Storage Area Networking (SAN) Overview

David Lytle, BCAF
Brocade Communications, Inc.

Mike Blair
Cisco Systems, Inc.

Monday February 28, 2011
Session Number 8485

Agenda for Session 8485

- History
- Connectors, Cables, and Wavelengths
- Mainframe Protocol Evolution
- Types and Components of Storage
- Let’s talk Fibre Channel
- FC Flow Control
- Fabric Routing / Virtual Fabrics / Partitioning
- Security / Zoning

Agenda for Session 8485

- **History - (Lytle)**
- **Connectors, Cables, and Wavelengths**
- **Mainframe Protocol Evolution**
- **Types and Components of Storage**
- **Let's talk Fibre Channel**
- **FC Flow Control**
- **Fabric Routing / Virtual Fabrics / Partitioning**
- **Security / Zoning**

Mainframe and Open Systems Time Lines

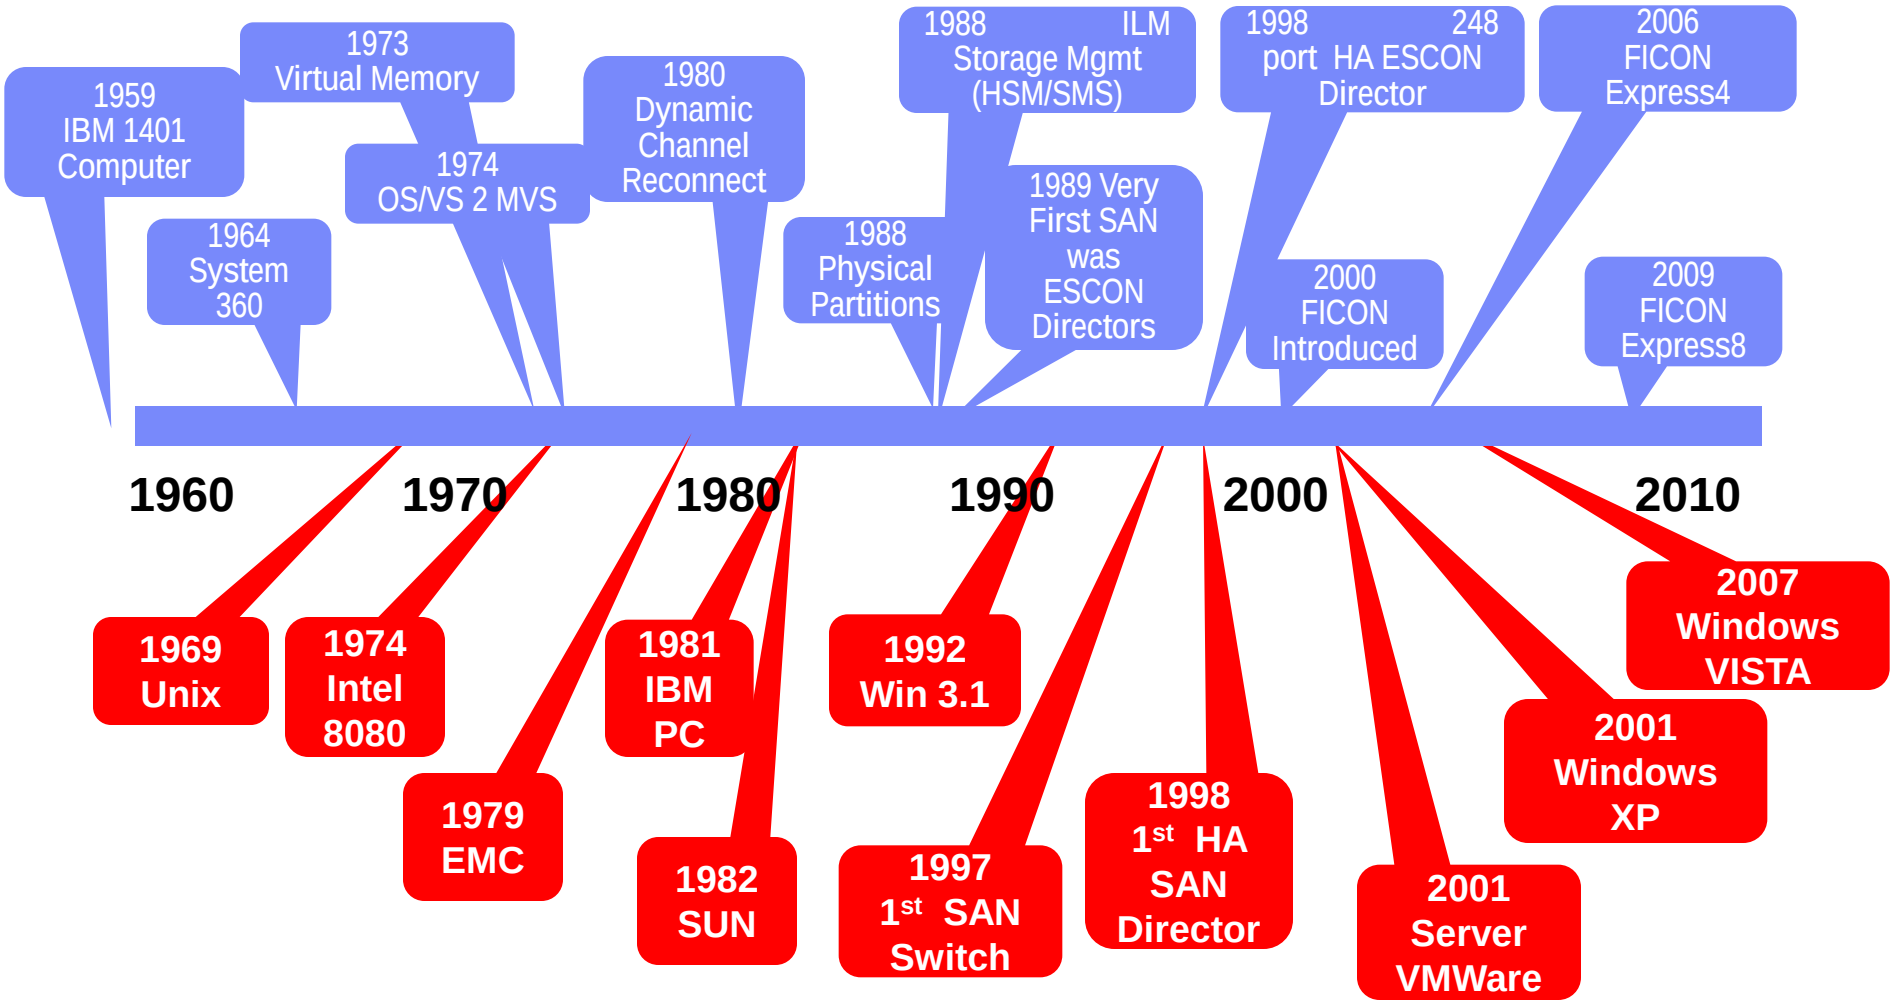

Material based on " I/O Architecture & Performance Analysis" class developed by IntelliMagic

Agenda for Session 8485

- History
- **Connectors, Cables, and Wavelengths - (Lytle)**
- Mainframe Protocol Evolution
- Types and Components of Storage
- Let’s talk Fibre Channel
- FC Flow Control
- Fabric Routing / Virtual Fabrics / Partitioning
- Security / Zoning

Fibre Channel Link

- A link consists of
  - 2 unidirectional “fibers” transmitting in opposite directions
  - May be either:
    - Optical fiber or Copper

- Transmitters may be:
  - Long wave laser
    - There can be multiple distances for these – ie. 4km/10km
  - Short wave laser
  - LED
  - Electrical

Components of a SAN Connection

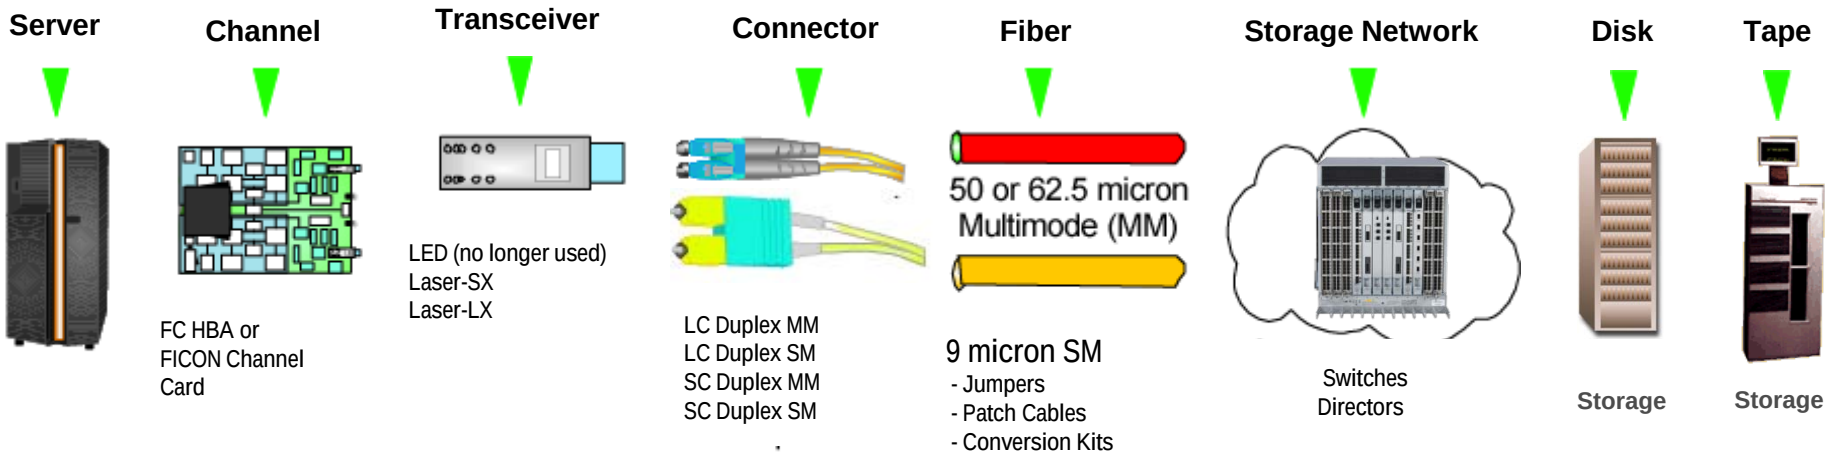

- **Many different combinations of components**
- **Connectors**
  - **SC is the very old style 1 Gbps connector**
  - **LC is the newer style 2/4/8 Gbps connector**
- **Fiber**
  - **Single-mode provides long distance connection (~10-80 km +)**
  - **Multi-mode provides local distance connection (~75-500 m)**

Open Systems compared to Mainframe

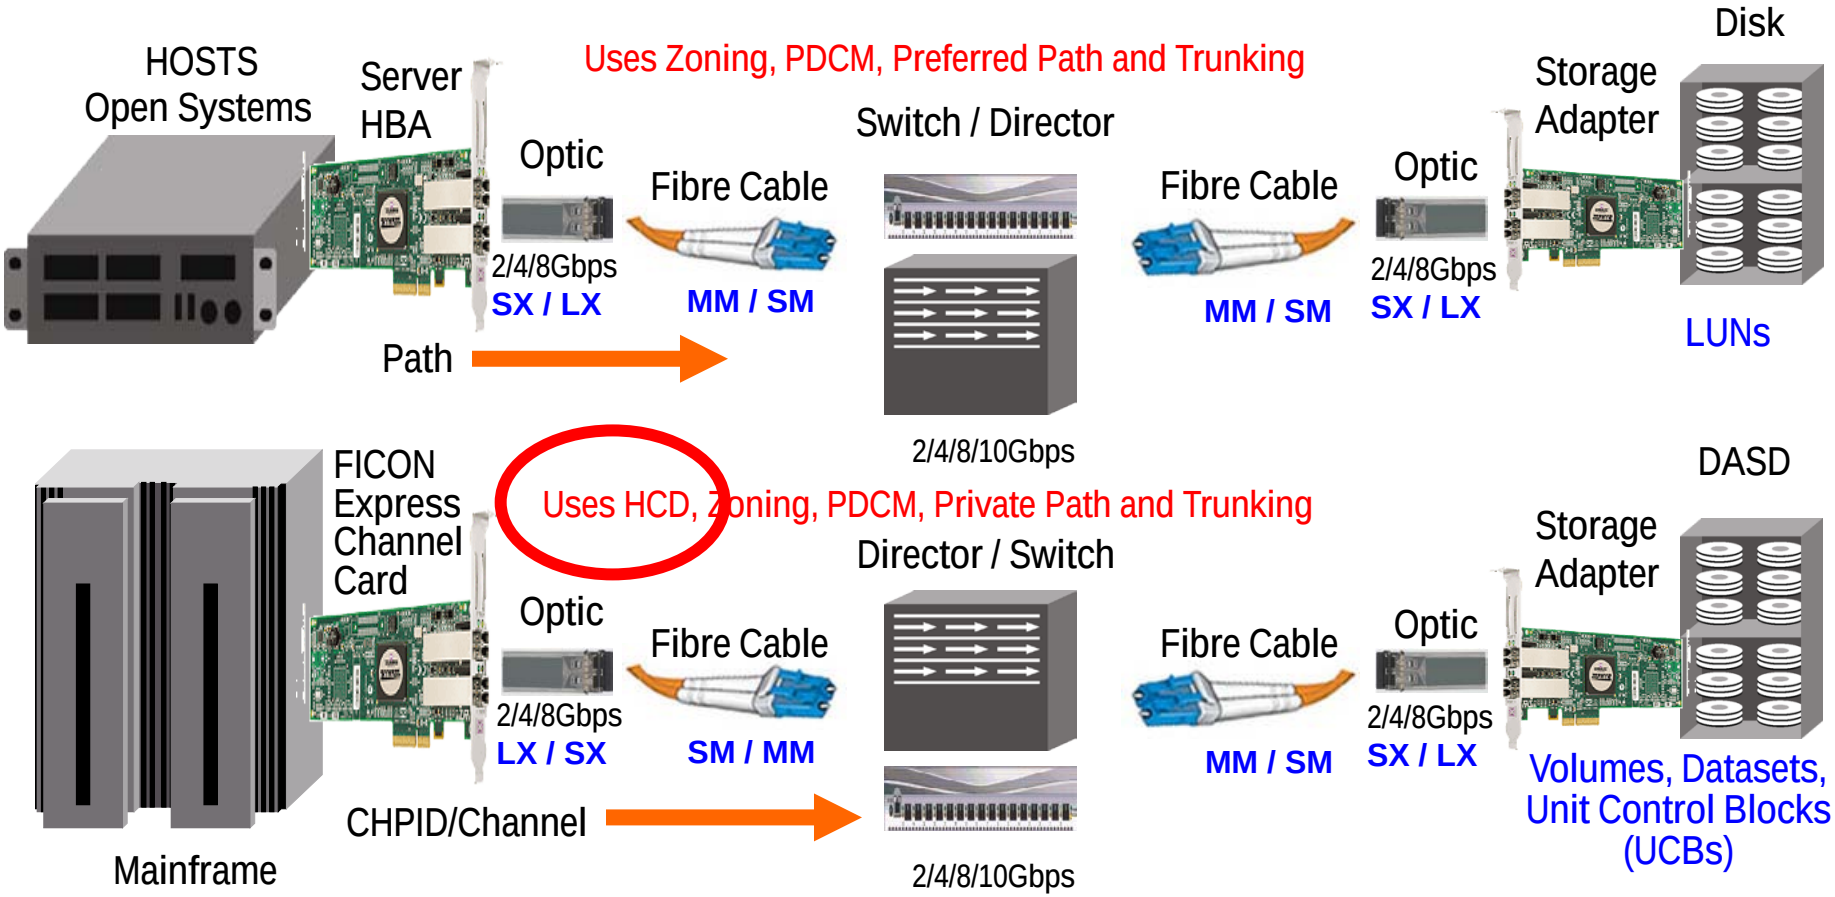

Multi-mode OM2
50 micron 500mHz
2/4 Gbps

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

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

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

Mainframe vs. Open System Terminology

Category	Mainframe Term	Open System Term
Switching	**FICON Fabric, Mainframe Infrastructure, Storage Network** **"SAN" not used in general mainframe usage**	**SAN**
Processor	**Host, CPU, zEnterprise (z196), System z, z10, z9, z990, z890, CPC (Central Processor Complex), CEC (Central Electronics Complex)**	**Server**
Operating System & Subsystems	**z/OS, MVS, zLinux, z/TPF, z/VM, Workload Manager (WLM)**	**Unix, Linux, Microsoft**
Host I/O	**Channel Adapter, CHIPID (Channel Path ID)**	**HBA, LUN, Initiator**
Device I/O	**UCB, Control Unit (CU), Device**	**Target**
Protocol	**FICON, ESCON, Bus & Tag, Channel Command Word (CCW), HPF**	**FCP, SCSI, iSCSI**
Storage	**DASD, VTS, VSM, Tape**	**Disk, Tape, VTL**

Agenda for Session 8485

- History
- Connectors, Cables, and Wavelengths
- **Mainframe Protocol Evolution - (Lytle)**
- Types and Components of Storage
- Let’s talk Fibre Channel
- FC Flow Control
- Fabric Routing / Virtual Fabrics / Partitioning
- Security / Zoning

ESCON, FICON, and HPF I/O Protocols

SHARE
Technology • Connections • Results

ESCON is:

- Very old (1989-present)
- Proprietary protocol
- Half-Duplex at ~10-14 MBps
- Maximum of about 1,200 I/O per second
- High handshakes for data transfer
- 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 is:

- Current (2000 - present)
- FC-SB3 Standard based
- Full-Duplex up to 600 MBps
- Maximum of about 30,000 I/O per second
- Less handshakes for data transfer
- Long distances up to 100km
- 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

HPF is:

- Very new (2009 – present)
- FC-SB4 Standard based
- Full-Duplex to ~800MBps
- Maximum of about 35,000 IO per second (more coming)
- Even fewer handshakes for data transfer
- Long distances up to 100km
- Only supported on System z10 and above
- Can run same channel as HPF to some CUs and FICON to other CUs
- Channel consolidation depends on how much HPF is used
- Really need to consider SM cables and long wave optics due to high speed

FICON (Fiber CONnection)

- **FICON connects IBM mainframe (System z and S/390) to storage subsystems**
  - Leverages fiber channel for mainframe cmds/data
  - Link rate increased to 1-2-4-8-10Gbps
  - Shortwave or Longwave (to 4km or 10km)
- **FICON "native mode"**
  - Supports communication with other FICON devices
- **Protocol intermix**
  - FICON and FCP on the same physical infrastructure
- **Switch cascading**
  - Two-switch cascading for DR/BC
- **Fabric binding**
  - Restricts ISLs to authorized switches
- **Control Unit Port (CUP) Management**
  - For in-band mgmt of FICON devices and RMF data

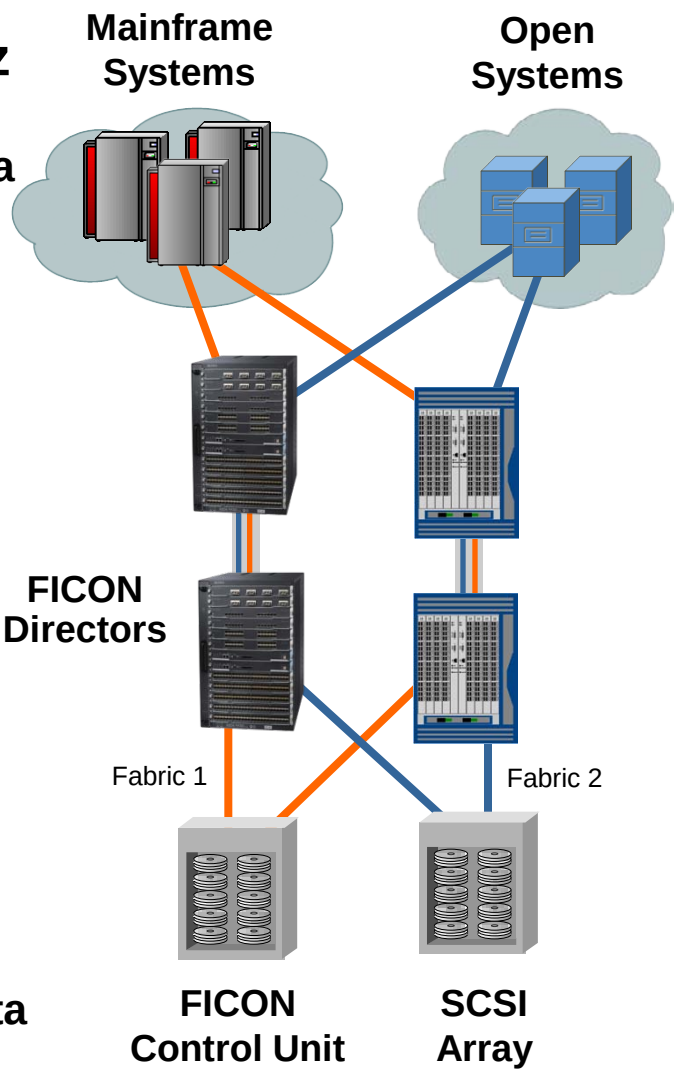

System z

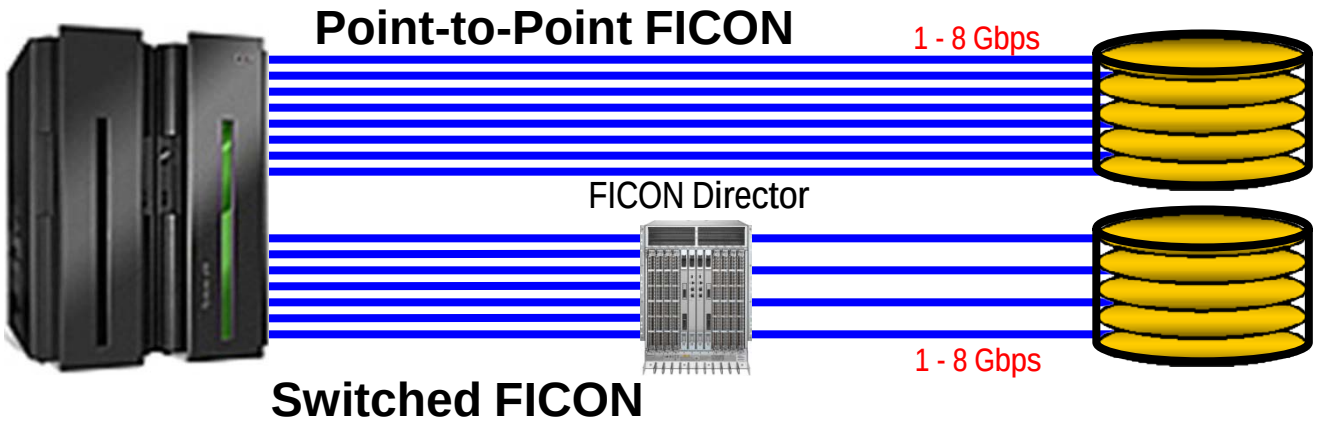

IBM supports 1 - 8 Gbps on the mainframe!

- **One of the main differences between a mainframe and other server platforms is the amount of data that can be exchanged between the mainframe and its external storage devices**
- **High I/O transfer rates can be achieved even while continuing to maintain a very high level of computing performance on all of the various workloads that are being executed**
- ***For the first time ever*, a mainframe user can deploy Fan In-Fan Out techniques to save on resources**

zHPF – High Performance FICON

z10/z196 (only)

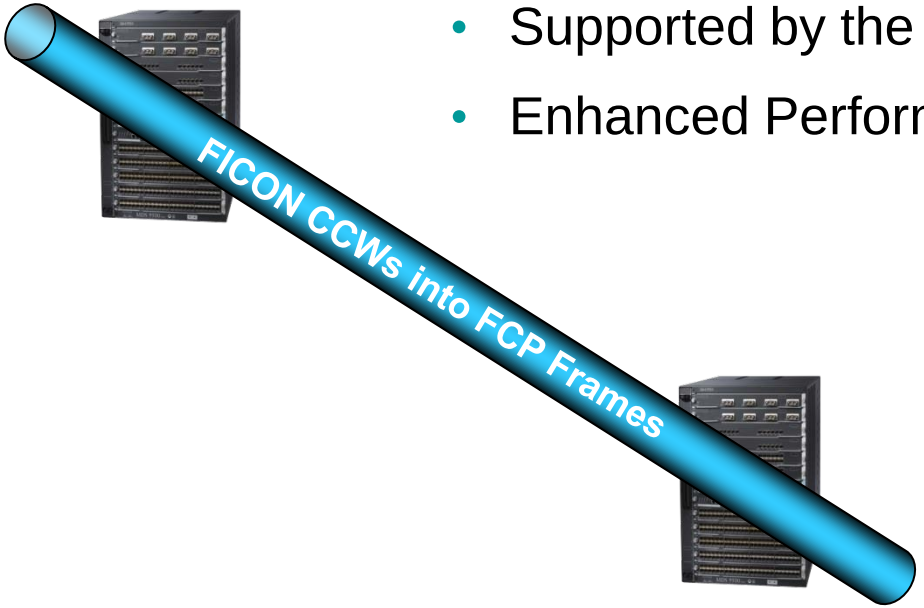

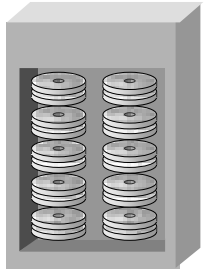

DASD

- Encapsulates FICON CCWs into FCP
- z10 or z196 required for host
- Supported by the major storage vendors
- Enhanced Performance

Node_Port ID Virtualization (NPIV)

- NPIV is standards-based (T11)
- Allows HBA port sharing between host partitions or virtual machines (VM)
- Separate fabric log-in by server partitions or VM enables application level
  - Zoning
  - Security
  - Traffic mgmt (e.g. QoS)
- Mainframe uses NPIV for Linux
- Brocade provides Access Gateway for blade server usage of NPIV in an open systems SAN

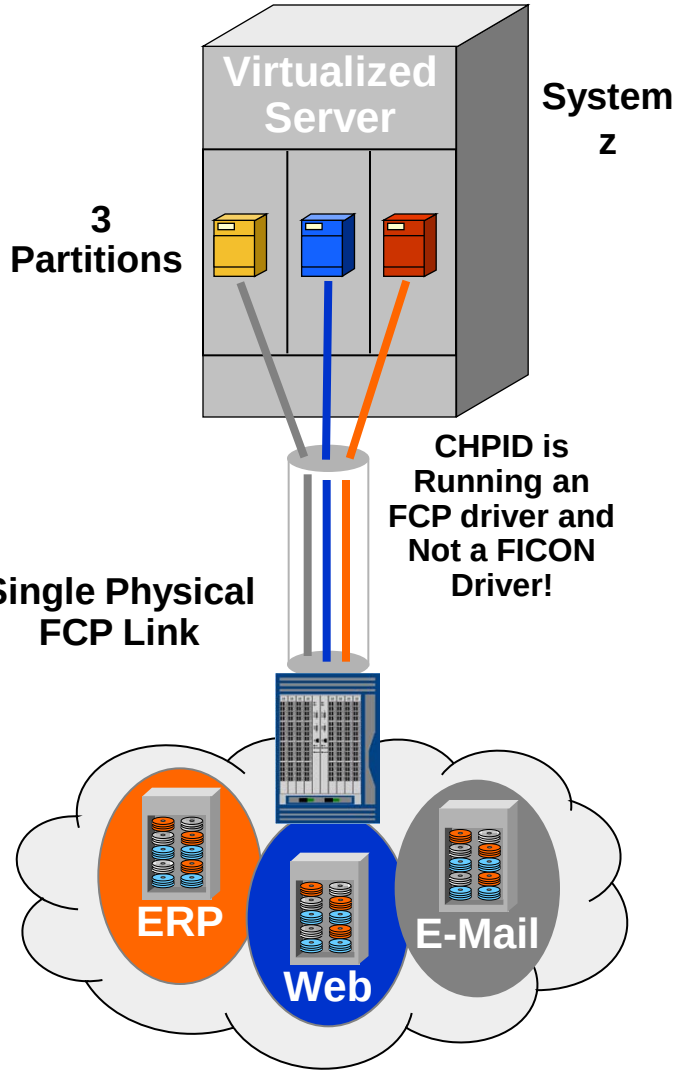

ESCON Statement of Direction

- **ESCON channels to be phased out**
  - It is IBM's intent for ESCON channels to be phased out. System z10 EC and System z10 BC will be the last server to support greater than 240 ESCON channels.
- Released April 28, 2009
- Currently, 1024 channels are supported on z10 EC and 480 channels are supported in the z10 BC
- Only 240 ESCON channels are supported on z196 …AND… z196 will be the last mainframe to natively support ESCON channels

FICON to ESCON Converter

Promotes FICON infrastructure modernization

Preserves ESCON infrastructure investments

Replaces ESCON directors, FICON bridge, and ESCON extenders

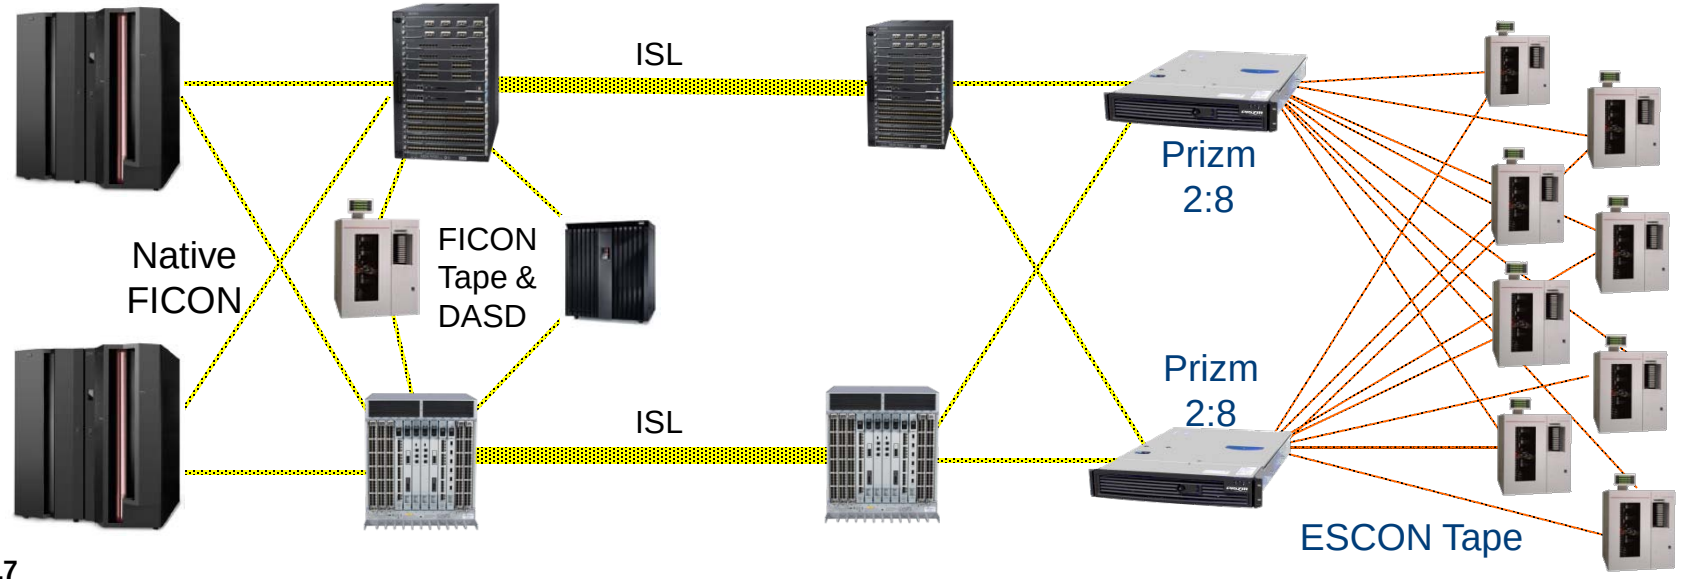

Agenda for Session 8485

- History
- Connectors, Cables, and Wavelengths
- Mainframe Protocol Evolution
- **Types and Components of Storage - (Blair)**
- Let’s talk Fibre Channel
- FC Flow Control
- Fabric Routing / Virtual Fabrics / Partitioning
- Security / Zoning

Storage Network Components

SHARE
Technology • Connections • Results

Multiple Infrastructure Vendors (eg. Brocade and Cisco)
Several components required to build a SAN

- Servers with Host Bus Adapters (HBAs)
- Mainframes with FICON/FICON Express Channels
- Storage systems
  - RAID (Redundant Array of Independent/Inexpensive Disks)
  - JBOD (Just A bunch of Disks)
  - Tape
  - VTS/VSM (Virtual Tape)
- Fibre Channel / FICON Switches or Directors
- Ethernet Switches (iSCSI)
- SAN management software












The Storage Industry
Storage Architectures

SHARE
Technology · Connections · Results

There are three types of storage architectures: DAS, NAS and SANs.

- **DAS – Direct Attached Storage is a server-centered storage architecture. Storage devices are linked directly to a server via a direct physical connection. The server is located between the storage devices and the LAN.**
- **NAS – A network-attached storage device is a server that is dedicated to file sharing, creating a shared connection to data and its storage using a file system. Its storage can be both internal to the server and external using DAS.**
- **SAN – A Storage Area Network is storage-network-centered. The storage devices and storage users are attached to a separate network. SCSI commands and data pass over FC.**
- **FICON – Same as SAN but mainframe channel commands and data pass over FC.**

Direct Attached Storage

- Direct Attached Storage (DAS)
- Storage is captive ‘behind’ the server, limited mobility
- Limited scalability due to limited devices
- No storage sharing possible
- Costly to scale
- Management can be complex
- Often cannot take full advantage of the technology

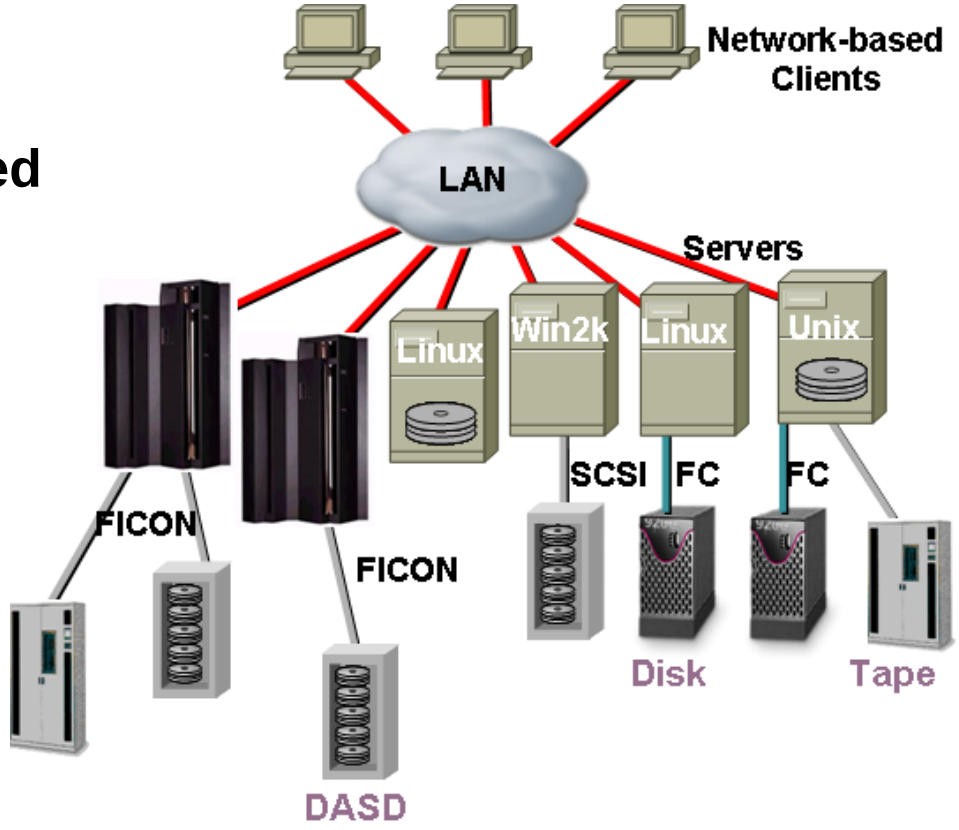

Network Attached Storage (NAS)

- Dedicated file server
- Optimized for file-based access to shared storage over an IP network
- Suitable for applications involving file serving/sharing
- High-performance access, data protection, and disaster recovery
- Capable of storage partitioning
- Uses network file system protocols such as NFS or CIFS

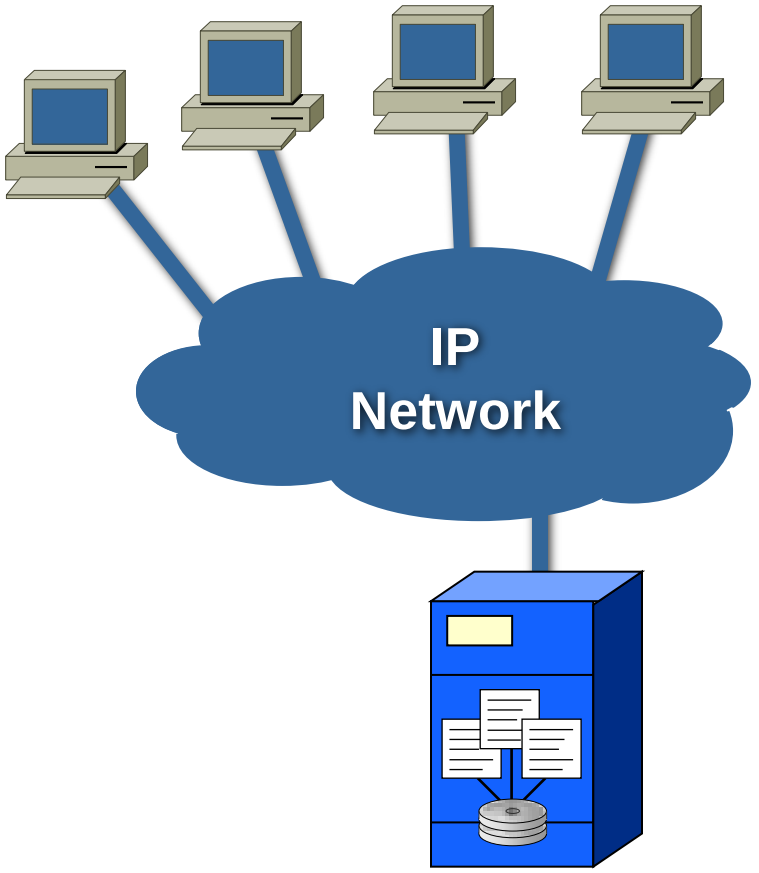

Storage Area Network (SAN)

Separation of Storage from the Server

- Storage is accessed Block-level via SCSI/FICON and can be in a switched environment
- High performance interconnect providing high I/O throughput
- Lower TCO relative to direct attached storage, storage can be shared
- Have to consider Vendor Interoperability / Qualifications
- More Complex management due to size/scale

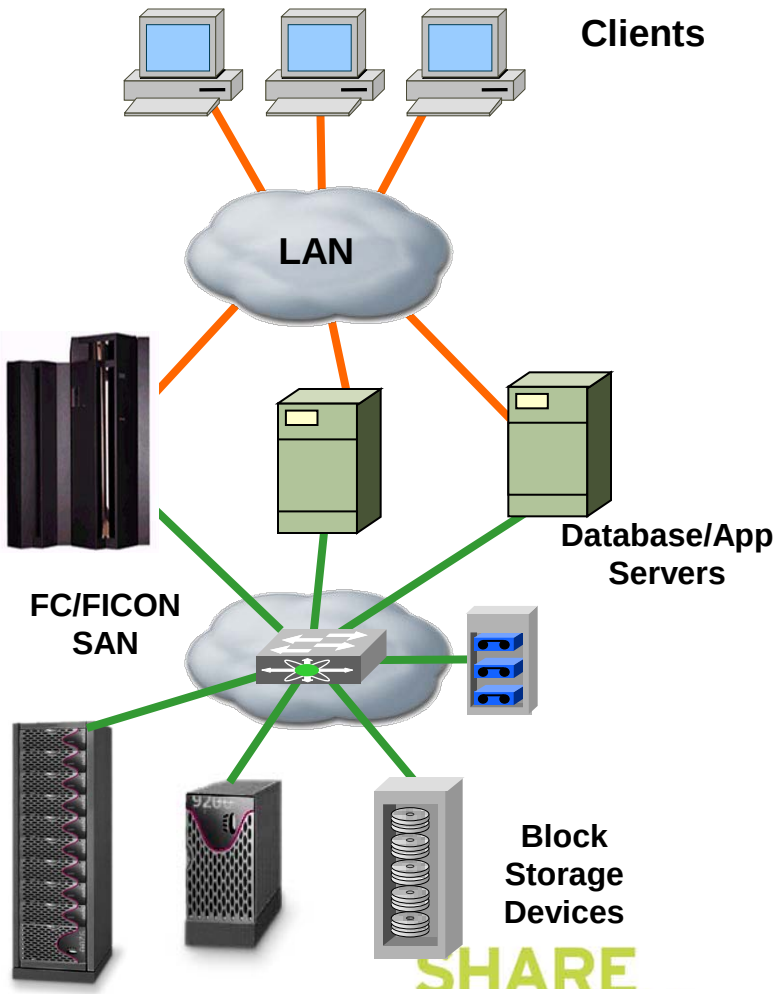

FICON or Storage Area Network (SAN)

Why Network Storage….?

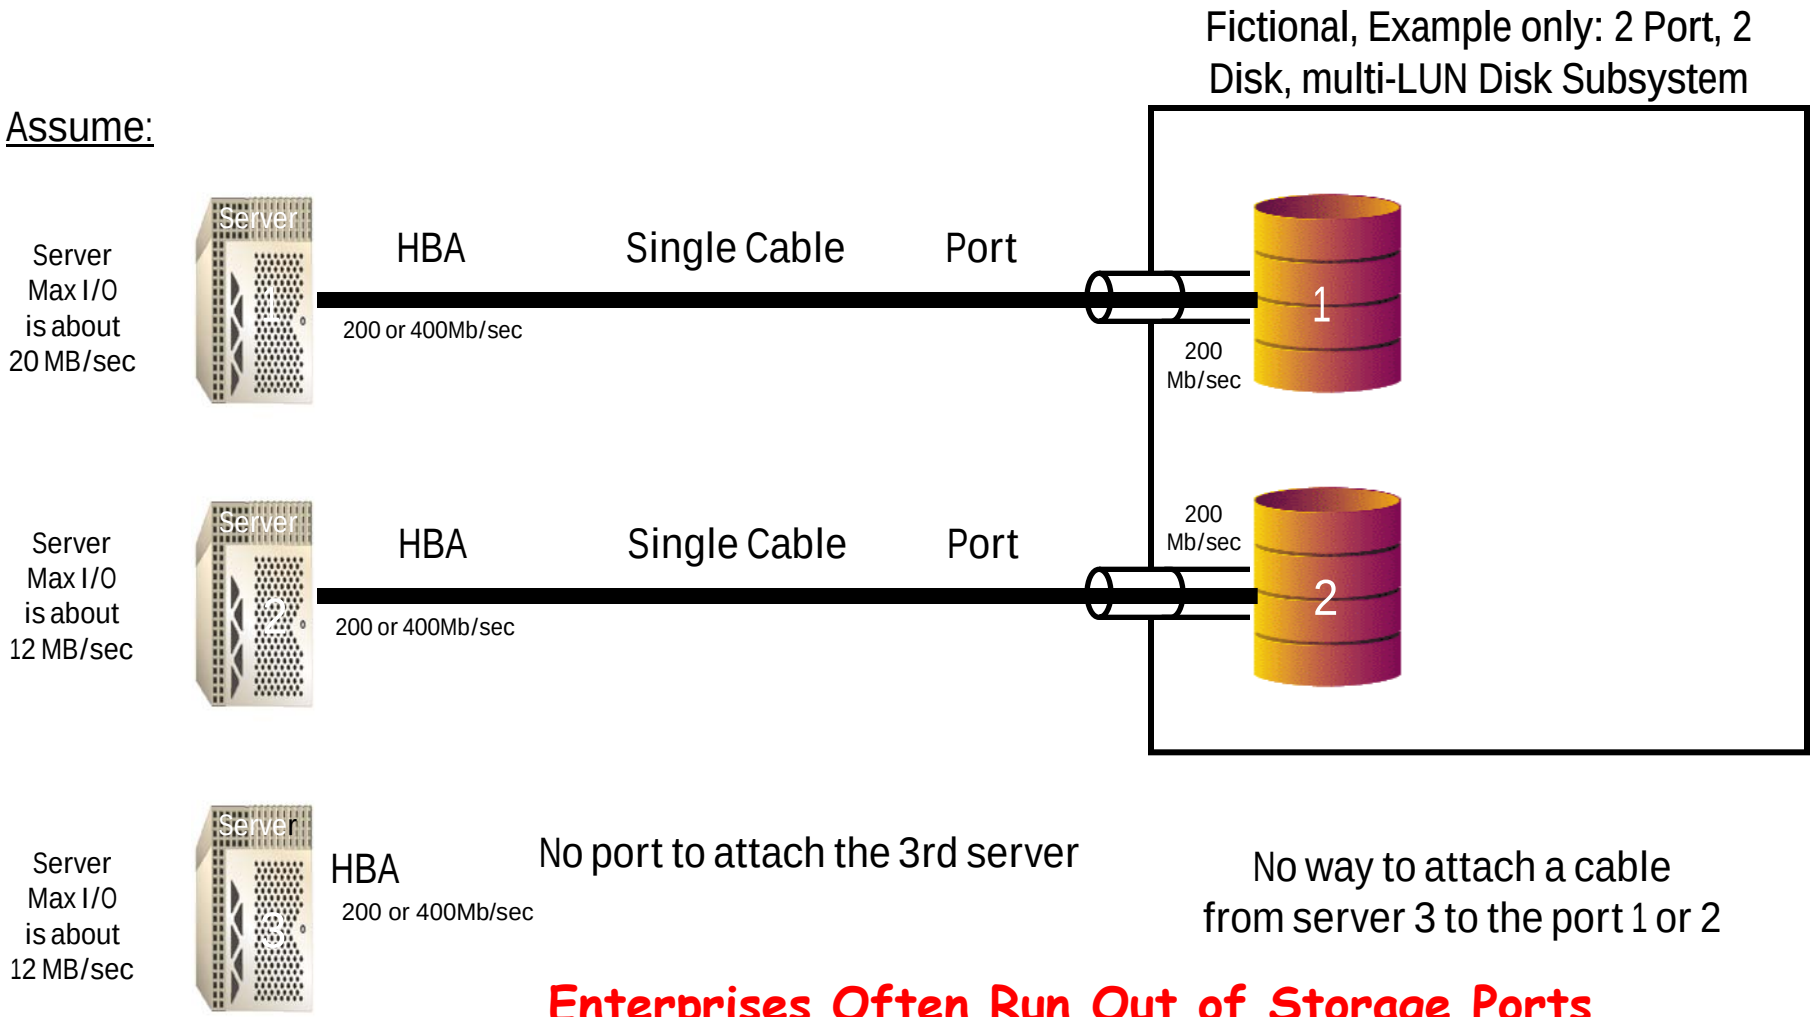

…To Share Storage Among Many Users!

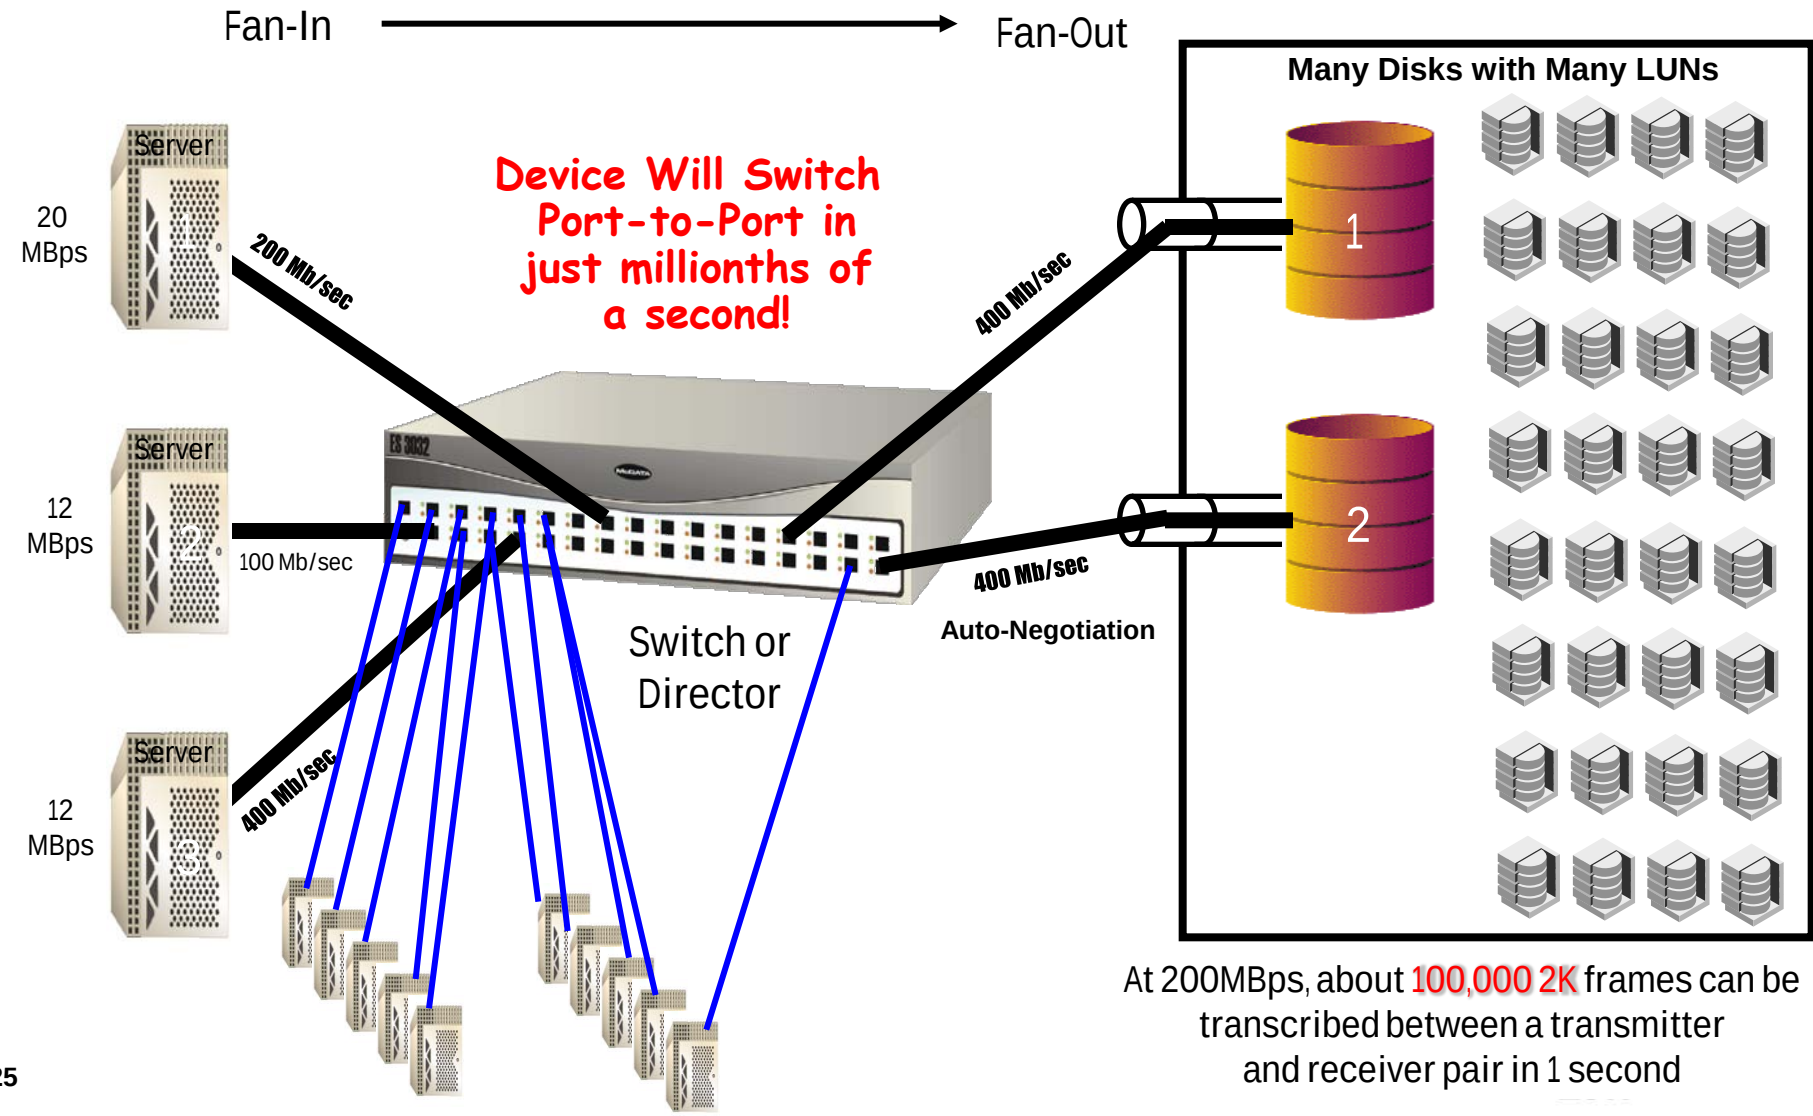

At 200MBps, about 100,000 2K frames can be transcribed between a transmitter and receiver pair in 1 second

Agenda for Session 8485

- History
- Connectors, Cables, and Wavelengths
- Mainframe Protocol Evolution
- Types and Components of Storage
- **Let's talk Fibre Channel - (Blair)**
- FC Flow Control
- Fabric Routing / Virtual Fabrics / Partitioning
- Security / Zoning

Fibre Channel Protocol

SHARE
Technology · Connections · Results

Fibre Channel (FC) provides high speed transport for Upper level (ie. FICON or SCSI) payloads

- **Some Confusing Terms**
  - **“Fiber” is the glass cable used to attach connections**
  - **“Fibre” is the protocol/architecture used to transport frames**
- **FC is the “protocol” for a Storage Network – attributes are:**
  - **Highly Scale - Addressing for up to 16 million nodes**
  - **Various Switched Topologies**
  - **High Speeds - 100, 200, 400 or 800 MBps …and… 10Gb ISLs**
  - **Segments of up to 100 Km between switches**
  - **Support for multiple protocols like FICON and OPEN (SAN)**
  - **Support for Security via Zoning and Prohibit/Allow Matrix**

Fibre Channel Overview

- **Fibre Channel is primarily utilized as a networked form of SCSI (open systems - SAN) or CCWs (System z - FICON)**
  - **Actually, the lower layers of Fibre Channel are generic**
    - **Able to transport multiple data types such as video on demand and Internet Protocol**
  - **But, most common deployment is Fibre Channel Protocol (FCP)**
    - **FCP is an upper layer protocol that provides for the transmission of SCSI commands and data over the Fibre Channel transport layers**
  - **Next most common deployment is Fiber Connection (FICON)**
    - **FICON is an upper layer protocol that provides for the transmission of mainframe CCW commands and data over the Fibre Channel transport layers**

The Fibre Channel Protocol

- FCP and FICON are just a part of the upper layer (FC-4) protocol
- They are compatible with existing lower layers in the protocol stack

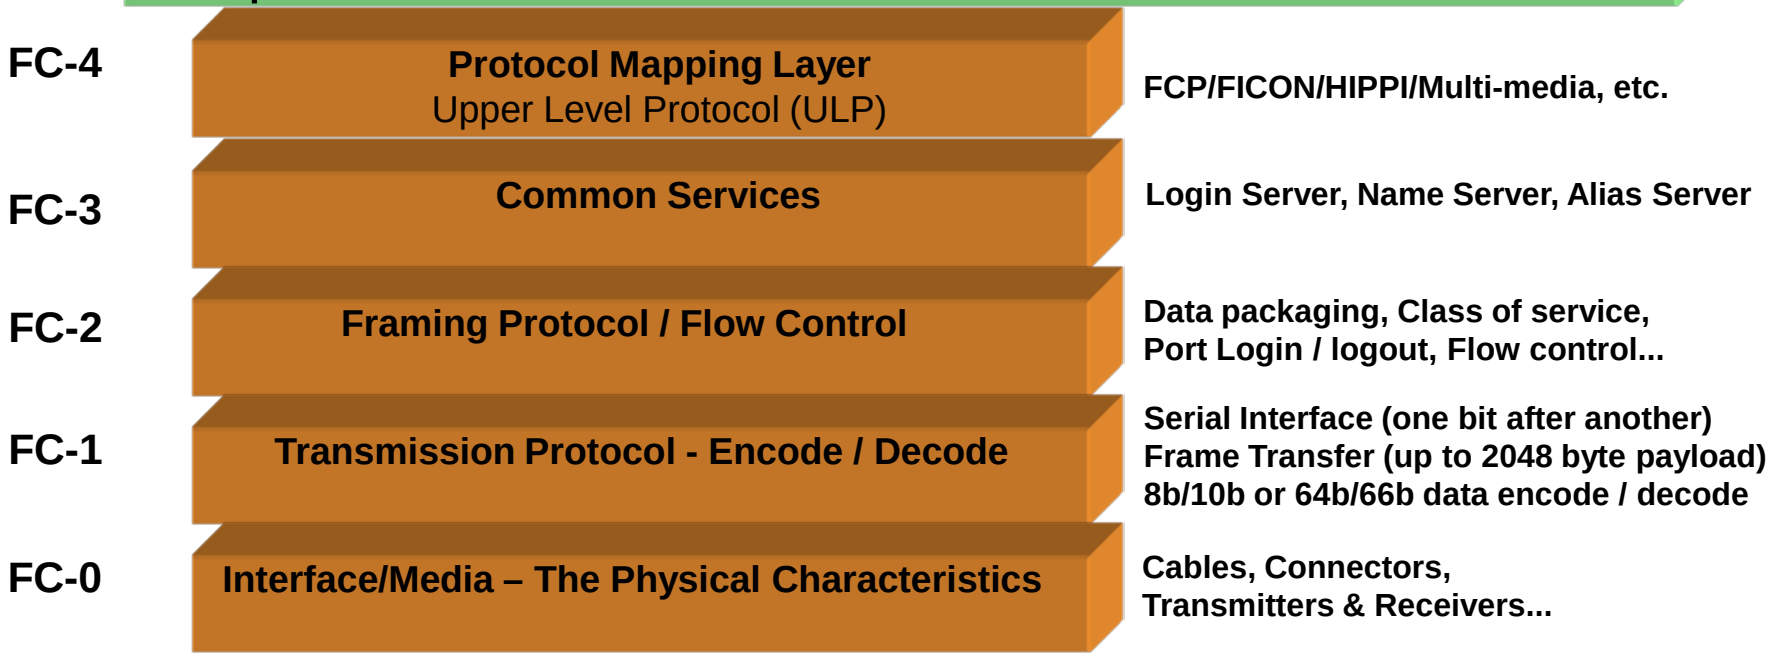

World Wide Names

- Each switch element is assigned a 64 bit WWN at time of manufacture
- Each switch port is assigned a 64 bit WWPN at the time manufacture
- During Fabric Logon (FLOGI) the switch identifies the WWN in the service parameters of the accept frame

These Address Assignments Can then Correlate Each Fabric Port with Switch Routing and the Fiber Channel ID (FCID)

Fabric Addressing

- The 24 bit FCID address is partitioned into 3 fields
  - Device or NPIV
  - Area
  - Domain
- This partitioning helps speed up routing
- Switch element assigns the address to N_Ports
- Address portioning is transparent to N_Ports

	8 bits	8 bits	8 bits
Switch/Open Topology	Switch Domain	Area	Device
FICON Topology	Switch Domain	Port Address	0 ….. or NPIV virtual addr.

Fibre Channel Frame Format

- All FC-2 frames follow the general frame format as shown below
- Idles are ‘Ordered Sets’ used for synchronization and basic signaling
- SOF – Start-of-Frame, EOF – End-of-Frame

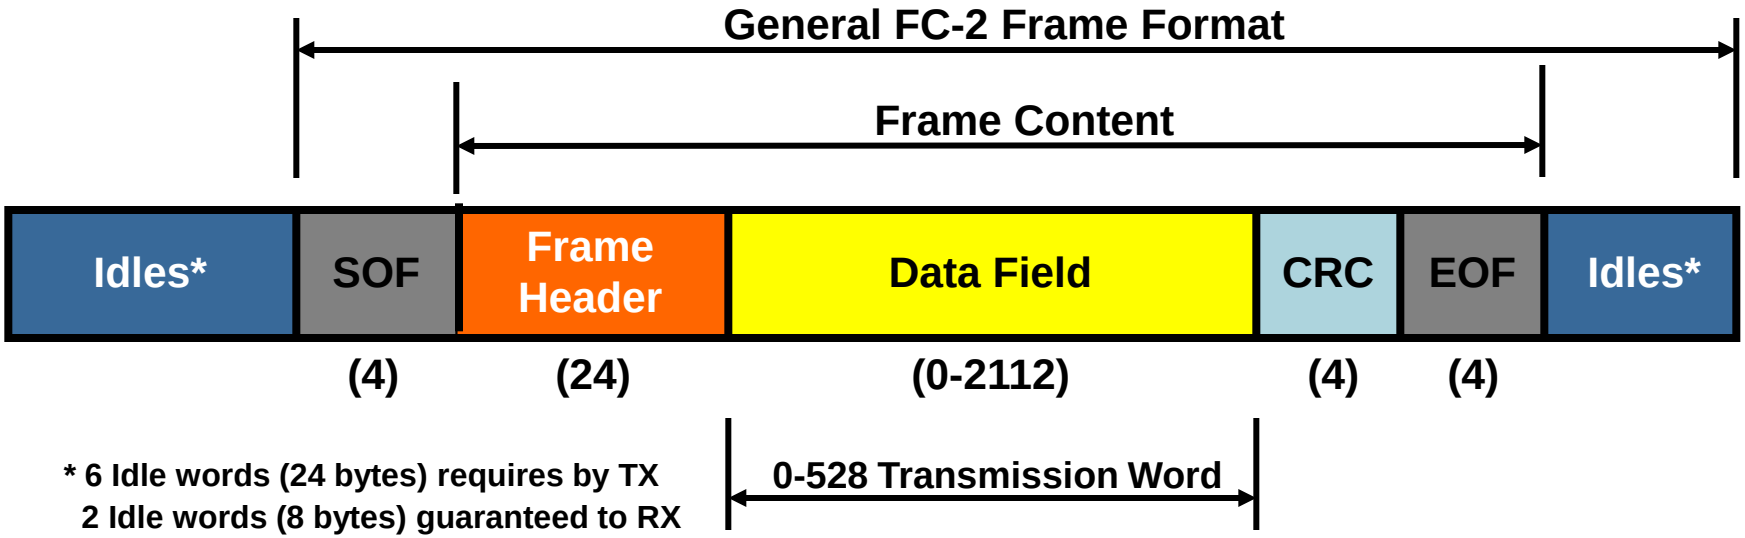

Fibre Channel Port Types

- **'N' port:** Node ports used to connect devices to switched fabric or point to point configurations.
- **'F' port:** Fabric ports residing on switches connecting 'N' port devices
- **'L' port:** Loop ports are used in arbitrated loop configurations to build networks without FC switches. These ports often also have 'N' port capabilities and are called 'NL' ports.
- **'E' port:** Expansion ports are essentially trunk ports used to connect two Fibre Channel switches
- **'G' port:** A generic port capable of operating as either an 'E' or 'F' port. Its also capable of acting in an 'L' port capacity. Auto Discovery.

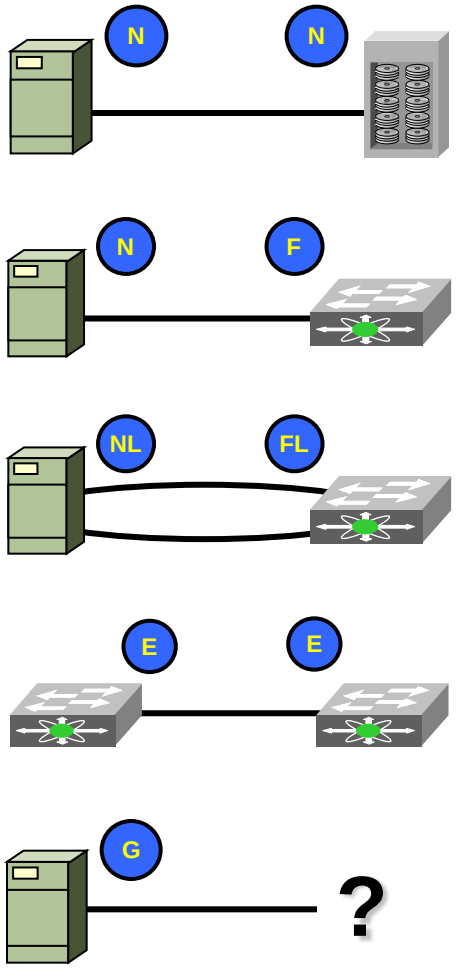

Fibre Channel Topologies

- **Point To Point** – communication between two devices only

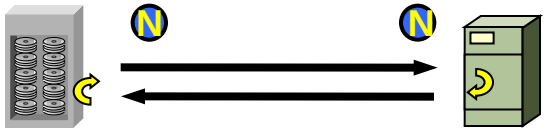

- **Arbitrated Loop** – loop consisting of up to 127 devices sharing total bandwidth (Mostly gone)

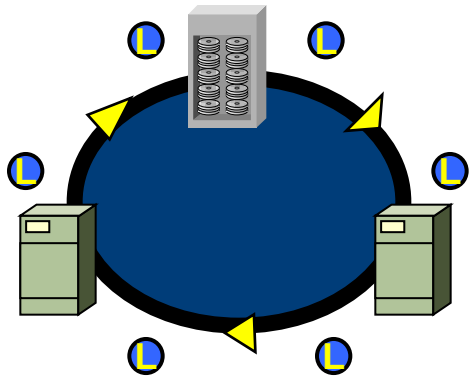

- **Switched Fabric** – switched network of devices similar to Ethernet/IP

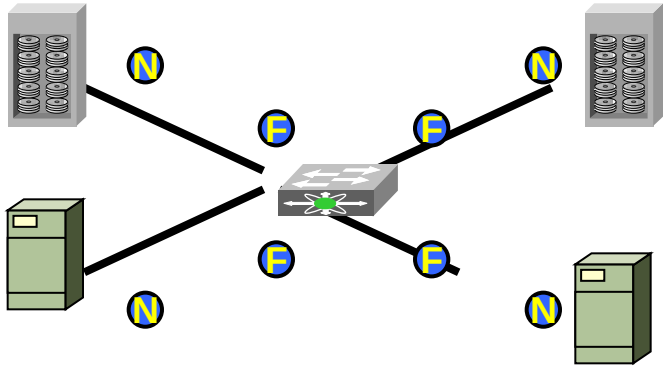

Switched Fabric

- Fabric may also interconnect public arbitrated loops using an FL_Port
- Dedicated bandwidth per port
- Provides routing based on destination address (Fibre Channel ID or FC_ID)
  - FC_ID includes Arbitrated Loop Physical Address (AL_PA) as the last
    octet if connecting to arbitrated loop
- Provides routing services via FSPF (Fibre Chanel Shortest Path First)
  - Only using only one area today

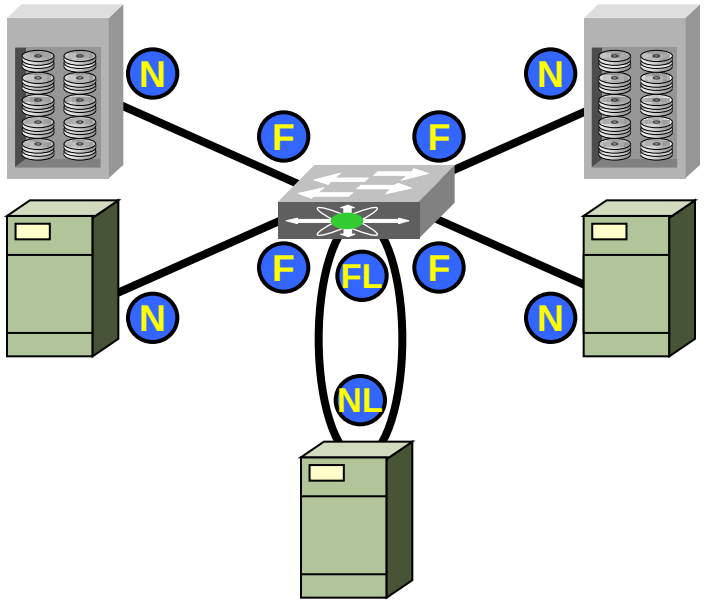

Inter-Switch Link (ISL)

- The interconnection between switches is called the Inter-Switch Link (ISL) or in FICON a Cascaded Link
  - E_Port to E_Port
  - For FICON, a 10Gbps link can ONLY BE a cascaded link (ISL)
- Supports all classes of service
  - Class 1, 2, 3, and a special Class F
- FC-PH permits consecutive frames of a sequence to be routed over different ISL links for maximum throughput

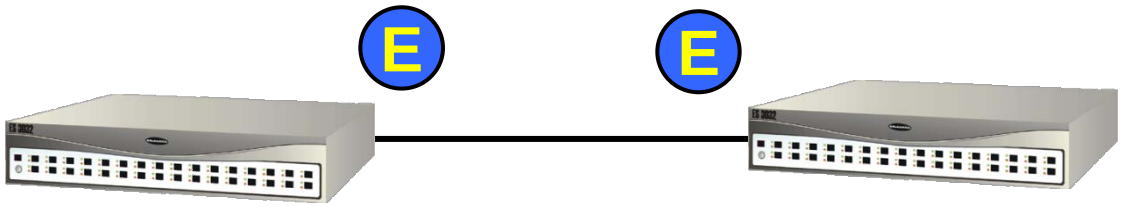

Typical SAN Topology
Dual Core (Non-Cascade for FICON)

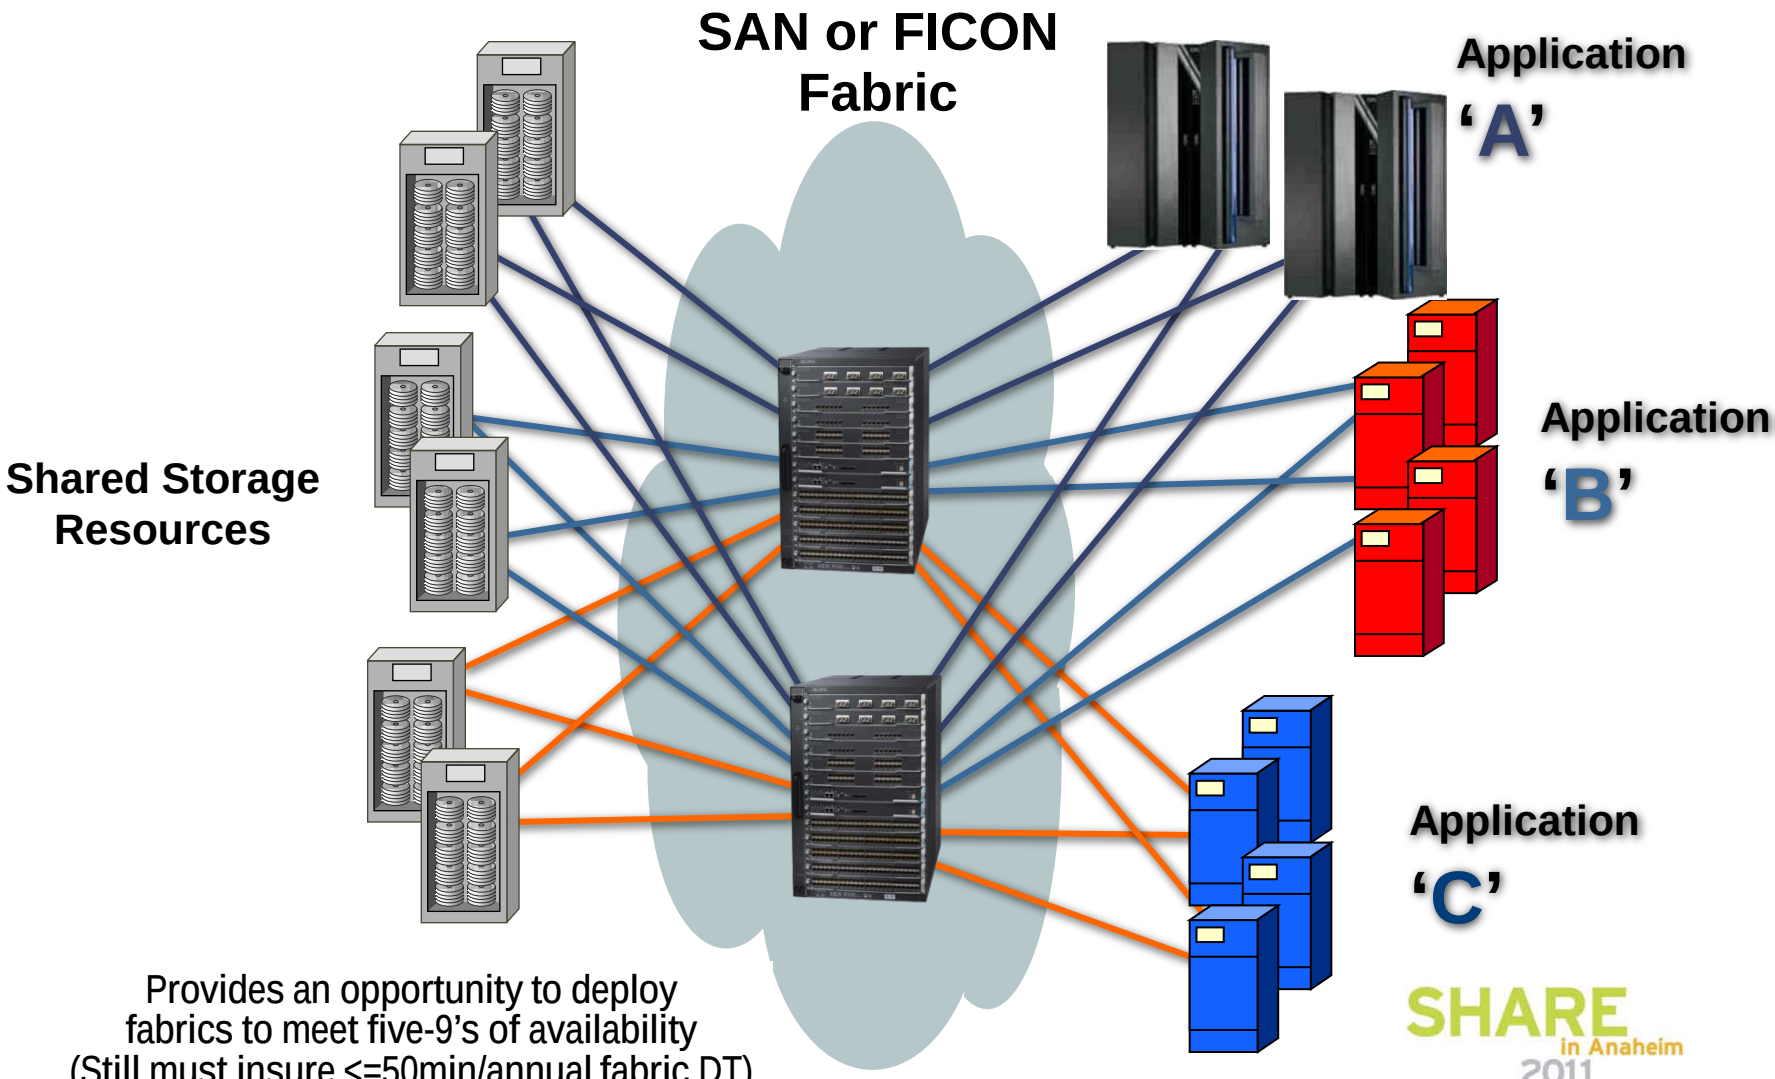

Provides an opportunity to deploy
fabrics to meet five-9's of availability
(Still must insure <=50min/annual fabric DT)

Typical SAN Topology
Core-Edge

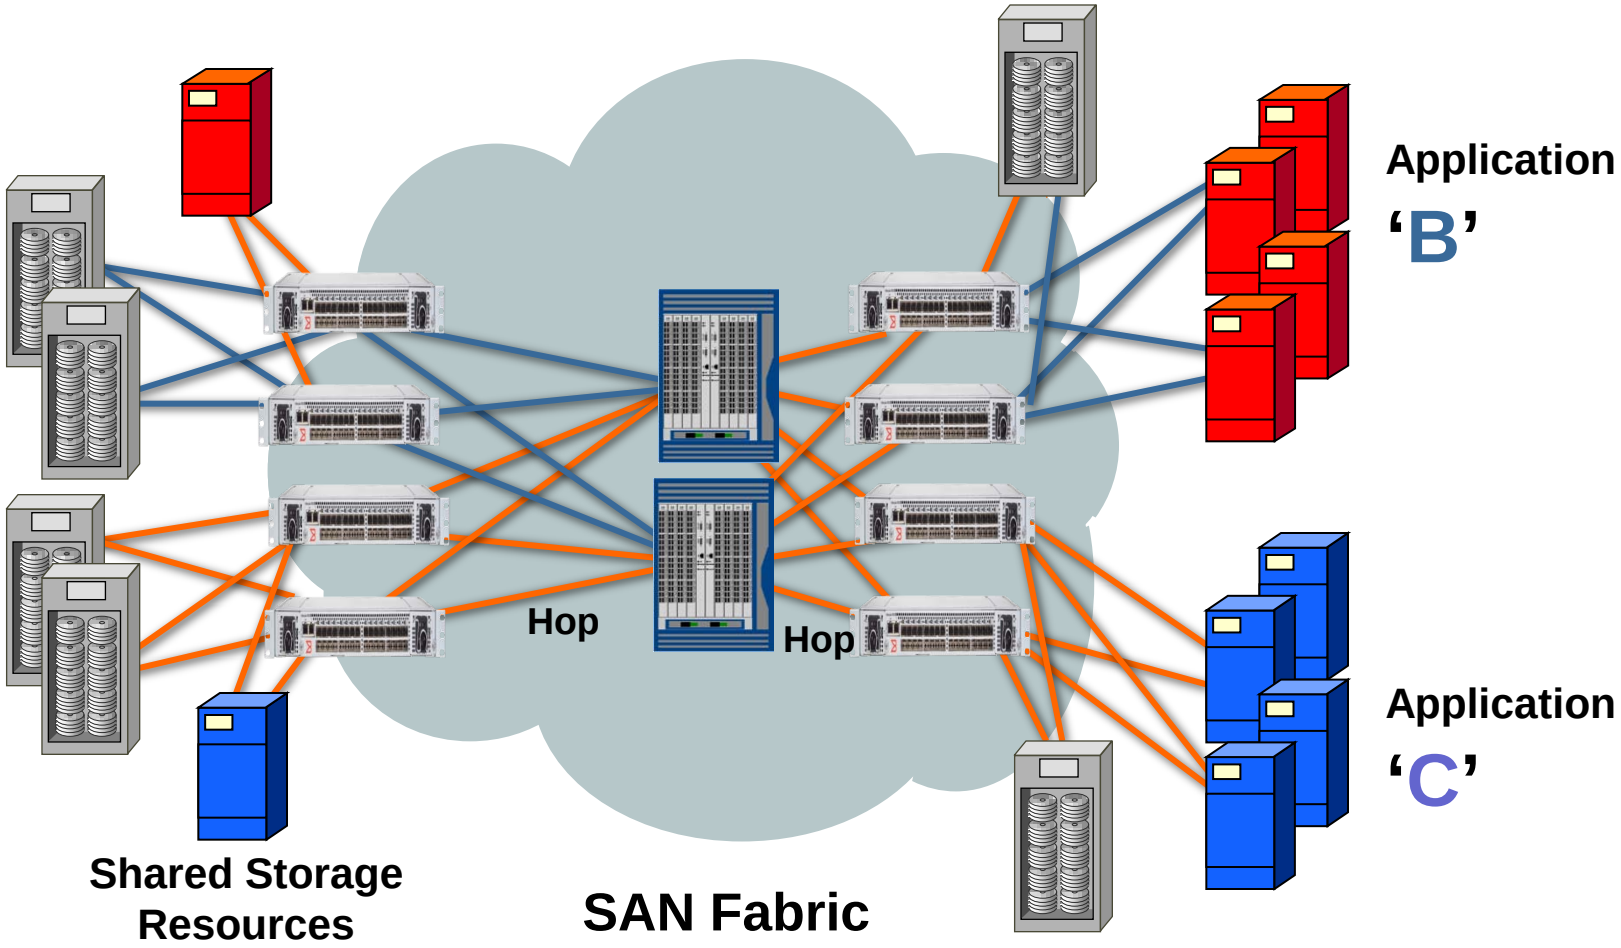

* Requires Multi-Hop So Not Currently Supported for FICON (due to IBM Qualification)

Typical SAN Topology (Cascaded FICON)

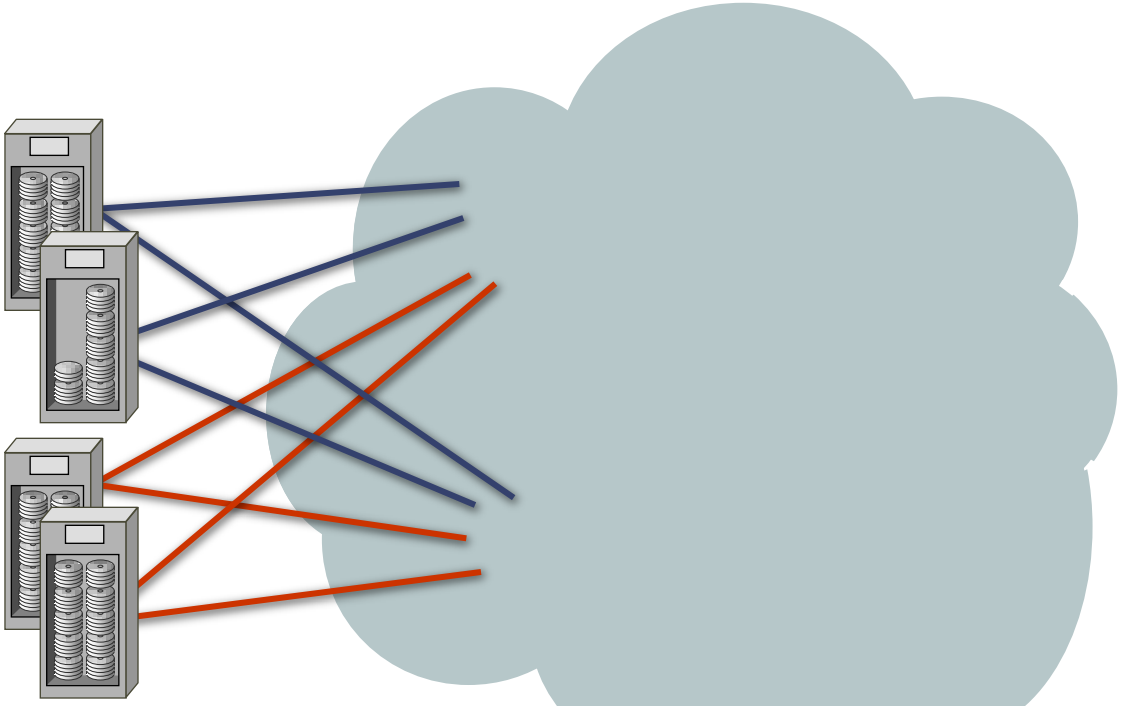

Shared Storage Resources

ISL Aggregation

Port Aggregation Is Used to Create a Single Logical ISL from multiple Physical ISLs

Different names depending upon the vendor

Brocade = Trunking
Cisco = Port Channel

- **Increases bandwidth and availability**
- **Simplifies Topology**
- **Usually some load balancing**
- **Interfaces can both be added and removed in a non-disruptive manner in production environments**
- **Preserves FC guarantee of in-order delivery (IOD)**

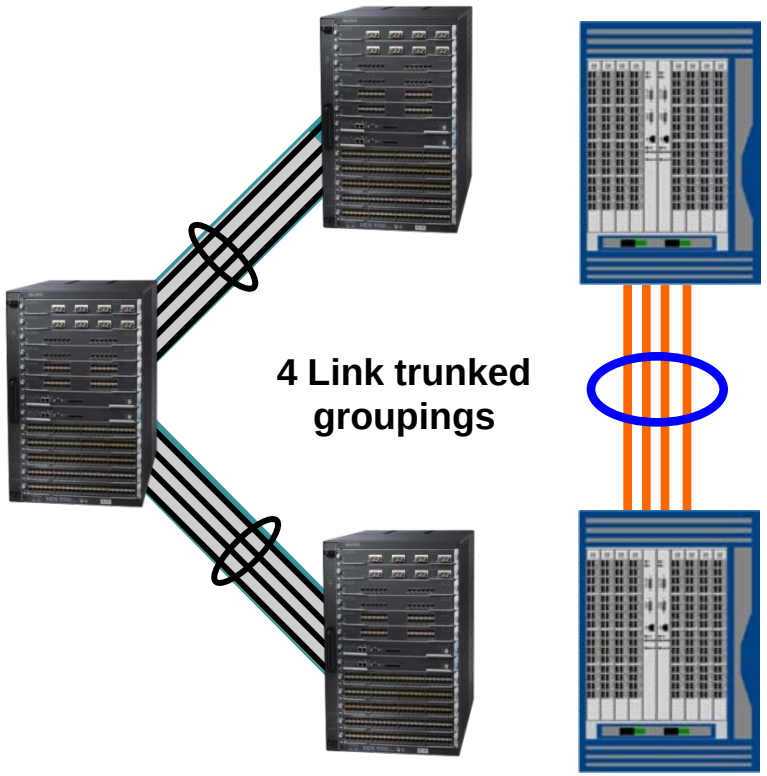

Agenda for Session 8485

- History
- Connectors, Cables, and Wavelengths
- Mainframe Protocol Evolution
- Types and Components of Storage
- Let’s talk Fibre Channel
- **FC Flow Control - (Blair)**
- Fabric Routing / Virtual Fabrics / Partitioning
- Security / Zoning

Buffer Credits in a Network

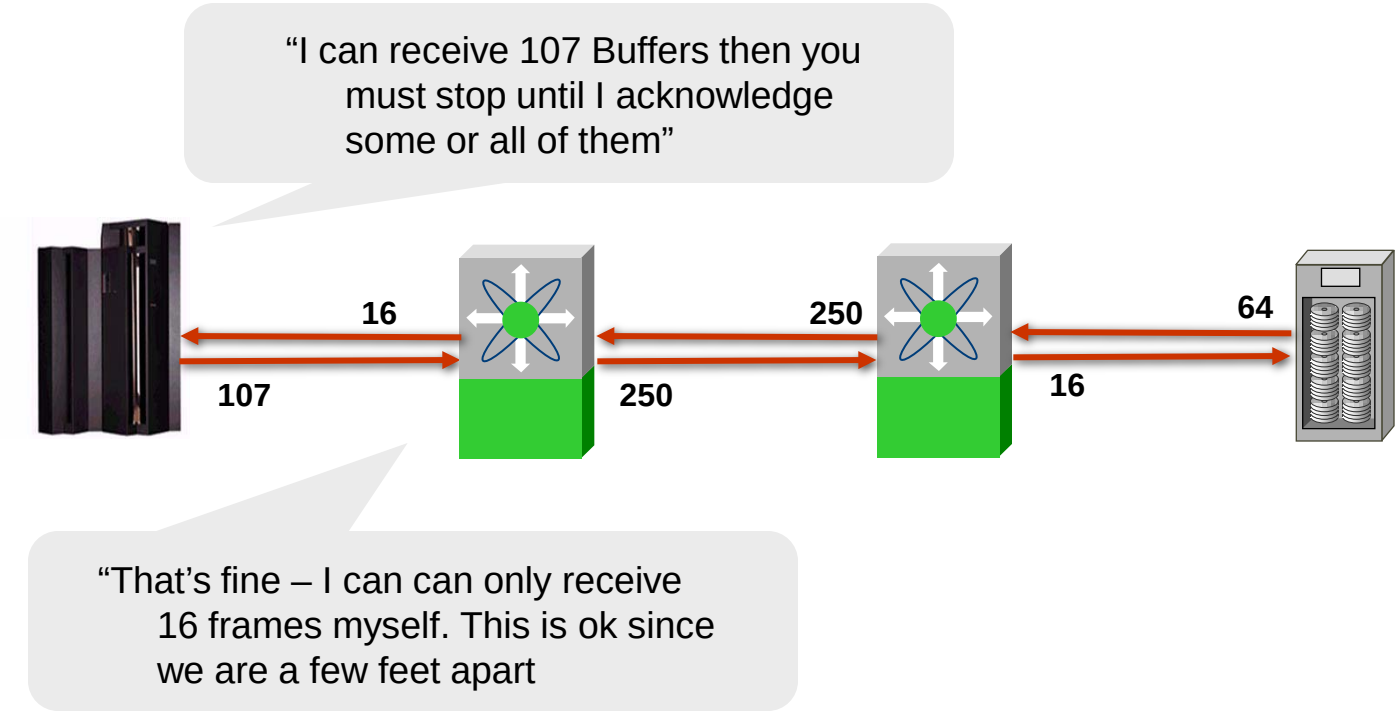

FC Buffer Credits and Flow Control

- BB_Credits are the “admission control” mechanism in FC to ensure that FC switches don’t run out of buffers (FC Switches cannot drop frames)
- For Devices operating at FC Class 3 (most devices), Buffer Credits are determined at login.
- BB_Credits are the **only** flow-control mechanism for FC Class 3.

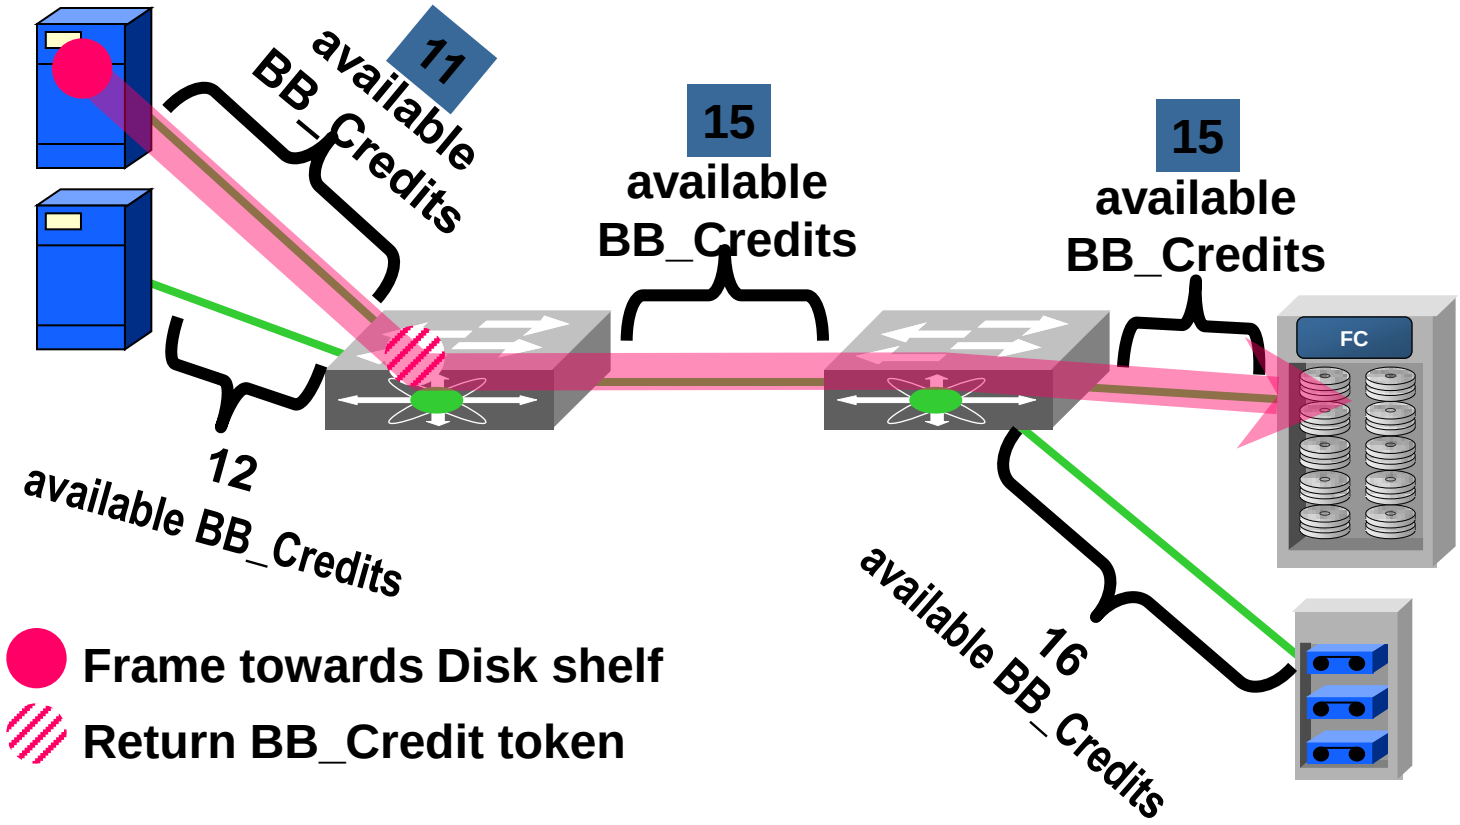

Buffer Credits (BB_Credits): Working Clean

Buffer Credits are a “Flow Control” mechanism to assure that frames are sent correctly

In an ideal FC network, there is no blocking in any device connected to the fabric. (all devices can process frames at the same rate and negotiate equal levels of BB_Credits)

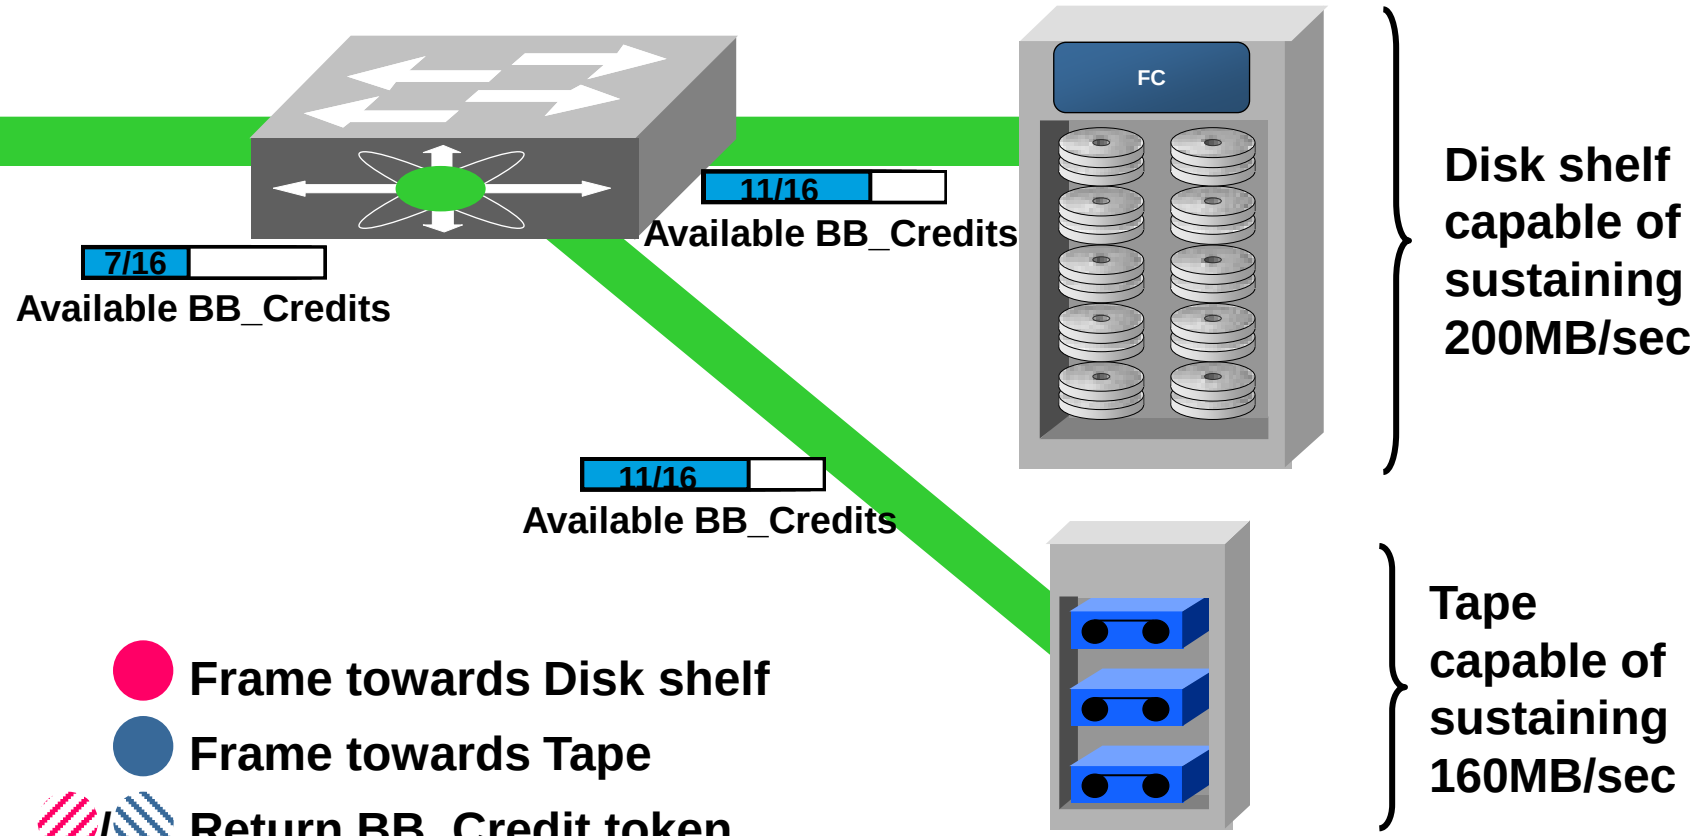

FC BB_Credits and Distance

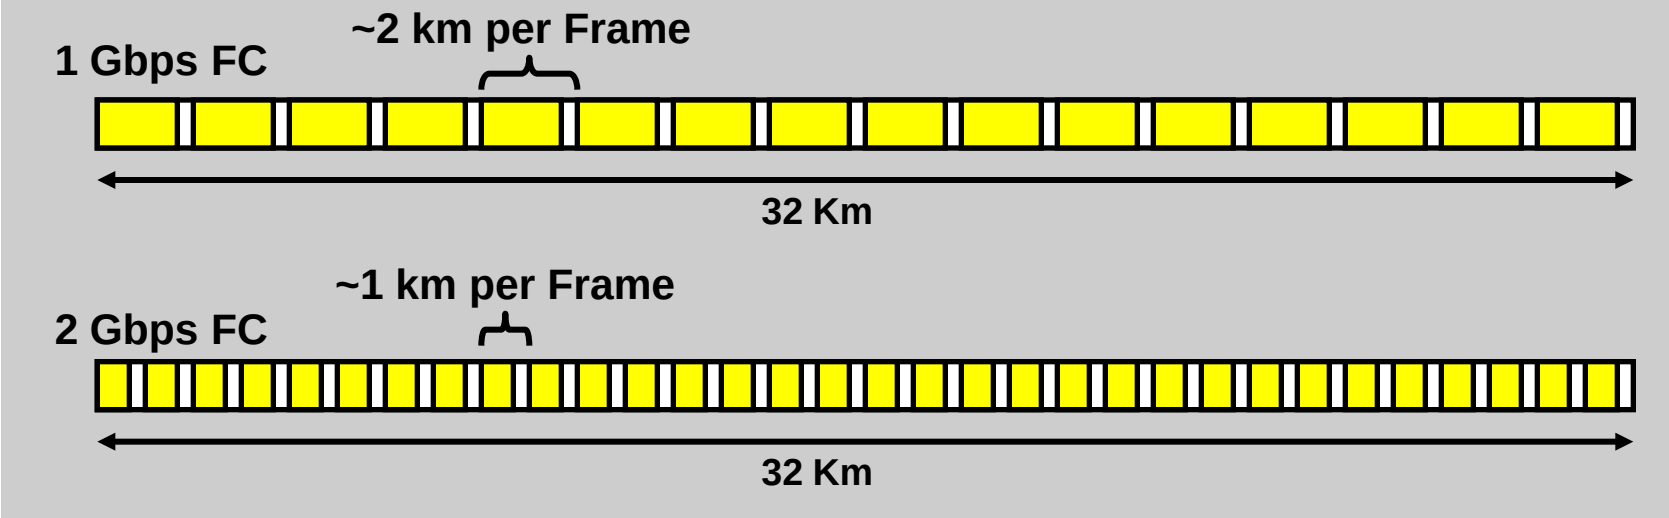

- BB_Credits are used to ensure enough FC frames in flight
- A full (2112 byte) FC frame is approximately 2 km long @ 1 Gbps and approximately 1 km long @ 2 Gbps and approximately .5 km long @ 4 Gbps
- As distance increases, the BB_Credits need to increase as well
- Shorter frames will require additional BB_Credits to ensure a "full pipe" – and ***USUALLY*** it is smaller frames that are sent!
- Insufficient BB_Credits will throttle performance—no data will be transmitted until R_RDY is returned

Number of Buffer Credits - Reality

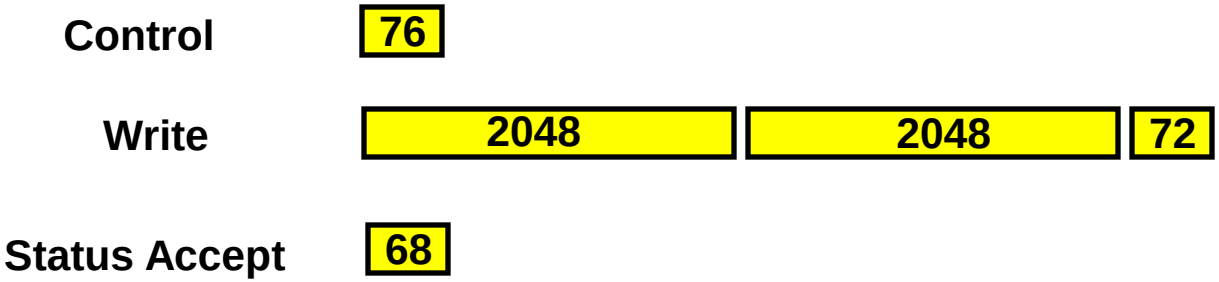

- Simple 4K write
- Will not fit into 2 buffers because of headers for FC as well as SB3

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

Buffer Credits Required
By Size of Frame and Link Speed

As distance link speed grows, so does the need for buffer credits!

FICON Director Activity Report (RMF 74-7)

(requires CUP Code)

FICON DIRECTOR ACTIVITY

z/OS V1R4 SYSTEM ID KS01 START 10/03/2005-13.55.00 INTERVAL 000.05.00
RPT VERSION V1R2 RMF END 10/03/2005-14.00.00 CYCLE 1.000 SECONDS
IODF = 0C CR-DATE: 07/07/2005 CR-TIME: 18.47.41 ACT: POR
MAN: MCD PLANT: 01 SERIAL: 0000013A2

	UNIT	ID	AVG FRAME PACING	READ		DTH (MB/SEC) -- WRITE --	ERROR COUNT
	CHP-H	47	0	433		1.52	0
05	CU	----	0	830	1195	5.81	0
06	CU	----	0	140	69	0.00	0
07	CU	C052	0	591	85	0.00	0
	CU	C050					
08	CHP-H	45	0	400	577	1.55	0
	CU	----	0	1355	37	1.00	0
	------	----	0	0		0.00	0
			0			0.00	0
			0			1.56	0
			0			1.41	0
			0			0.00	0
0F	CU	C053	0	1773	78	0.08	0.00	0
	CU	C051					
10	CHP-H	46	0	366	716	1.06	2.32	0
11	CU	----	0	1099	393	4.62	0.97	0
12	------	----	0	0	0	0.00	0.00	0
13	------	----	0	0	0	0.00	0.00	0
14	CHP-H	50	0	533	832	0.32	0.58	0
15	CU	----	0	868	1223	2.25	5.55	0
16	CU	----	0	158	72	0.00	0.00	0
17	CU	C053	0	1761	77	0.09	0.00	0
	CU	C051					
18	CHP-H	49	0	378	745	1.04	2.33	0
19	CU	----	0	1118	399	4.83	0.99	0
1A	------	----	0	0	0	0.00	0.00	0
1B	------	----	0	0	0	0.00	0.00	0
1C	CHP-H	51	0	737	535	0.34	0.17	0
1D	CU	----	0	877	1230	3.22	7.91	0
1E	CU	----	0	0	0	0.00	0.00	0
1F	CU	C052	0	590	82	0.00	0.00	0
	CU	C050					
20	CHP-H	4A	0	374	756	1.04	2.40	0
21	CU	----	0	1472	413	3.51	0.36	0
22	------	----	0	0	0	0.00	0.00	0
23	------	----	0	0	0	0.00	0.00	0

Agenda for Session 8485

- History
- Connectors, Cables, and Wavelengths
- Mainframe Protocol Evolution
- Types and Components of Storage
- Let’s talk Fibre Channel
- FC Flow Control
- **Fabric Routing / Virtual Fabrics / Partitioning (Blair/Lytle)**
- Security / Zoning

Fabric Routing Across ISLs: FSPF

- For FSPF a domain ID identifies a single switch
  - This limits the max number of switches that can support in the Fabric to 239 when FSPF is supported
    - Each Cisco VSAN is treated as a separate Fabric
    - Each Brocade Virtual Fabric is treated as a separate fabric
    - Each Brocade Physical Partition is treated as a separate fabric
- FSPF performs hop-by-hop routing
  - Each Cisco VSAN runs it's own FSPF process
    - Routing between VSAN's is done with Inter VSAN Routing (IVR)
  - Brocade will use FSPF routing if TRUNKING is not used
    - Old CNT/McDATA devices use FSPF for ISL routing
- FSPF supports hierarchical path selection
  - Provides the scalable routing tables in large topologies

Virtual Fabrics / Virtual SANs

- **Multiple Virtual SANs in a single hardware entity**
  - **Used to Converge SAN Islands**
  - **IETF RFC 4747**
  - **Although it is a Standard – Implementations are different per Vendor**

Brocade Virtual Fabrics for FICON

- Virtual Fabrics (VF) is the technology (and standard)
- OpenVSANs are the enabling feature on the Mi10K for the VF technology
- FICON and OpenVSAN on an Mi10K are mutually exclusive
- Administrative Domains (AD) are the enabling feature on the 48000 for the VF technology
- FICON can be connected into ports residing in an AD on an 48000
- Logical Switch (FOS 6.2) creates Virtual Fabrics on DCX family
- Port-by-Port assign ports to from 1 to 8 Virtual Fabrics on a DCX or DCX-4S chassis
- If you do not need it, do not use it
- Just another layer of management that can make FICON more complex
- If you do not need administrative separation of a physical fabric then do not implement Administrative Domains
- If you do need to carve a single fabric into portions that different administrators can manage, then Virtual Fabrics is perfect

Virtual Fabrics (VF) for DCX Family

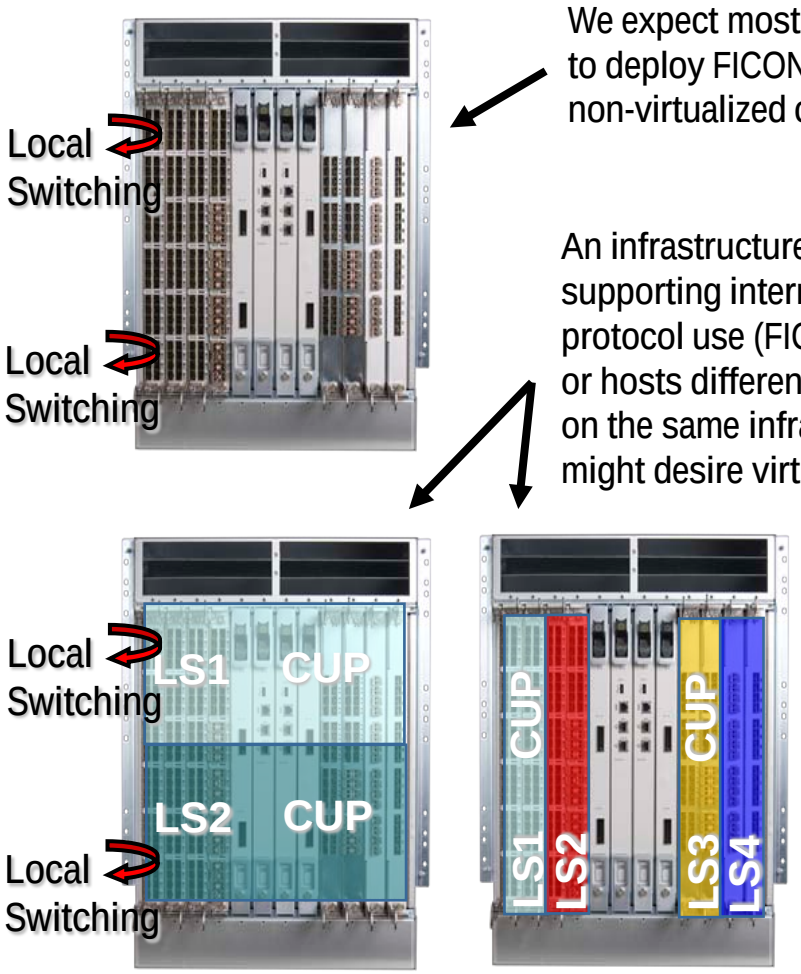

- VF is typically an OPTIONAL feature on the Brocade DCX family of FICON/SAN Directors
- Creating a Logical Switch, which enables Virtual Fabric, requires FOS 6.2 or later
- Creates new services for each Virtual Fabric
  - Name Server, etc
- Provides additional isolation of ports within a chassis
- On a DCX, 48-port blades for FICON can be used only when Virtual Fabrics enabled
  - DCX-4S never requires VF

Cisco's VSANs (Virtual SAN)

A way to Partition a Switch or SAN into a Virtual/Logical environment

- **Virtual SANs created from larger cost-effective redundant physical fabric**
- **Reduces wasted ports of the older "island" approach**
- **Hardware-based isolation**
- **Statistics can be gathered per VSAN**
- **Management per VSAN**
- **Unique Serial Number / CUP per FICON VSAN**

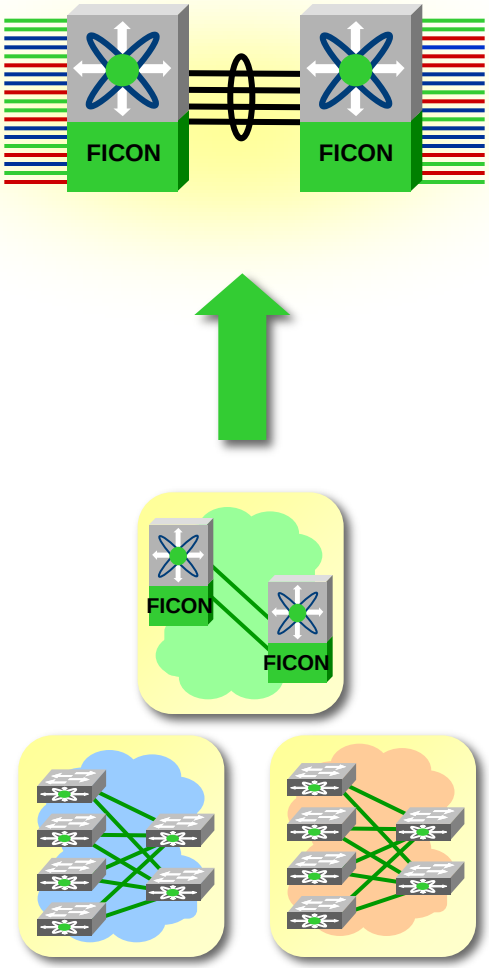

Cisco FICON VSANs – Workload Isolation

To build a cost saving fabric virtualization solution, 7 key services are required:

- Virtual Fabric Attachment – the ability to assign virtual fabric membership at the port level
- Multiprotocol Extensions – the ability to extend virtual fabric service to iSCSI, FCIP, FICON, etc.
- Virtual Fabric Services – the ability to create fabric services per virtual fabric (Login, Name, RSCNs, QoS, etc.)
- Virtual Fabric Diagnostics – the ability to troubleshoot per virtual fabric problems
- Virtual Fabric Security – the ability to define separate security policies per virtual fabric
- Virtual Fabric Management – the ability to map and manage virtual fabrics independently

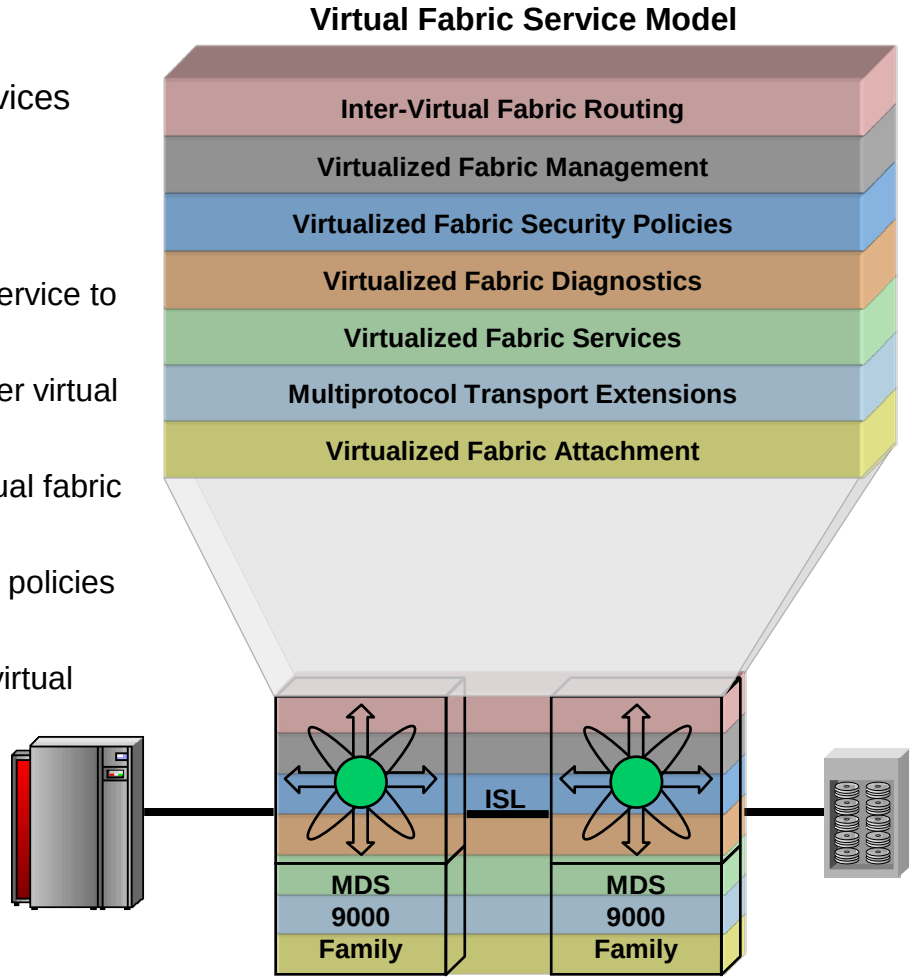

Agenda for Session 8485

- History
- Connectors, Cables, and Wavelengths
- Mainframe Protocol Evolution
- Types and Components of Storage
- Let’s talk Fibre Channel
- FC Flow Control
- Fabric Routing / Virtual Fabrics / Partitioning
- **Security / Zoning - (Blair)**

Zoning

- A logical grouping of fabric connected devices within a SAN (or virtual fabric)
- Zoning establishes access control
  - Devices within a zone can access each other
- Zoning increases security
  - Limiting access prevents unauthorized access
- Zone membership might be configured by:
  - Port World Wide Name (pWWN)—device
  - Fabric World Wide Name (fWWN)—fabric
  - Fibre Channel Identifier (FCID)
  - Fibre Channel Alias (FC_Alias)
  - IP address
  - Domain ID/port number
  - Interface

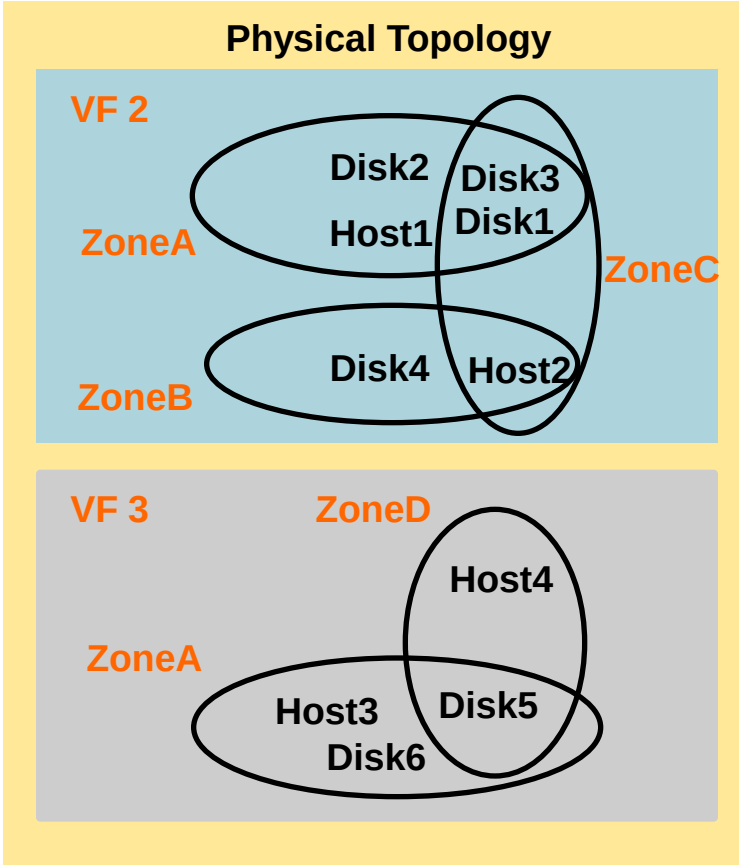

FICON Port Security

- Zoning
  - Able to be used with FICON (some vendors require it, others don't)
- Single Domain
  - Prohibit / Allow Matrix (ala. ESCON like port blocking/unblocking)
- Cascaded Configurations
  - Access to Cascaded resources controlled from IOCDS
  - Can also be done with Zoning but most choose not to

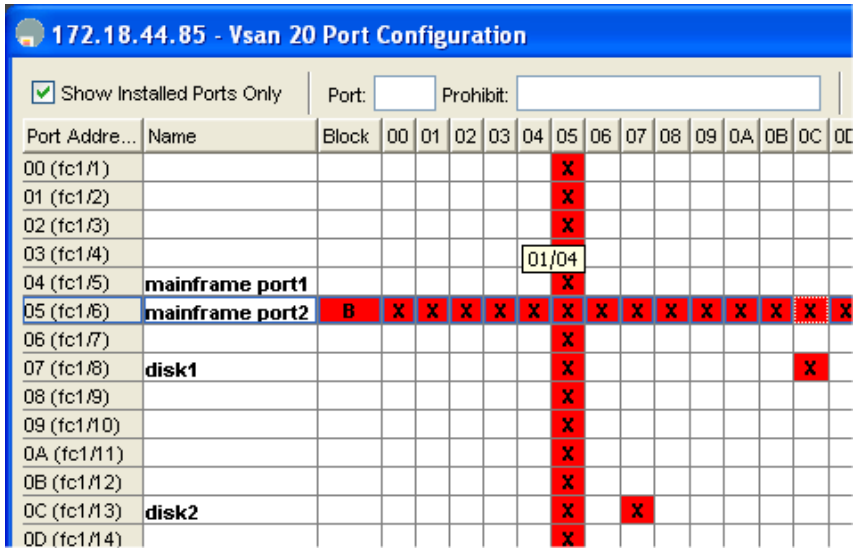

You can block or prohibit ports, eg:
- 05 is blocked and prohibited on all
- 04 can reach 07 and 0C
- 07 is prohibited from 0C

Please Attend Our Other SAN Sessions

BROCADE and CISCO SESSIONS

Tuesday, March 1, 2011: 3:00 PM-4:00 PM Room 211B
zSeries FICON and FCP fabrics - intermixing best practices

March 2, 2011: 4:30 PM-5:30 PM Room 211B
Customer Deployment Examples for FICON Technologies

Please Attend Our Other SAN Sessions

BROCADE SESSIONS

Monday Feb 28: 12:15-1:15 PM Room 207A:
Do as I Do, Not as I Say! Principled Leadership by Example

Monday Feb 28 : 1:30-2:30PM Room 211B:
FICON Buffer to Buffer Credits, Exchanges and Urban Legends

Tuesday March 1: 3:00-4:00 PM Room 211B:
zSeries FICON and FCP fabrics - intermixing best practices

Wednesday March 3: 1:30-2:30 PM Room 203A:
Planning and Implementing NPIV for System z

Wednesday March 2: 3:00-4:00 PM Room 211B:
Understanding FICON Performance

Thursday March 3: 3:00 PM-4:00 PM Room 208A:
Improving Your Ethernet Services with Virtual Cluster Switching

Please Attend Our Other SAN Sessions

CISCO SESSIONS

Tuesday, March 1, 2011: 3:00 PM-4:00 PM Room 212A
Using Enterprise Extender In the Network with Cisco SNA Switch

Tuesday, March 1: 4:30 PM-5:45 PM Room 212B
SAN Security Overview

Wednesday, March 2: 8:00 AM-9:00 AM Room 208A
ESCON Channels Will Be Phased Out - So, What Should I Do?

Wednesday, March 2, 2011: 1:30 PM-2:30 PM Room 208A
Datacenter Networking Convergence - Trends and Directions

Thursday, March 3: 9:30 AM-10:30 AM Room 212A
Mainframe OSA Connectivity and Routing in a Cisco Environment

Thursday, March 3: 1:30 PM-2:30 PM Room 211B
FICON Over IP - Technology and Customer Use

THANK YOU FOR ATTENDING!

Session 8485

Please fill out your evaluation forms