requirements are much wider and deeper
  - IBM support for ESCON is waning

Mainframe Terminology

1971 / 1974
S/390 Mainframe (M/F)

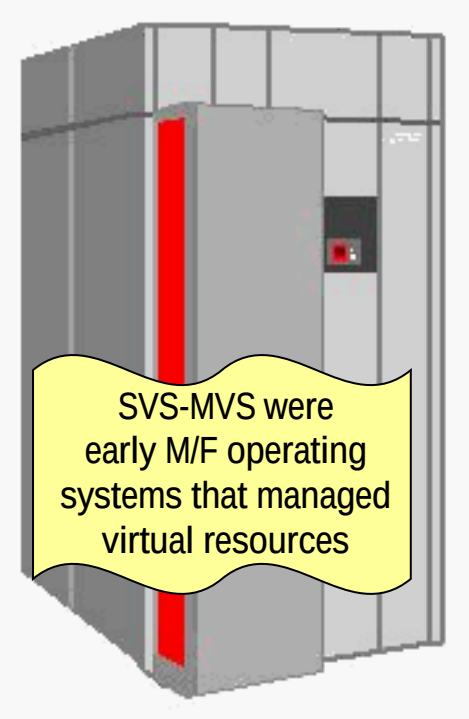

MVS could be run in up to 15 partitions on a M/F or could be intermixed with other operating systems running in other LPARs

Single Virtual Storage (SVS): Multiple Virtual Storage (MVS)

MVS was introduced in 1974, and it is often mentioned, even today, by old mainframer's as their primary operating system

MVS has become the de facto name for the mainframe operating system

This z/OS ancestor was a 24-bit, virtualized, batch processing-oriented operating system that managed lots of memory and DASD space for its time.

Mainframe Terminology

System z Mainframe (M/F)

z/OS can run in up to 60 partitions on a M/F or it can be intermixed with other operating systems running in other LPARs

Z Operating System (z/OS)

Was introduced in 2000 when the zSeries mainframes became available.

z/OS is a 64-bit server operating system, the latest IBM mainframe operating system, combining MVS and UNIX System Services (formerly known as MVS Open Edition, or OpenMVS).

Starting in 2010, on z196 and z114, z/OS can directly connect and manage an IBM BladeCenter with Power 7 and IBM System x blades

Mainframe Terminology

System z Mainframe (M/F)

TPF is a risk adverse system that demands the highest performance and availability levels!

Transaction Processing Facility (TPF)

This is a mature real-time operating system that processes many requests very quickly and runs on mainframes like *zSeries* and *System z.*

Any company that needs to process very high volumes of transactions (hotels, airlines, cruise lines, etc.) often utilizes this operating system which requires very high availability.

Mainframe Terminology

System z Mainframe (M/F)

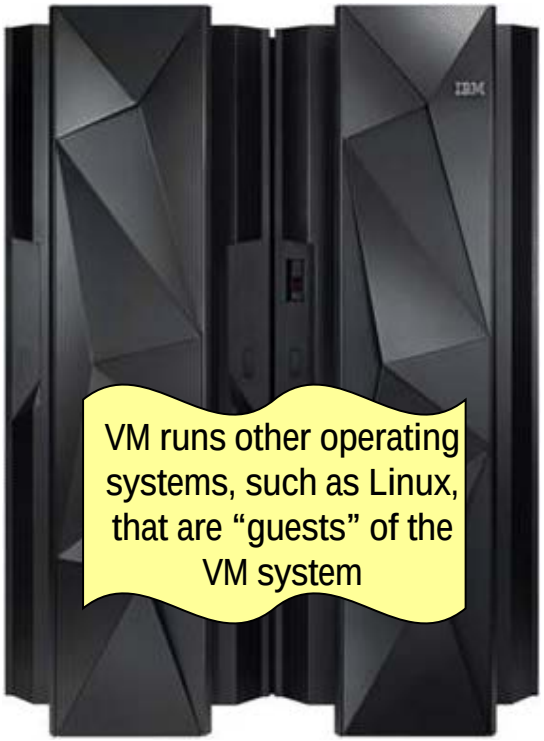

Linux on the System z is most often run as a “guest” under VM and VM actually does the I/O on behalf of Linux

Virtual Machine (VM)

Officially called VM/ESA (Enterprise Systems Architecture), it is an operating system for mainframes that can host other operating systems, including z/OS and Linux.

Each of the guest OS’s seems to have its own complete system of software and hardware resources (data storage, processor,…) but are actually sharing resources via VM services.

Mainframe Terminology

System z Mainframe (M/F)

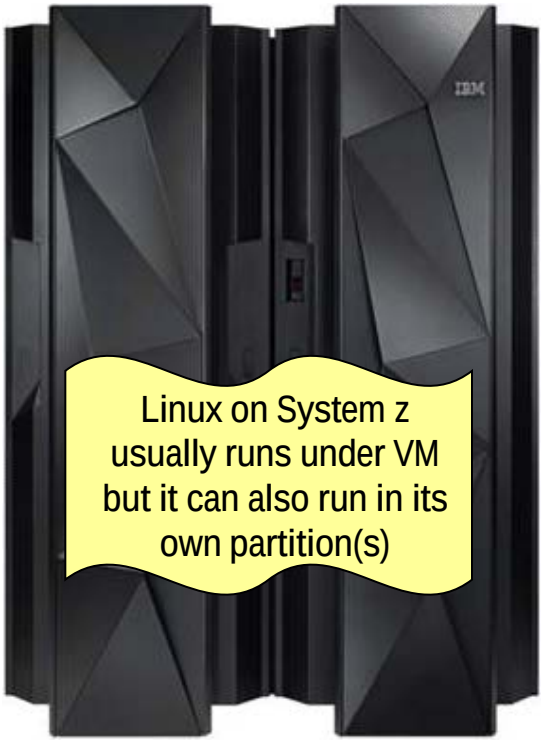

FICON channels in FCP mode for use with zLinux use the Queued Direct Input/Output (QDIO) I/O for communication with the operating system.

Linux On System z (zLinux)

Linux on System z allows a customer to leverage their highly available, reliable and scalable mainframe along with their highly available and powerful FCP and FICON infrastructure capabilities

A Linux administrator now simply administers Linux on a “Big Server”

Linux has been running on System z since 1999 – thirteen years ago!

Mainframe Terminology

Hardware Configuration Definition (HCD)

Define the mainframe's computing and I/O environment

Example FICON IOCP

```
ID

RESOURCE PARTITION=(LP1,1),(LP2,2),(LP3,3)
CHIPID PATH=(AD),SHARED, SWITCH=01,TYPE=FC
CHIPID PATH=(B2),SHARED, SWITCH=01,TYPE=FC
CHIPID PATH=(B6),SHARED, SWITCH=01,TYPE=FC
CHIPID PATH=(BA),SHARED, SWITCH=01,TYPE=FC

CNTLUNIT CUNUMBER=5000, PATH (AD,B2,B6,BA),
     UNITADD=((00,256)),LINK=(0114,0116,011A,011C),
     CUADD=0,UNIT=2105

     *

CNTLUNIT CUNUMBER=5700, PATH (AD,B2,B6,BA),
     UNITADD=((00,256)),LINK=(0104,0106,010A,010C),
     CUADD=7,UNIT=2105

IODEVICE
```

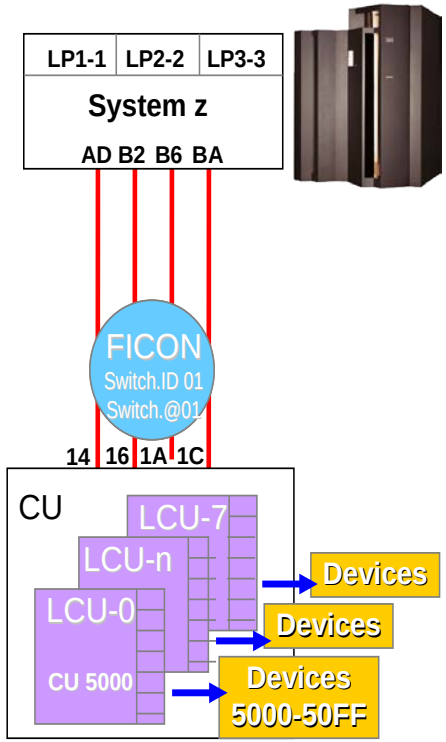

HCD is an element of z/OS that provides the interactive tool which is used to define the hardware configuration for both a processor's channel subsystem and the operating system running on the processor.

Hardware Configuration Definitions (HCD)

HCD provides an interactive interface that allows customers to define the hardware configuration for both a processor's channel subsystem and the OS running on the processor.

There really is not a comparable facility in the distributed world.

Mainframe Terminology

Sysplex and Parallel Sysplex

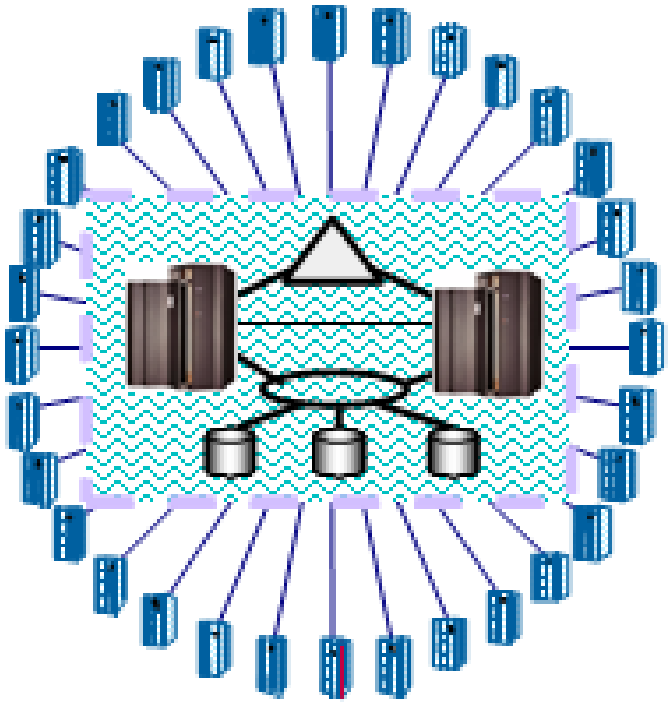

1 to 32 Mainframes

A Parallel Sysplex provides a clustered environment to provide both parallelism for application processing as well as better reliability/availability

Sysplex

IBM's SYStems ComPLEX was introduced in 1990 as a platform for the MVS/ESA operating system for IBM mainframe servers. The sysplex consists of the computer or multiple computers that make up the Computer Electronics Complex (*CEC*).

Parallel Sysplex is a Sysplex evolution providing a clustering architecture that has improved communication capabilities and supports more connected Central Processing Complexes (*CPCs*) and more copies of the operating system and is often deployed with a *Workload Manager capability.*

Parallel Sysplex Processing – mainframe clustering

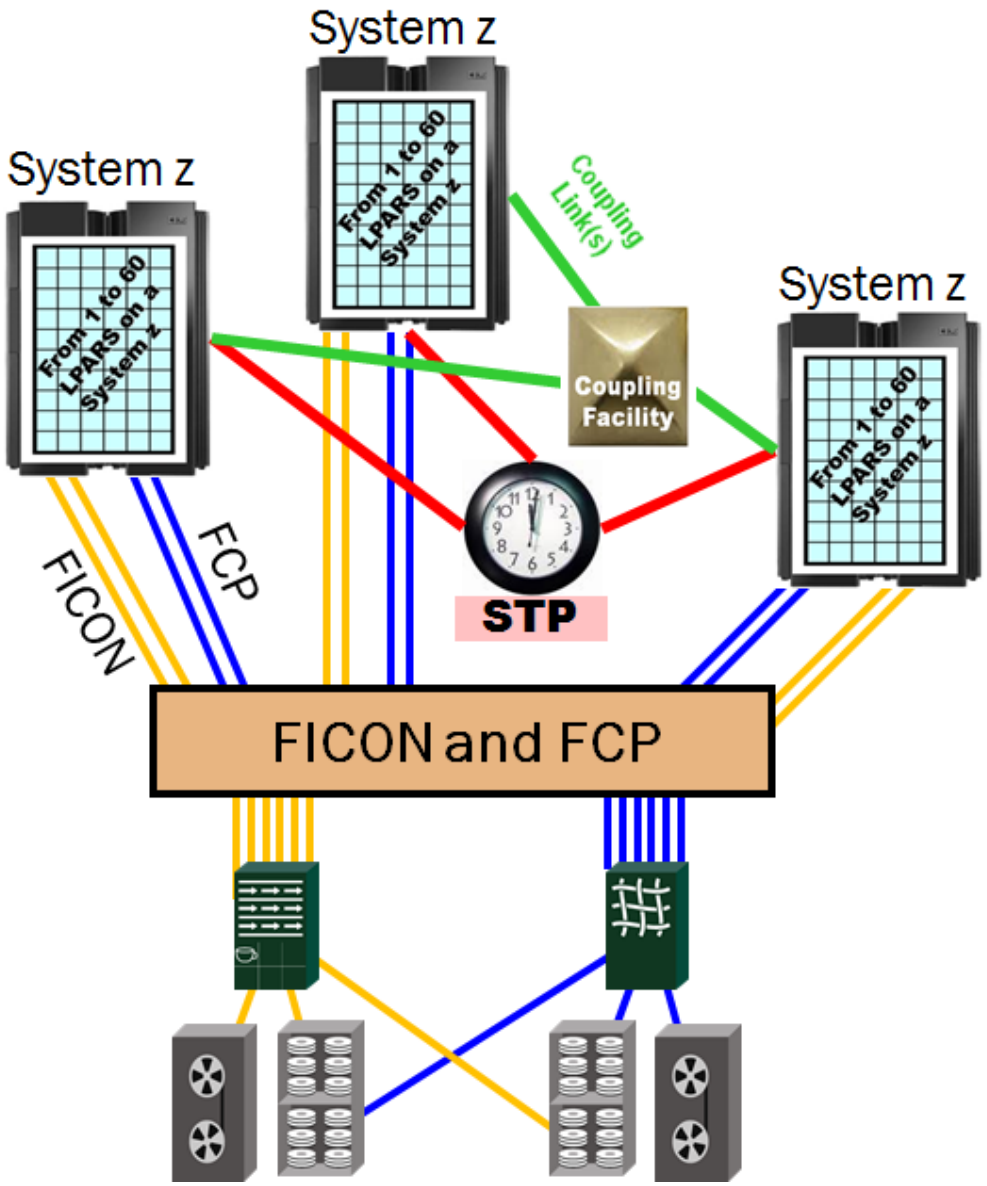

- A sysplex is an innovative multi-system data-sharing technology
- It is a collection of z/OS systems that cooperate, using certain hardware and software products, to process work
- Provides direct concurrent read/write access to shared data from all processing nodes
  - No loss of data integrity
  - No performance penalty
- Transactions and queries can be distributed for parallel execution based on available capacity and not restricted to a single node

Open Systems compared to Mainframe

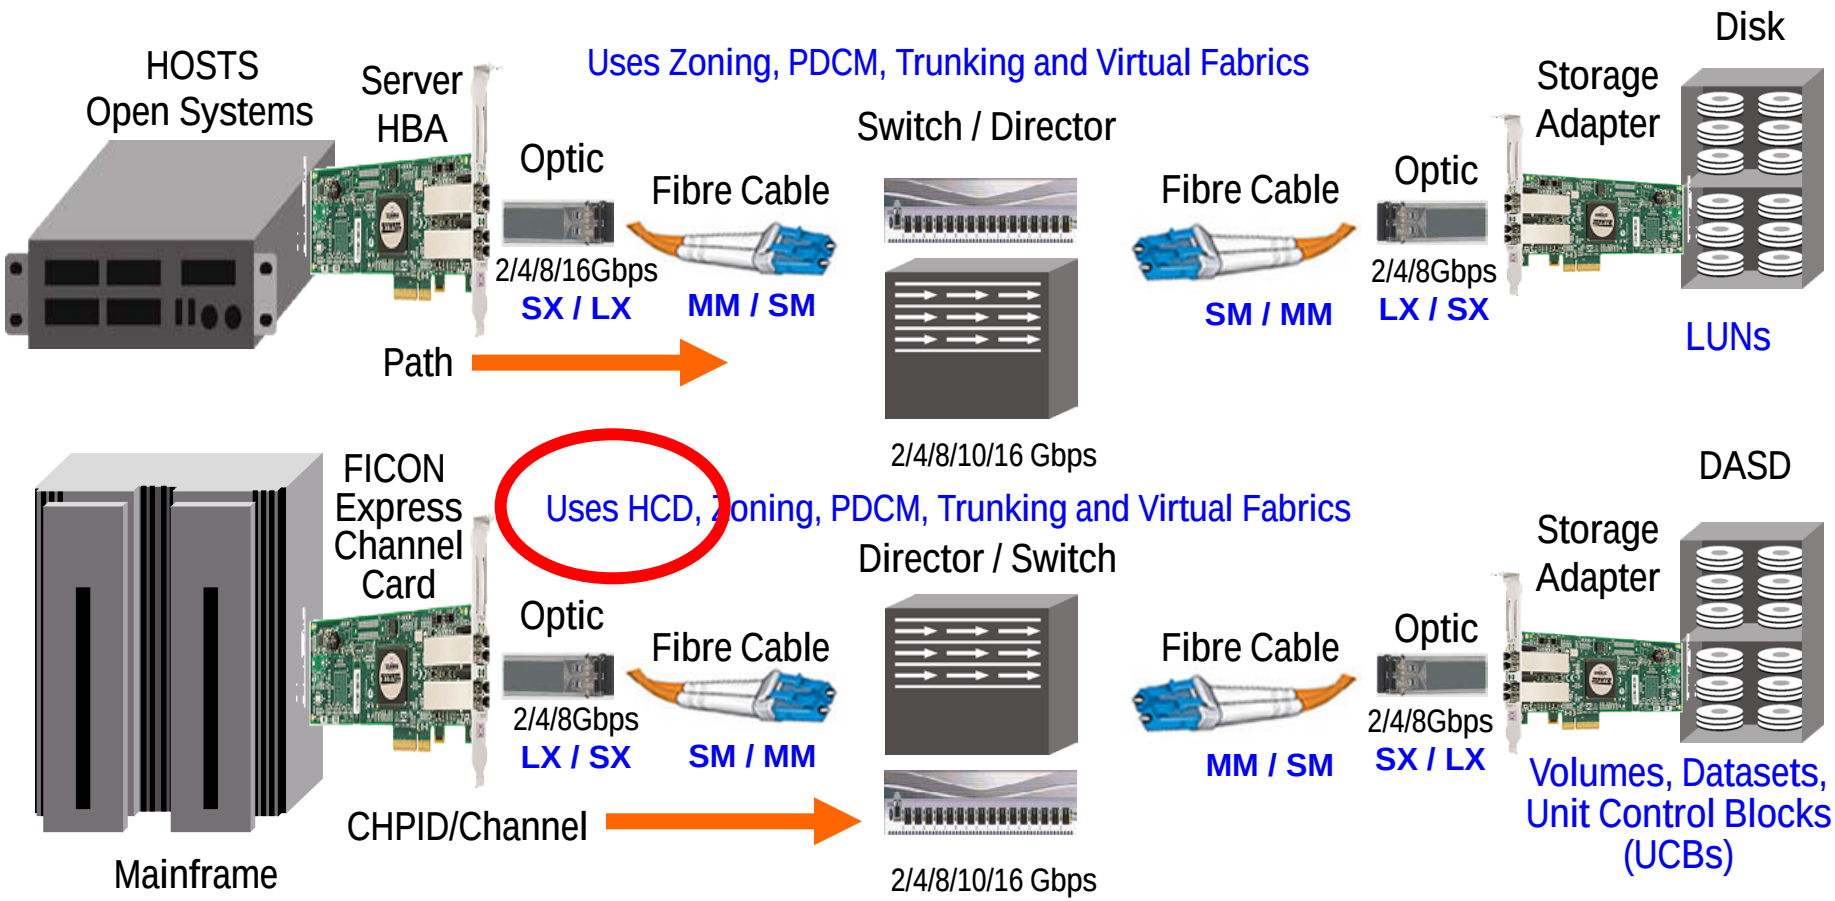

Multi-mode OM2
50 micron 500mHz
2/4 Gbps

Multi-mode OM3
50 micron 2000mHz
2/4/8/10/16Gbps

Multi-mode OM4
50 micron 4700mHz
2/4/8/10/16Gbps

Single-mode OS1
9 micron
2/4/8/10/16Gbps ~ 10k

Agenda for Session 13011

Session 13011

- History
- Terminology, Connectors, Cables, and Wavelengths
- **Addressing in FICON (MIKE)**
- ESCON Status, zHPF and NPIV
- Buffer Credits, CUP, RMF, BC/DR Solutions

FICON Traffic Routing – Single Switch

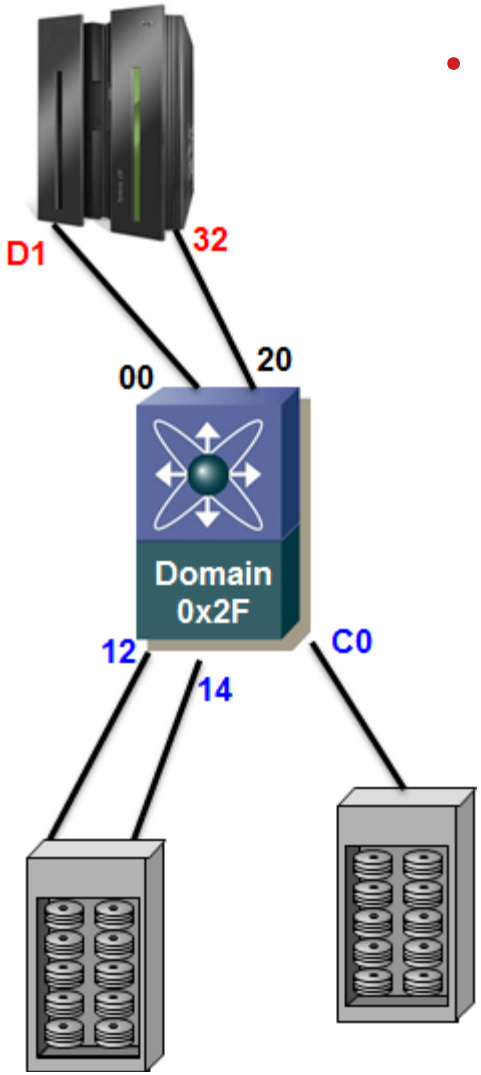

- Mainframe uses static FCID-based routing for devices
  - CHPIDs 32 and D1 are connected to a single switch
  - Switch has statically defined domain of 0x2F
  - Host IOCDS or HCD defines route to devices
  - Specifying the CHPID(s) and corresponding output port

```
CHPID PATH=(CSS(0),D1),SHARED,                                  *
      PARTITION=((LPARMVSY),(LPARMVSX,LPARMVSA,LPARMVSB,VMLPAR*
      02)),SWITCH=2F,PCHID=1B1,TYPE=FC
CHPID PATH=(CSS(0),32),SHARED,                                  *
      PARTITION=((LPARMVSY),(LPARMVSX,LPARMVSA,LPARMVSB,VMLPAR*
      02)),SWITCH=2F,PCHID=1B3,TYPE=FC
```

FICON Traffic Routing – Single Switch

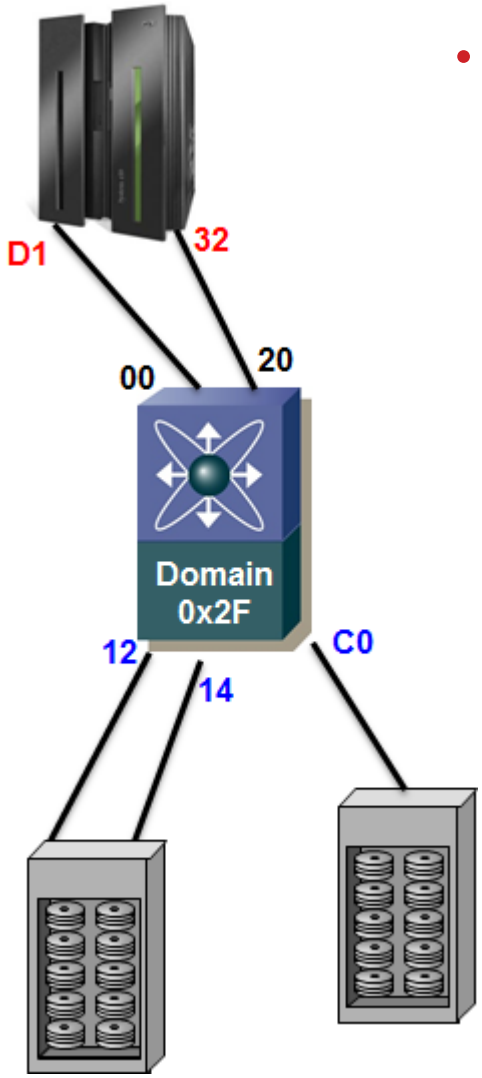

- Mainframe uses static FCID-based routing for devices
  - CHPIDs 32 and D1 are connected to a single switch
  - Switch has statically defined domain of 0x2F
  - Host IOCDS or HCD defines route to devices
  - Specifying the CHPID(s) and corresponding output port

```
CHPID PATH=(CSS(0),D1),SHARED,                                    *
      PARTITION=((LPARMVSY),(LPARMVSX,LPARMVSA,LPARMVSB,VMLPAR*
      02)),SWITCH=2F,PCHID=1B1,TYPE=FC
CHPID PATH=(CSS(0),32),SHARED,                                    *
      PARTITION=((LPARMVSY),(LPARMVSX,LPARMVSA,LPARMVSB,VMLPAR*
      02)),SWITCH=2F,PCHID=1B3,TYPE=FC


CNTLUNIT CUNUMBR=0CC0,PATH=((CSS(0),D1,32)),                      *
      LINK=((CSS(0),12,14)),                                      *
      UNIT=2105,CUADD=B,UNITADD=((00,032))

CNTLUNIT CUNUMBR=0C00,PATH=((CSS(0),D1)),                         *
      LINK=((CSS(0),C0)),                                         *
      UNIT=2105,CUADD=C,UNITADD=((00,032))
```

FICON Traffic Routing – Cascade

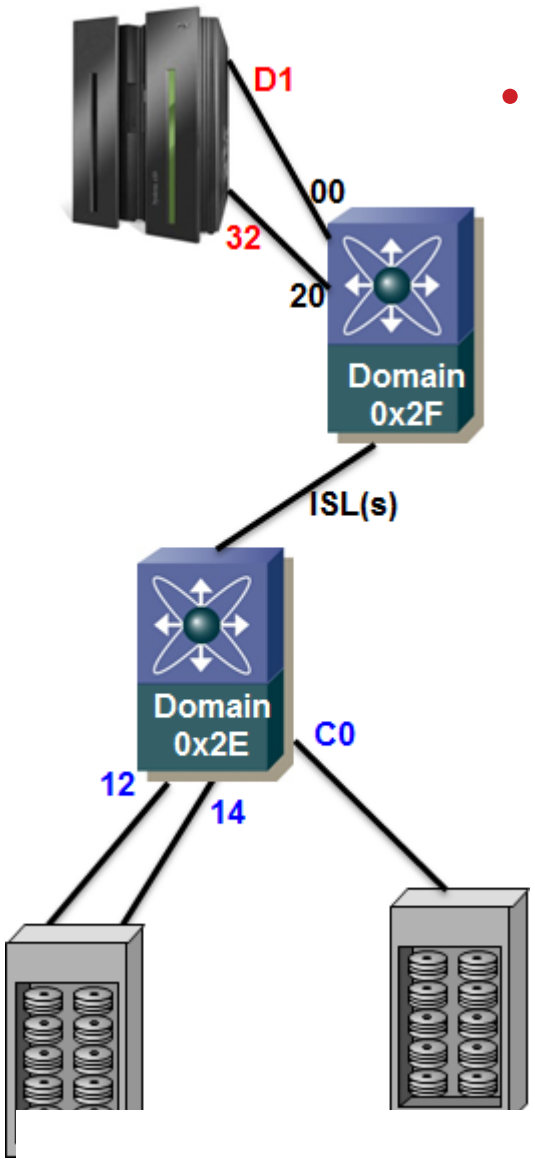

- Mainframe two switch routing
  - CHPIDs 32 and D1 connected to a host-side switch (0x2F)
  - DASD is attached to a different switch (0x2E)
  - Host IOCDS or HCD defines route to devices
    - Specifying the CHPID(s) and corresponding output domain/port
  - Mainframe is “blind” to the the ISL

```
CHPID PATH=(CSS(0),D1),SHARED,                                  *
      PARTITION=((LPARMVSY),(LPARMVSX,LPARMVSA,LPARMVSB,VMLPAR*
      02)),SWITCH=2F,PCHID=1B1,TYPE=FC
CHPID PATH=(CSS(0),32),SHARED,                                  *
      PARTITION=((LPARMVSY),(LPARMVSX,LPARMVSA,LPARMVSB,VMLPAR*
      02)),SWITCH=2F,PCHID=1B3,TYPE=FC

CNTLUNIT CUNUMBR=0CC0,PATH=((CSS(0),D1,32)),                    *
      LINK=((CSS(0),2E12,2E14)),                                *
      UNIT=2105,CUADD=B,UNITADD=((00,032))

CNTLUNIT CUNUMBR=0C00,PATH=((CSS(0),D1)),                       *
      LINK=((CSS(0),2EC0)),                                     *
      UNIT=2105,CUADD=C,UNITADD=((00,032))
```

Agenda for Session 13011

Session 13011

- History
- Terminology, Connectors, Cables, and Wavelengths
- Addressing in FICON
- **ESCON Status, zHPF and NPIV**
- Buffer Credits, CUP, RMF, BC/DR Solutions

What Is Happening With ESCON?

ESCON channels are being phased out!

- ***It is IBM's intent for ESCON channels to be phased out!***
- Only 240 ESCON channels are supported on z196 and z114
- And System z196 and z114 were the last mainframes to natively support ESCON channels
- System zEC12 does not allow attachment of ESCON CHPIDs

It is time to move to a FICON Infrastructure!!

…but…

What about the old ESCON devices that you still use?

FICON to ESCON Converter

Promotes **FICON infrastructure modernization**

Preserves **ESCON infrastructure investments**

Replaces **ESCON directors, FICON bridge, and ESCON extenders**

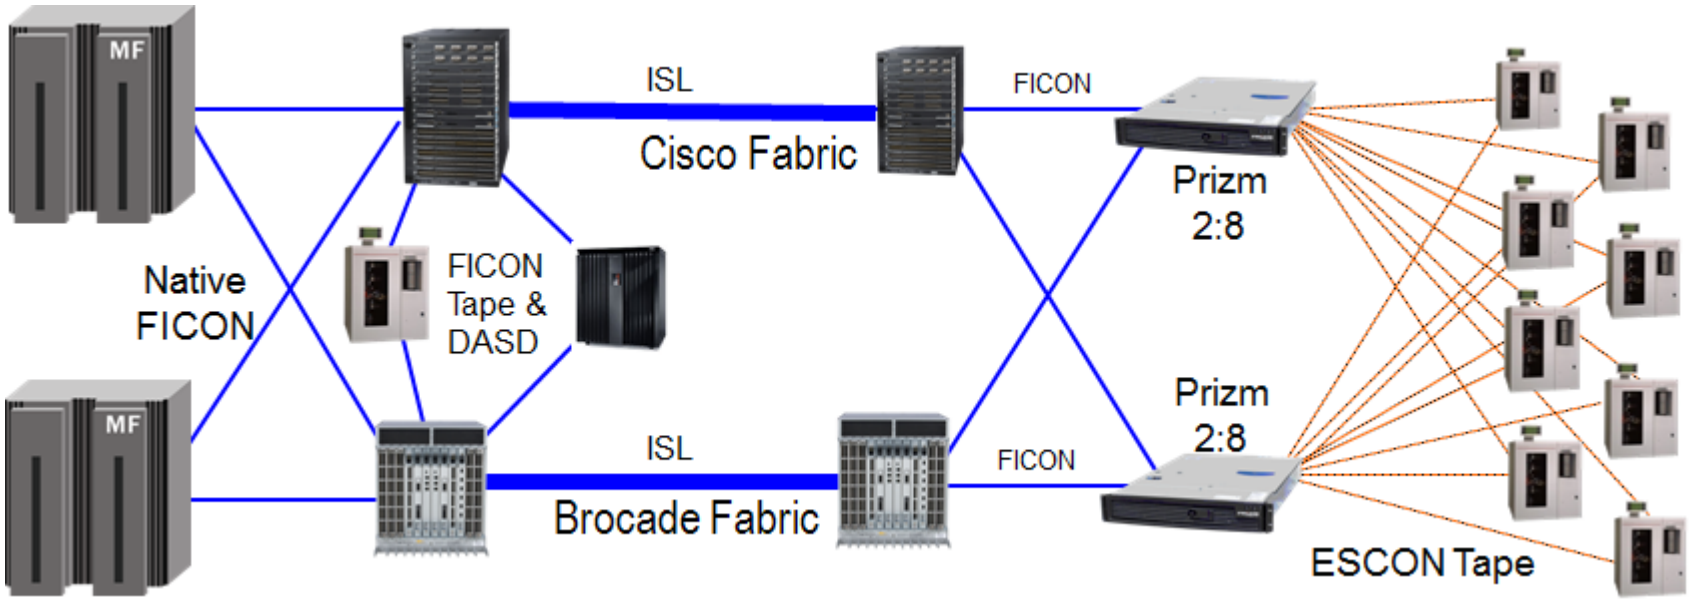

zHPF – High Performance FICON

It was first introduced in 2008

It continues to be enhanced

- Currently used for extended format data sets managed by Media Manager or by EXCP (DASD only – not tape)
- Encapsulates multiple FICON CCWs into a single frame
- z10, z196, z114 or zEC12 must be the host
- Supported by the major storage vendors
- Goal: Enhanced Performance

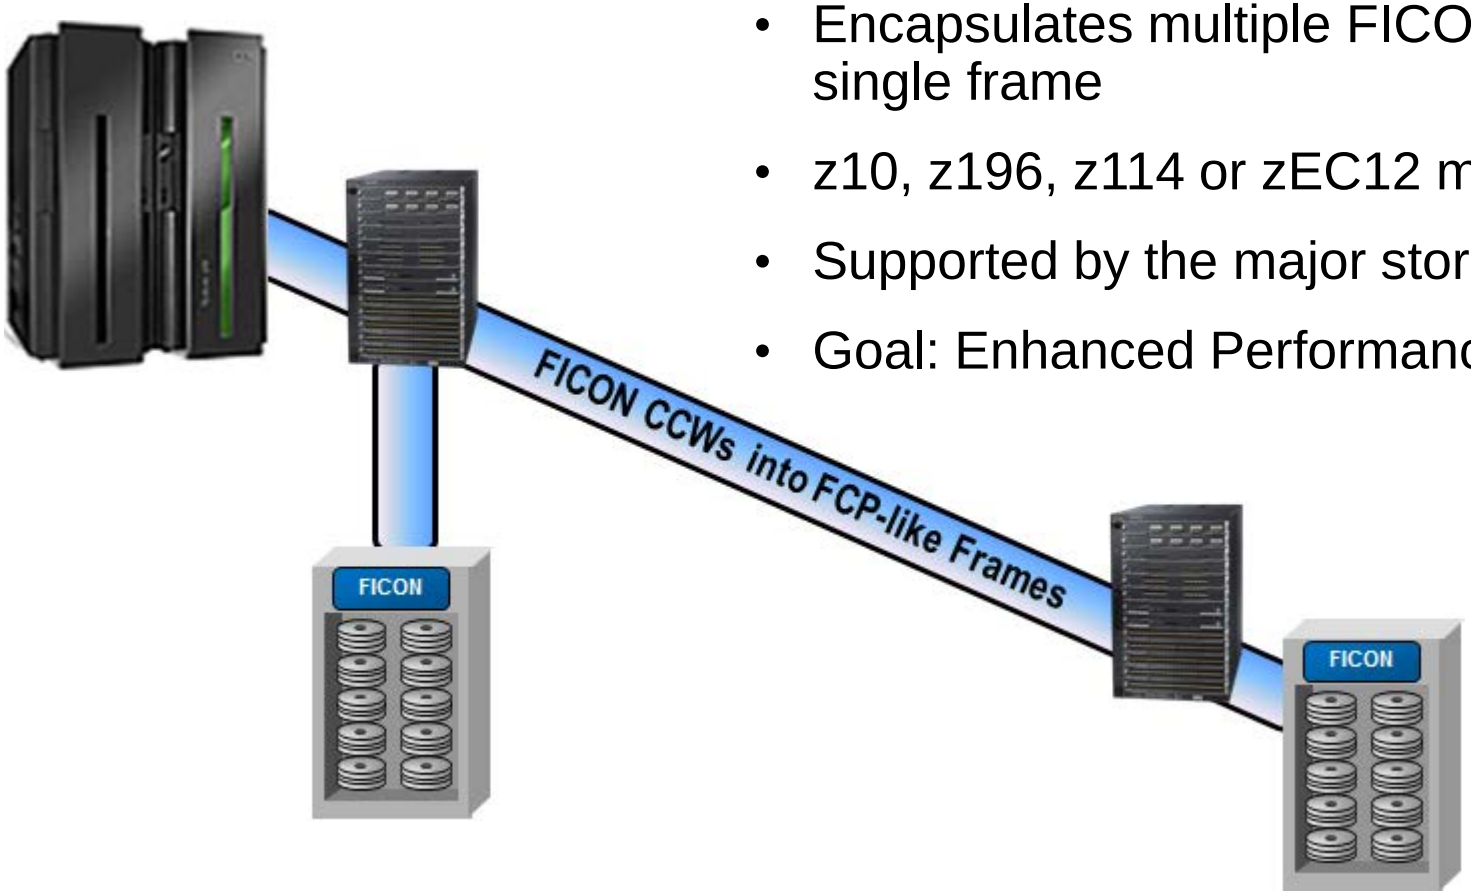

Node_Port ID Virtualization (NPIV)

- NPIV extends virtualization to the channel path just as virtualization has been extended into so many other areas of the mainframe.
- NPIV requires the use of switched-FICON fabrics that support NPIV standards.
- NPIV will allow the sharing of a single physical FCP channel among many operating system images. (CHPID N_Port connected to a switch F_Port)
- Once a physical path has been logged in, additional Fabric Discovery (FDISC) commands provide virtual addresses on that physical path which Linux can utilize to interleave I/O flow from many Linux guests onto a single physical path without creating serialization on that path
- NPIV allows for full support of LUN masking and zoning by virtualizing the Fibre Channel identifiers.
- IBM announced NPIV on Linux on System z in 2005 and today NPIV is supported on the System z9, z10, z196, z114 and zEC12.

Node_Port ID Virtualization (NPIV)

An example of System z when using NPIV

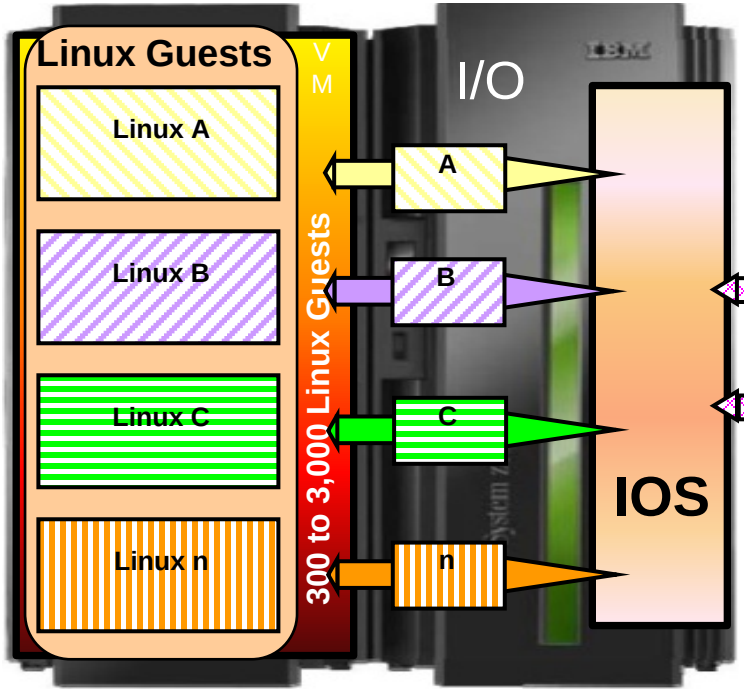

NPIV is a Standards-based Feature

NPIV works only when using switched-FICON

One FCP channel for many Linux guests

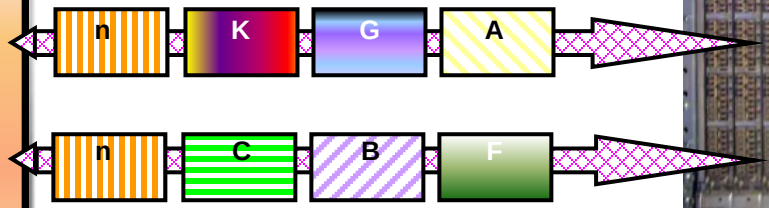

Lots of Parallelism

Fewer switch ports required!

FICON Switch NPIV enabled

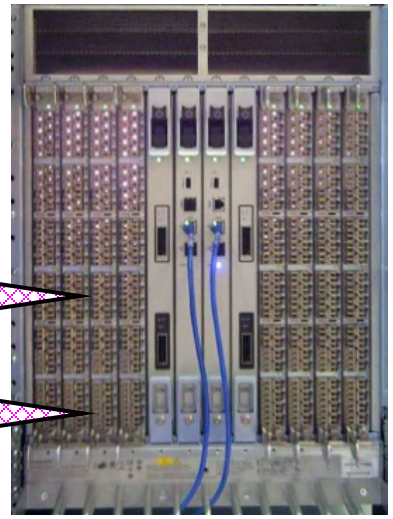

Much better I/O bandwidth utilization per path

8Gbps Is Great For NPIV!

NPIV – Node_Port ID Virtualization

- NPIV is standards based
- NPIV is used on the mainframe when Linux utilizes FCP Ports

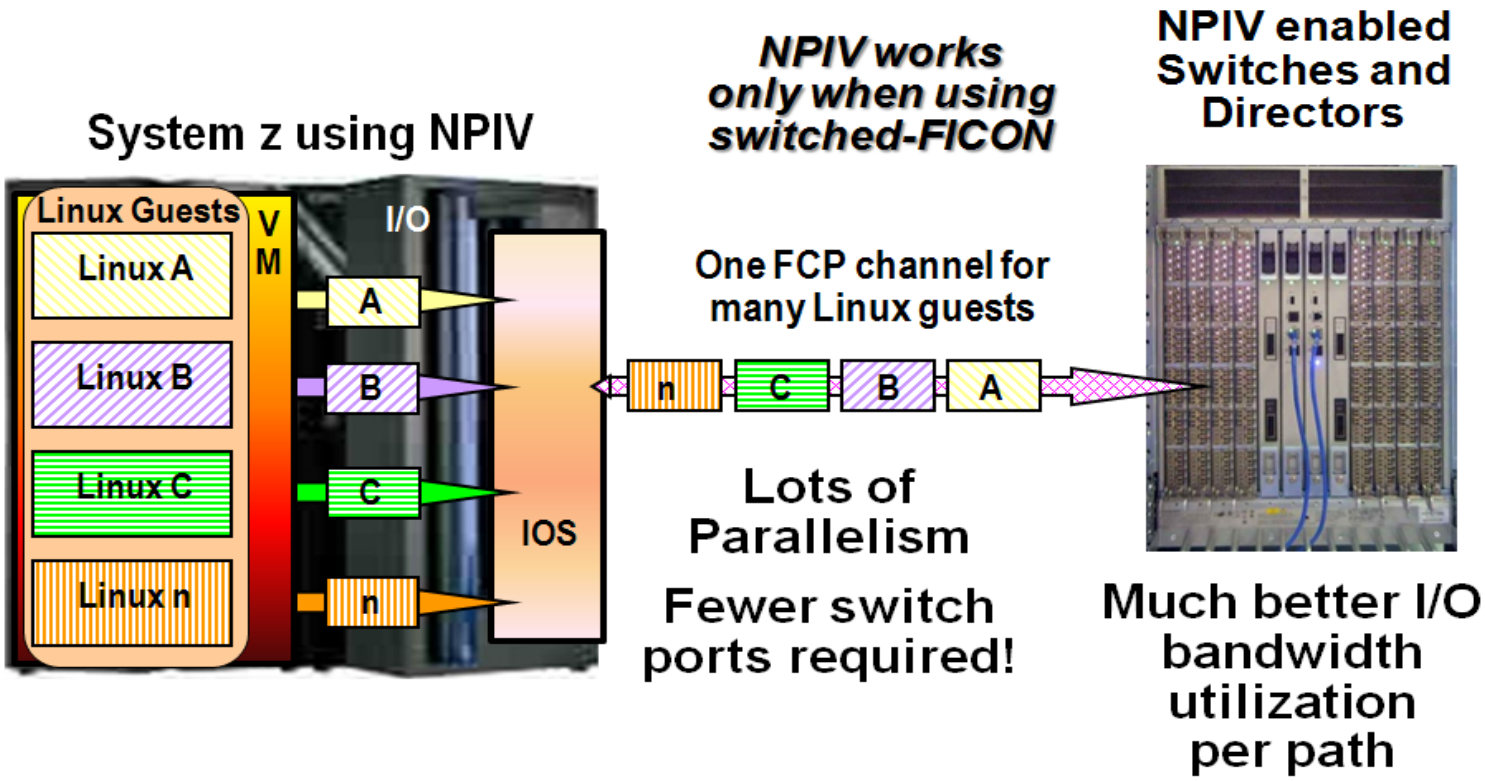

Agenda for Session 13011

Session 13011

- History
- Terminology, Connectors, Cables, and Wavelengths
- Addressing in FICON
- ESCON Status, zHPF and NPIV
- **Buffer Credits, CUP, RMF, BC/DR Solutions**

Mainframe
Number of Buffer Credits - Reality

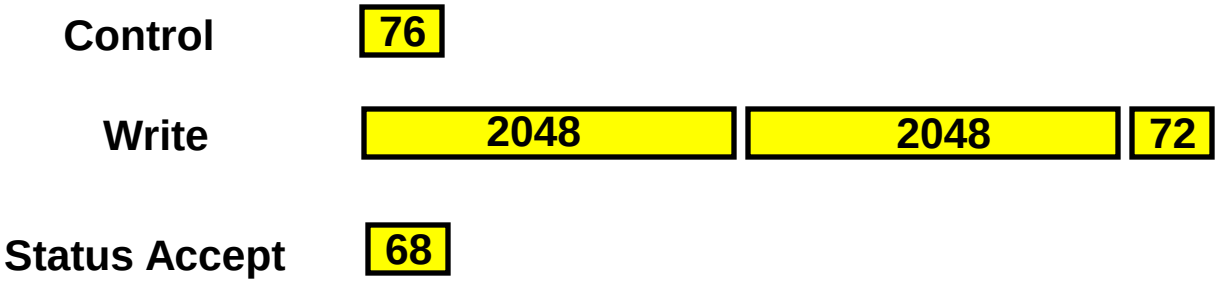

- Simple 4K write
- Will not fit into 2 buffers because of headers for FC as well as SB3

Average = (76+2048+2048+72+68) / 5 = 862 Bytes

Mainframe Terminology

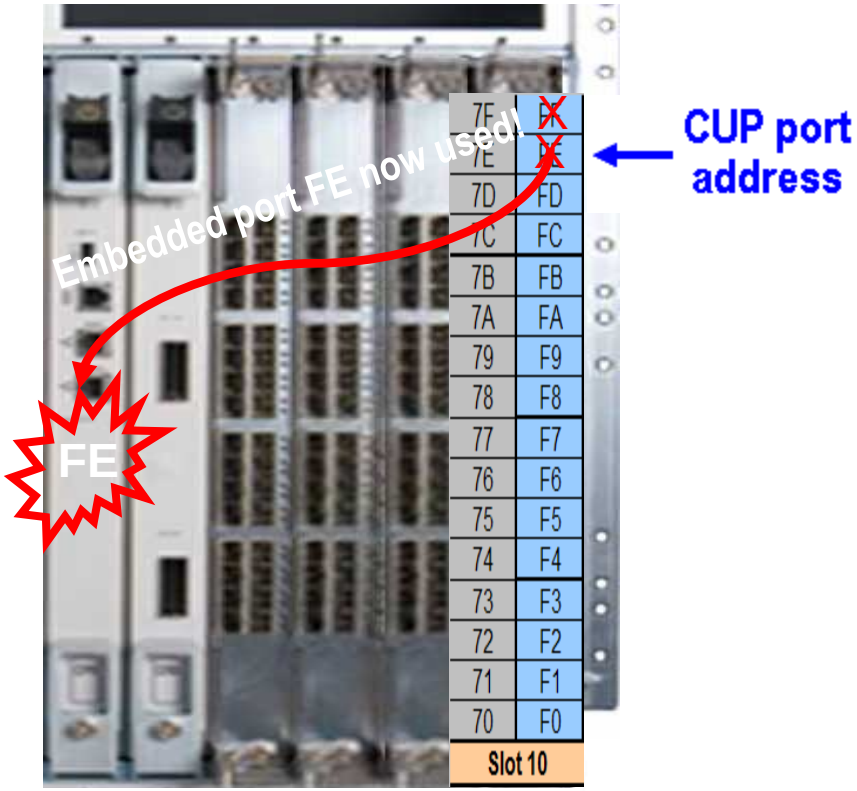

CUP should be deployed on every FICON Director! It allows you to provide the M/F environment with port performance information and buffer credit flow control information.
2 ports could be lost when CUP is used!

Control Unit Port (CUP)

- CUP is a licensed feature for Brocade and comes with the firmware for Cisco
- CUP is utilized for in-band management and FICON Director performance reporting up to the mainframe environment.
- Port address 0xFE is always the port address exclusively defined for the CUP – but it can also be a physical port which cannot be used for connectivity when CUP is enabled.

FICON Director Activity Report With Frame Delay

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

```
                                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 ABCD              START 04/12/2009-04.30.00  INTERVAL 000.15.00
                                         RPT VERSION V1R8 RMF        END   04/12/2009-04.45.00  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: 0000HIJKLMN
PORT     -CONNECTION-     AVG FRAME     AVG FRAME SIZE      PORT BANDWIDTH (MB/SEC)      ERROR
ADDR     UNIT       ID      PACING       READ   WRITE      -- READ --   -- WRITE --     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        998     789         50.32         10.56          0
 2B      CHP        70           0         69    2022          0.00          0.71          0
```

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 Buffer Credit Starvation

Fabric with zHPF Enabled

Data Replication Services for
Disaster Recovery and Business Continuance

Local/ Short Metro Distance Continuous Availability of Data within a Data Center	Continuous Availability within a Metropolitan Region	Protecting Data Centers at Extended Distance	Multi-tiered Solution for Protecting Data Centers
Single Data Center Applications remain active Continuous access to data in the event of a storage subsystem outage	Two Data Centers Systems remain active Multi-site workloads can withstand site and/or storage failures	Two Data Centers Rapid Systems Disaster Recovery with "seconds" of Data Loss Disaster recovery for out of region interruptions	Three Data Centers High availability for site disasters Disaster recovery for regional disasters
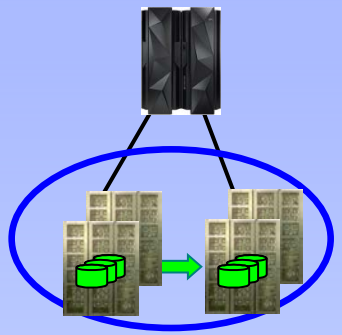	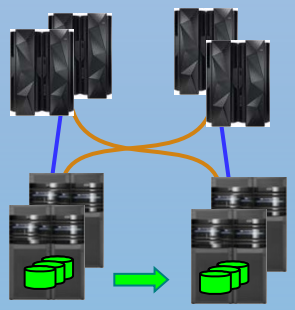	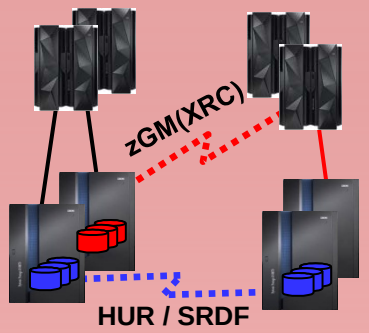	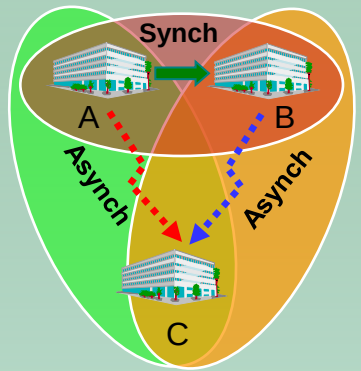
0 – 25 km	25 – 200 km	Local to 1,000s km	Automation of BC/DR

Provisioning for Distance

- The distance between two enterprise data centers could be within a campus, metro, or wide-area span. SAN extension solutions include:

- FC Attachments:
  - Direct attached (dark fiber) connections
  - Wavelength services such as Dense Wavelength Division Multiplexing (DWDM) and Course Wavelength Division Multiplexing (CWDM)
  - Synchronous Optical Networking (SONET) and Synchronous Digital Hierarchy (SDH services)
  - Fibre Channel over IP (FCIP)

- Direct fibre (up to 10km), DWDM (up to 300km with amplification), and CWDM (up to 100km) are typically meant for interconnecting SAN segments within metro or regional spans. SONET/SDH (thousands of kilometers) is meant for the national level and FCIP (ten of thousands of kilometers) is targeted toward global WAN ranges.

FICON Connectivity Between Data Centers

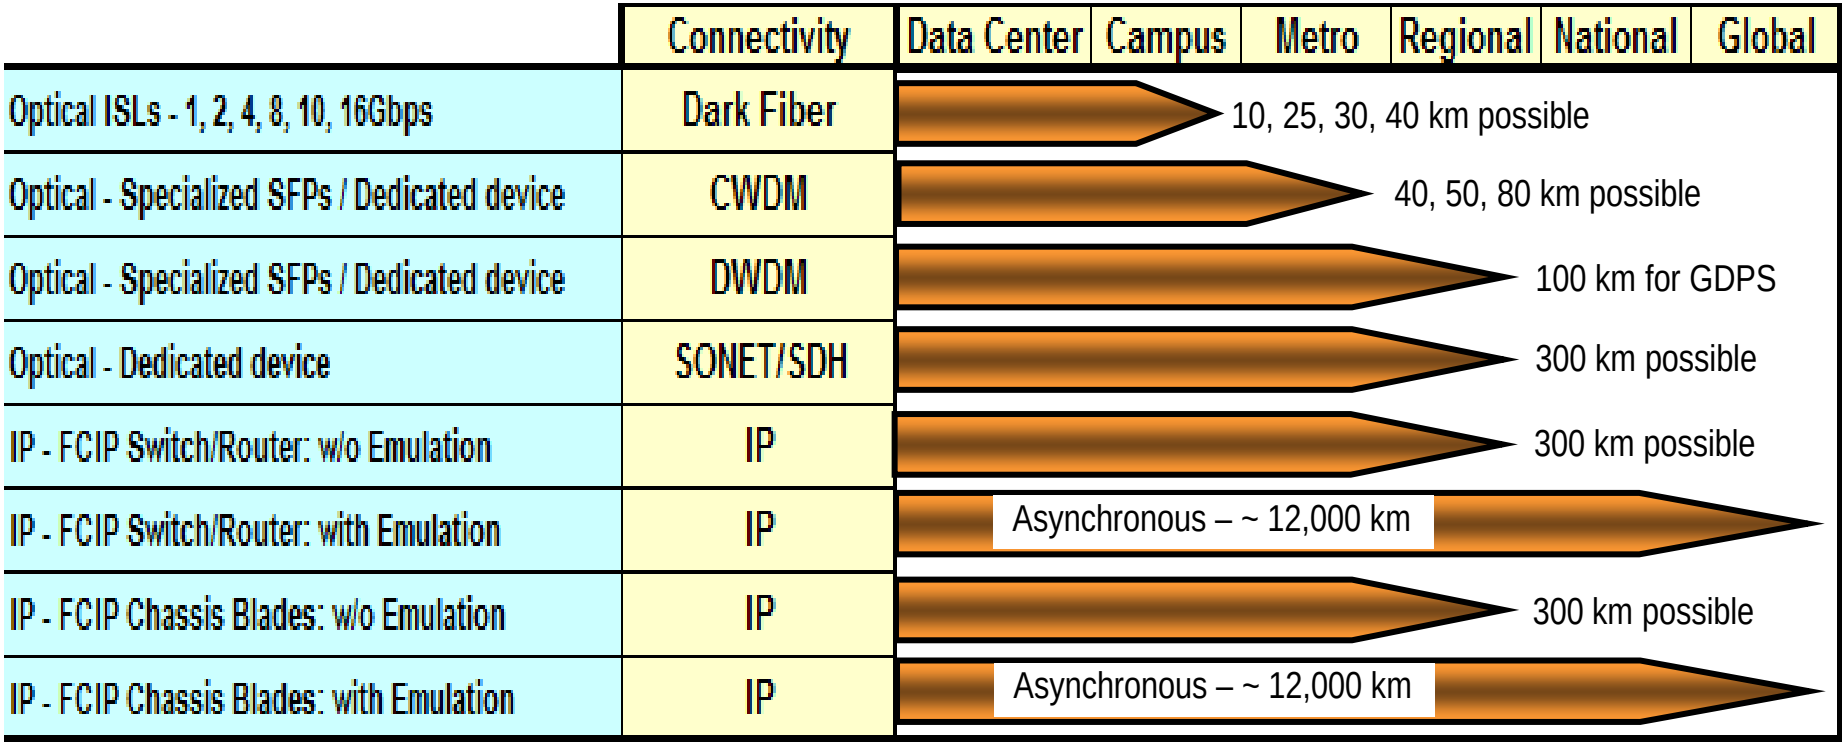

Maximum IBM Supported Distances

- Buffer Credits and link speed help determine real maximum distances for all distance extension functionality except IP
- DWDM with the RPQ (300km) requires special functionality
- Vendors provide emulation capabilities for link performance

This Is The End Of Part 2 and the end of the SAN 101 and 102 Training

We Hope You Have Been Informed and Entertained By This Presentation

SAN Sessions at SHARE this week

Tuesday:

Time-Session

1100 – 12166: What Every Mainframer Needs to Know About Networking

Wednesday:

Time-Session

0800 - 13062: FICON Channel Extension
0930 – 13013: Datacenter SAN & LAN Networking Convergence
1100 – 13117: Best Practices For SAN Management - For Both Open and FICON
1700 - 12734: Enhanced Availability and IT Resilience: An Integrated TS7700 Grid

Thursday:

Time-Session

0800 – 13010: A First Look at the Inner Workings and Hidden Mechanisms of FICON
0930 – 13009: A Deeper Look Into the Inner Workings and Hidden Mechanisms of FICON Performance
1300 – 13012: Buzz Fibrechannel - To 16G and Beyond

Mainframe/SAN Resources For You To Use

Visit Brocade's Mainframe Blog Page at:

http://community.brocade.com/community/brocadeblogs/mainframe

Visit Brocade's New Mainframe Communities Page at:

http://community.brocade.com/community/forums/products_and_solutions/mainframe_solutions

Visit Cisco's Storage Networking Page at:

http://www.cisco.com/en/US/products/hw/ps4159/index.html

#SHAREorg

SHARE
Technology · Connections · Results

Please Fill Out Your Evaluation Forms!!

This was session:

13011

Thank You For Attending Today!

5 = "Aw shucks. Thanks!"
4 = "Mighty kind of you!"
3 = "Glad you enjoyed this!"
2 = "A Few Good Nuggets!"
1 = "You Got a nice nap!"

And Please Indicate On Those Forms If There Are Other Presentations You Would Like To See In This Track At SHARE.

QR Code

SHARE in San Francisco 2013