IBM STG

SAN Basics For Mainframers
Share Boston

Scott Drummond
spd@us.ibm.com

07/28/2010

Abstract

- The speaker will present the basics of SAN using mainframe references to explain the technologies. He will explore the Fibre Channel standard, SAN hardware, SAN software and other appropriate items related to SAN.

Alert, Alert, Alert!

- This presentation primarily covers that other stuff besides mainframes - yeah, you know - AIX, Windows, Solaris, Linux on RISC and z/Series and Intel, etc ...
- The industry is trying to take what we learned from mainframes and apply it to the non-mainframe systems under the name of Storage Area Network (SAN)
- z/OS is already a very good homogeneous, multivendor SAN
- This session will show how non-mainframe SANs are rolling out, using MVS (z/OS) terminology as a "Rosetta Stone" for translation purposes

Storage Management Challenge

- Costs of Storage Management
  - What is a Well Managed SAN?
    - The customer needs to assure that storage tasks performed include backup/archive, DR planning, problem, change, performance and capacity management, etc. – More than backup!
  - Non-centralized storage usually adds twice the people costs over centralized storage
    - This is assuming that Direct Attached Storage (DAS) is managed at all!

Data Size!

- Gigabytes (10*9) (Billion) of disk storage - small customers
- Terabytes (10*12) (Trillion) of disk storage - most customers and many individuals
- Petabytes (10*15) (Quadrillion) of disk storage – many large customers
- Petabytes (10*15) (Quadrillion) of tape storage – many large customers
- Exabytes (10*18) (Quintillion) of tape storage - a few accounts
- Zettabytes (10*21) (Sextillion) WW digital data – 1 Zettabyte - IDC
- Yottabytes (10*24) (Septillion) - How Long will it take?

IBM SAN Evolution

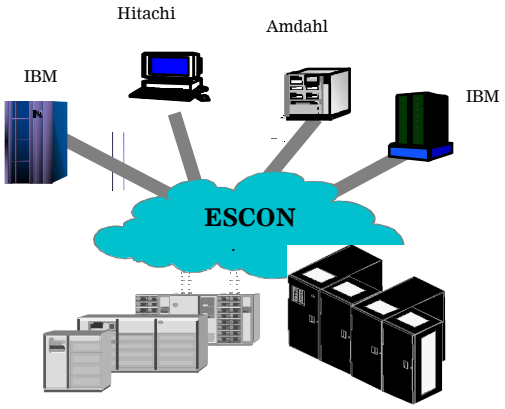

Mature technology
with Automated Policy Mgt and Device/Data Sharing Technology

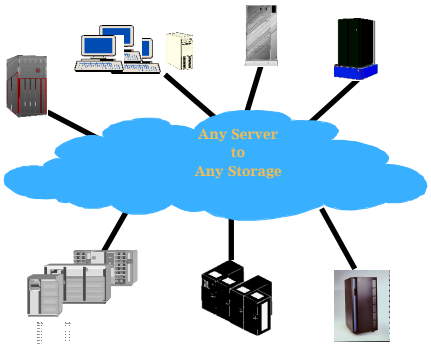

Heterogeneous Server and Storage SANs being built with technology from S/390s and LANs/WANs

LAN/WANs
(Ethernet and Token Ring)
Sophisticated Mgt Capabilities

IBM
- ☑ Experience
- ☑ Road map

Storage Area Network (SAN)

- *SAN* - ***Centrally managed*** high speed networks of ***multivendor*** storage subsystems, applications servers, clients and networking hardware that allow companies to exploit the ***value of their business information*** *via* ***universal access and sharing of resources.***

Information used to belong to the server....
NOW it belongs to the Enterprise !

- **SAN** - A network whose primary purpose is the transfer of data between computer systems and storage elements and among storage elements. A SAN consists of a communications structure, which provides physical connections, and a management layer, which organizes the connections, storage elements and computer systems so that data transfer is secure and robust. ***(SNIA)***

SAN - A Technical Definition

- A Storage Area Network is a highspeed dedicated network (usually Fibre Channel based) that offloads the backup/archive/restore/retrieve data I/O stream off the LAN.
- The SAN de-couples the ownership of the storage control unit from the attached servers in order to provide "storage pooling" allowing a higher utilization of the storage. It does not break the relationship of the file systems from the operating systems without special software such as virtualization functions.
- Because no one server is in control of the "pooled storage", the SAN needs to provide a "traffic cop" function between the servers when they access the pool of storage. The functions the SAN provides are called Zoning (path blocking) and LUN masking (LUN assignment). These mechanisms work in conjunction with each other to provide orderly access to data.

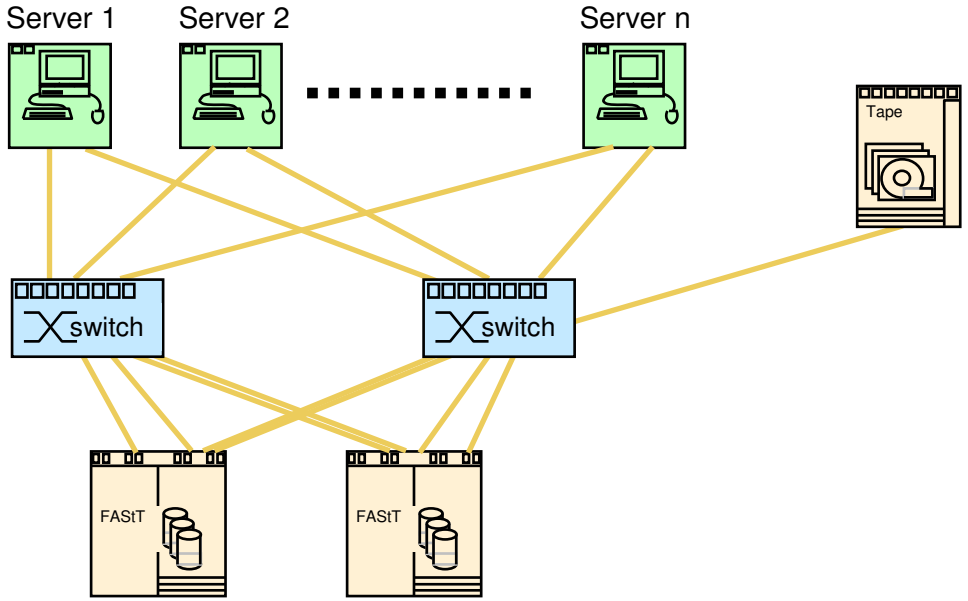

What Does A SAN Enable?

- Any-server to any-storage connectivity
- Resource sharing
- Offload the Ethernet LAN
  - Server-free data transfer
  - LAN-free backups
- Decoupled processor/storage growth
- Single Point of Control for Management

Storage Pooling Benefits via Fibre Channel SANs

- Enables Long Term Benefits
  - Leverage Information
  - Enables consistent management of data
    - Centralized, dedicated storage management team
    - Consistent policies, practices & procedures
    - Consistent security
    - Enables problem and change management
    - Prerequisite for Storage Virtualization & Information Life Cycle Mgt (ILM…really HSM!)
  - Enables better planning (capacity mgt, storage resource management (SRM), etc.)
  - Better positioned for Business Continuity solutions
  - Better positioned to implement Service Level Agreements & Charge Back facilities

z/OS ESCON/FICON SAN

- z/OS is the most advanced SAN in the industry:
  - The "Switch" is the ESCON or FICON Director
  - Device Sharing (including the equivalents of "Zoning" and "LUN Masking") was built into the operating system during the 70's and 80's (DASD Sharing)
  - z/OS Dynamic Path Reconnect (DPR) is superior to any dual path software for open systems storage (SDD, Power Path)
  - Storage Pooling came along with System-Managed Storage in the late 80's and early 90's - also supports multivendor storage devices
  - Generalized Data Sharing and additional Device Sharing came into being with the Parallel Sysplex during the mid 90's
  - z/OS SAN provides workload balancing, priority queuing and the industry leading continuous availability and data integrity
  - z/OS SAN (FICON) will be able to share fabric resources with other SANS - FICON/FCP Intermix
  - z/OS SAN - FICON Cascading now supported

SAN Evolution

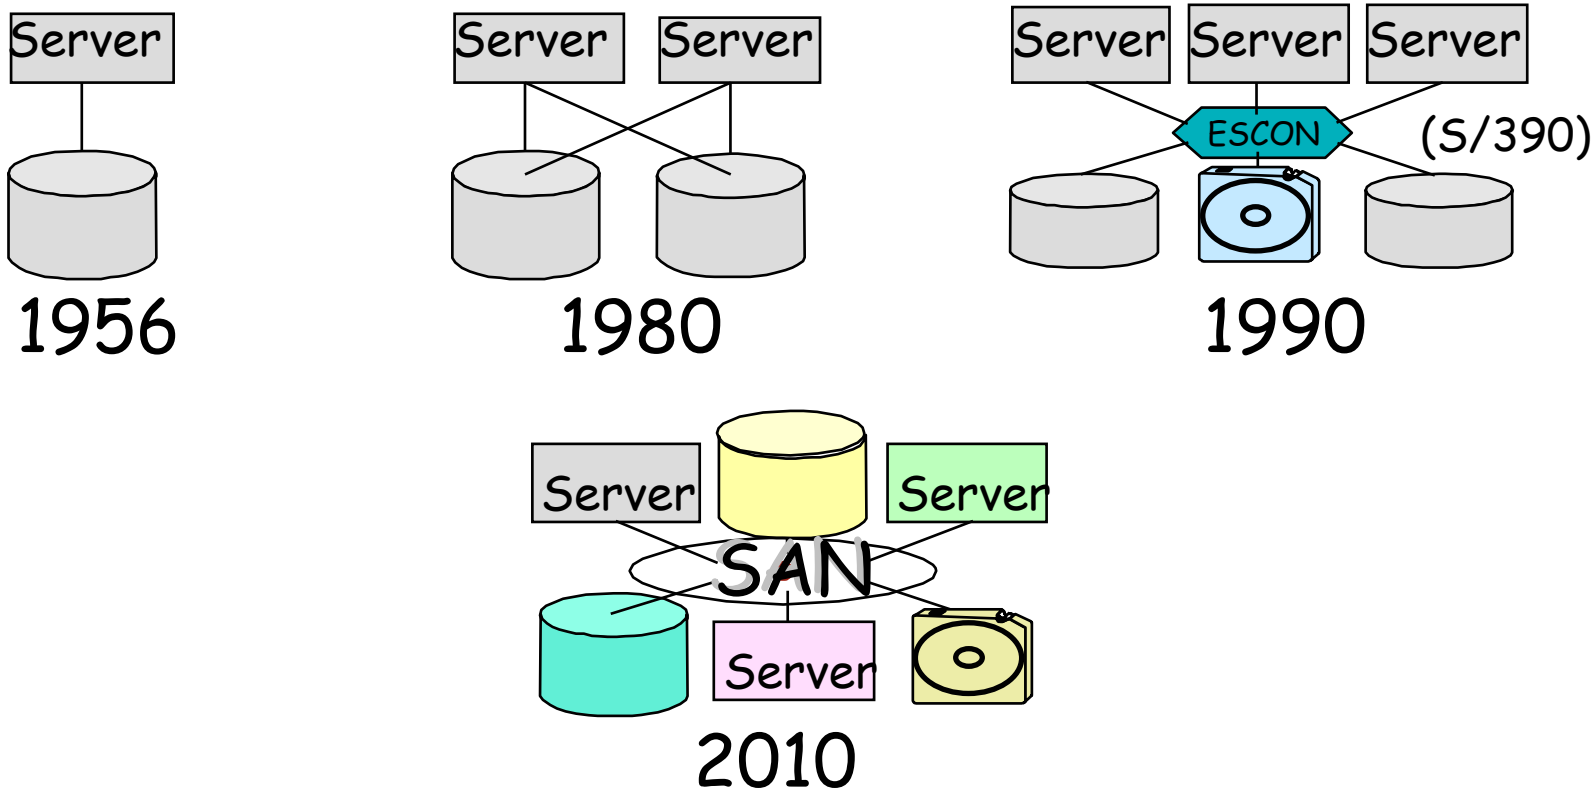

The Six "S"s of SAN

- **Servers**
  - Operating Systems, File Systems, Device Drivers/Dual Path Device Drivers, HBA's & Microcode
- **SAN Fabric Components**
  - Switches, Directors (and Blades), SAN Routers
- **Storage Virtualization**
  - Management, Flexibility and Efficiency
- **Storage**
  - Disk Systems, Tape Drives, Tape Libraries
- **Software**
  - Storage Resource Management, SAN Exploitation, ITIL and CMDB input and exploitation
- **Services**
  - Planning, Testing and Implementation, Education

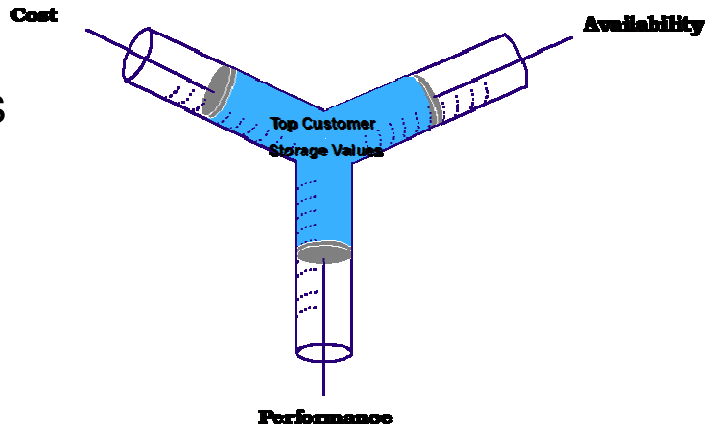

The Three Things That Matter

SAN Topologies

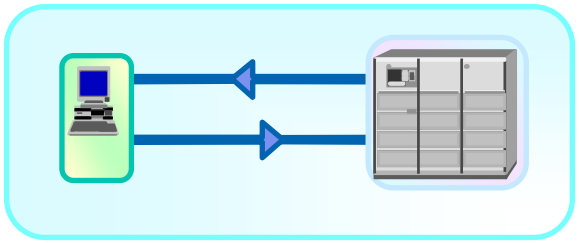

FC Point to Point

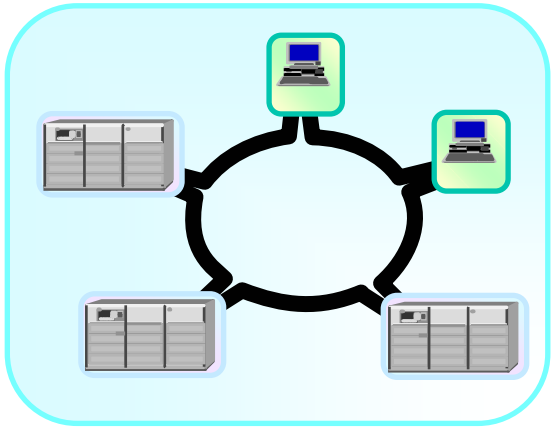

FC Arbitrated Loop

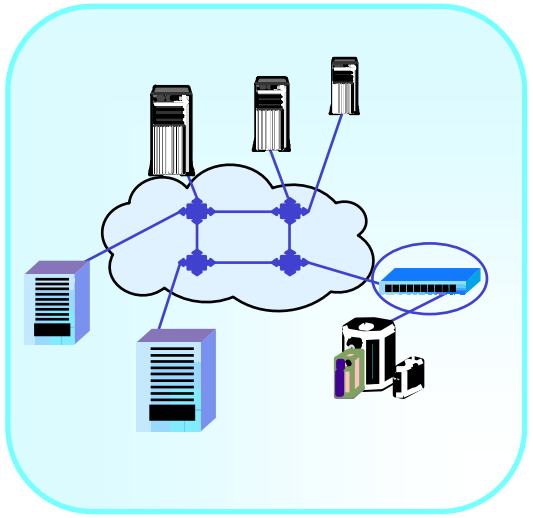

Switched Fabric

SCSI Roadmap

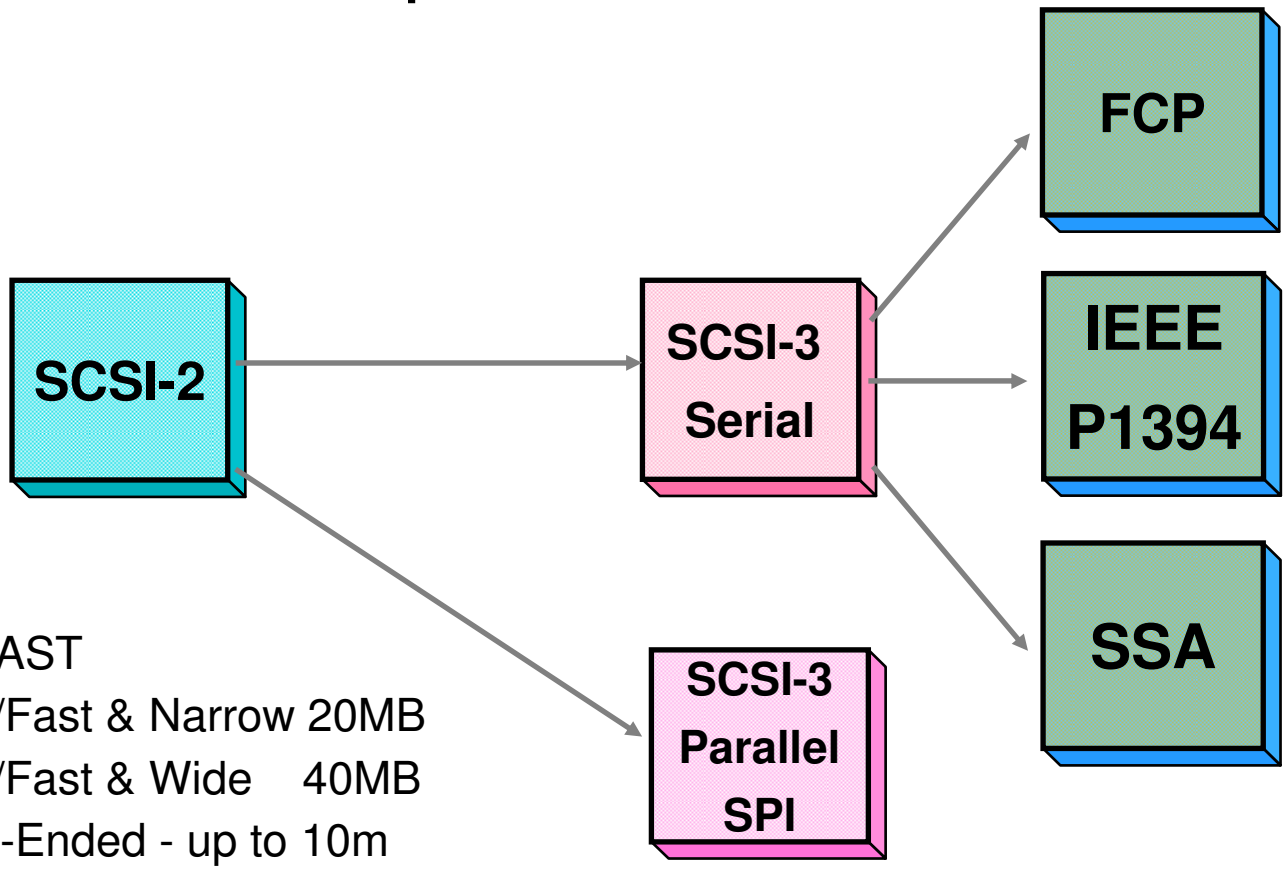

SCSI-FAST
SCSI-II/Fast & Narrow 20MB
SCSI-II/Fast & Wide 40MB
- Single-Ended - up to 10m
- Differential - up to 25m

Ultra2SCSI 80MB
Ultra3SCSI 160MB
Ultra4SCSI 320MB

The Fibre Channel Architecture

- Fibre Channel Architecture
  - An integrated set of rules (FC-0 thru FC-4) for serial data transfer between computers, devices and peripherals developed by INCITS (ANSI)

  See also www.t11.org
- FICON
  - IBM zSeries™ implementation of Fibre Channel Architecture
  - An industry standard under the name FC-SB-2
- FCP
  - Fibre Channel Protocol for SCSI
  - Mapping of the SCSI command protocol onto the Fibre Channel Architecture

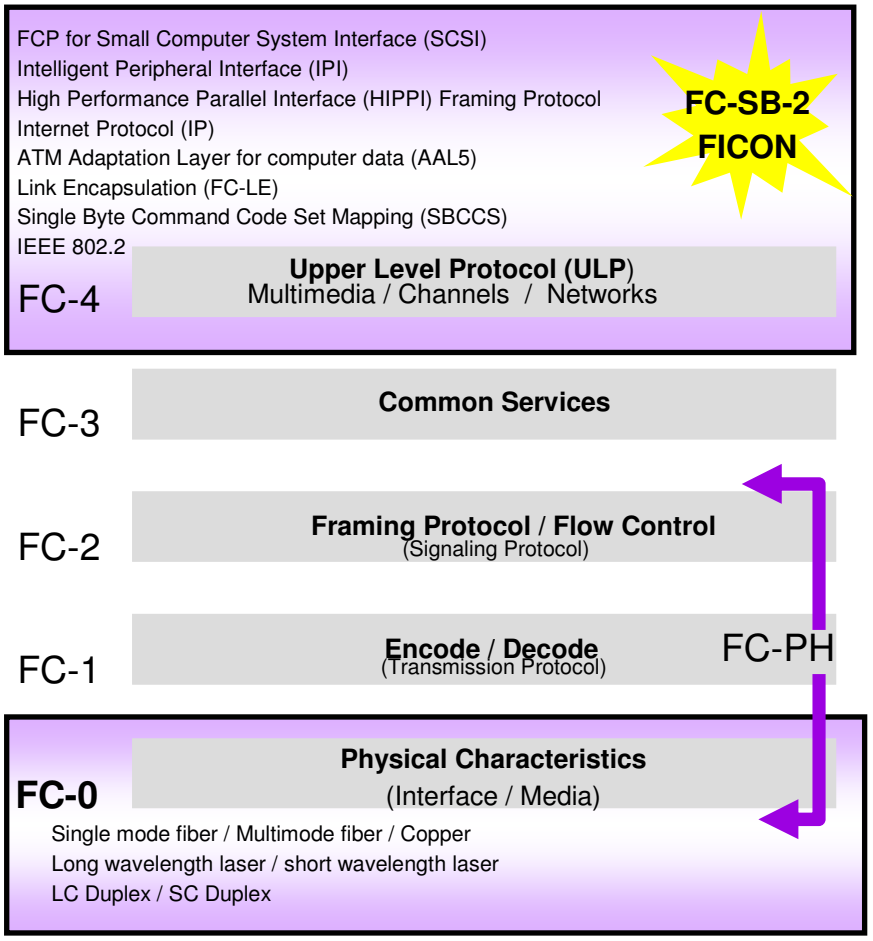

FCP Topologies

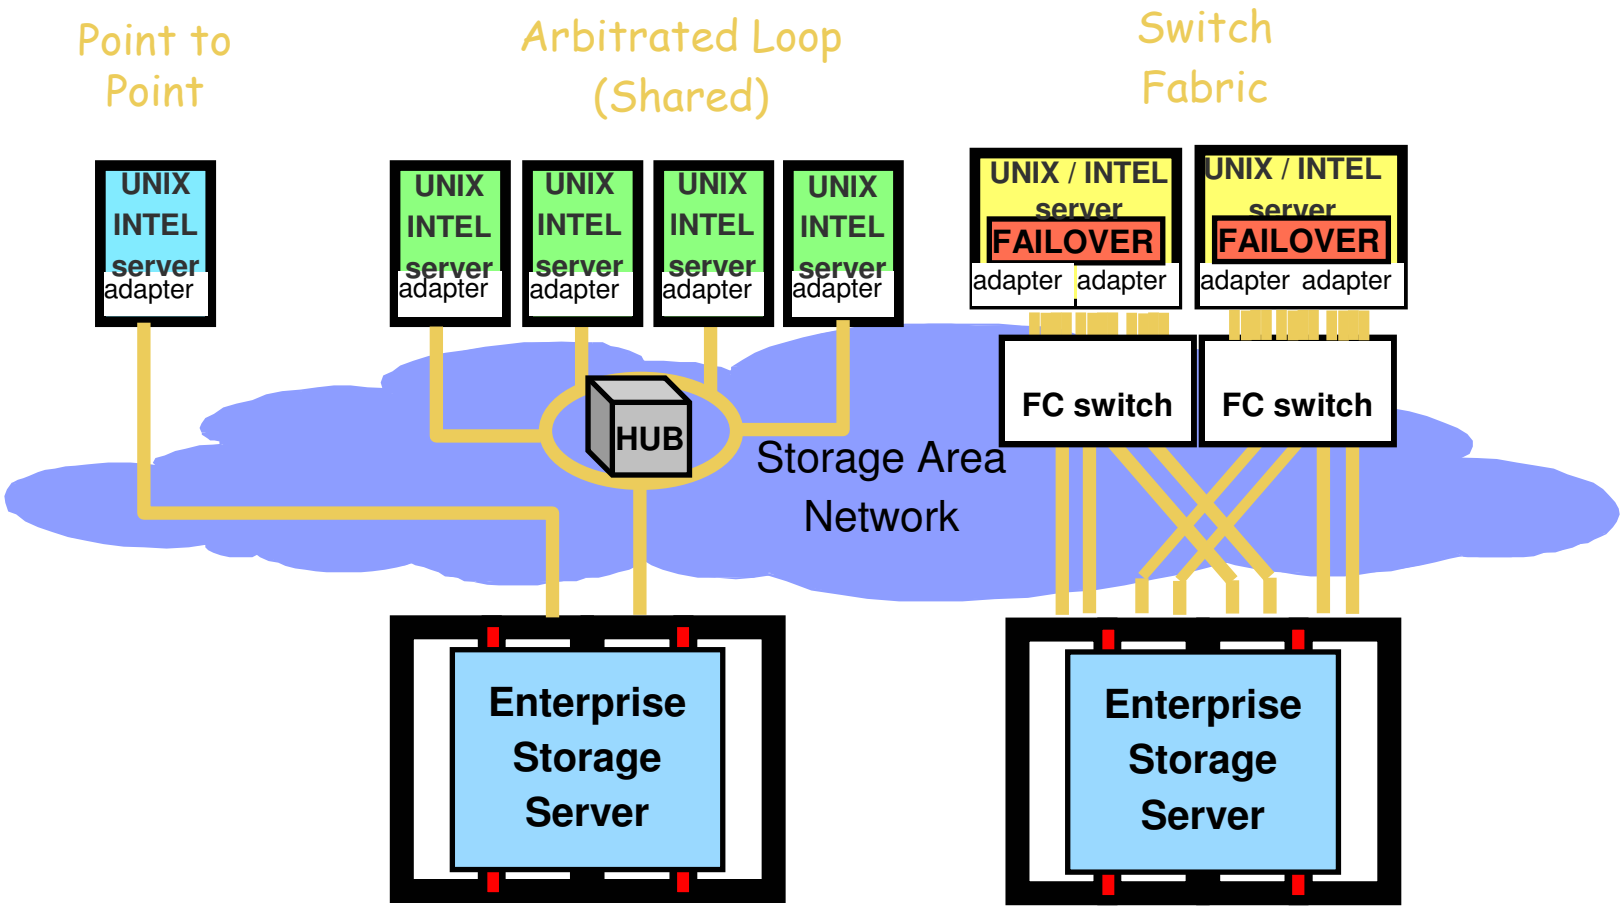

Internet SCSI (iSCSI) A Low Cost Technology

- **SCSI over IP Networks - "SAN" with IP fabric**
- **The iSCSI protocol provides for the most efficient "packing" of storage data into TCP/IP packets**

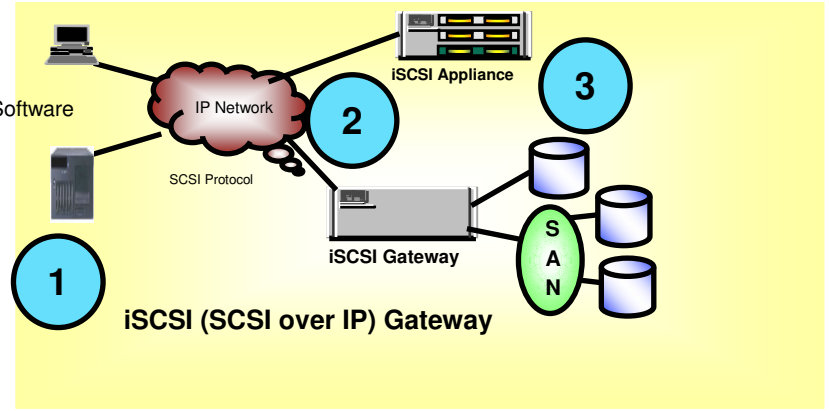

Two industry approaches--

<u>iSCSI appliances</u> (with embedded storage)
<u>Gateways</u> (IP/Fibre Channel bridges)

1. iClient (initiator) code reroutes SCSI commands over IP network
2. iSCSI target code receives SCSI commands from IP network
3. SCSI commands then either routed directly to embedded storage (iSCSI appliance) or routed to FC SAN (iSCSI gateway)

Why Use Native FICON Channels?

- FICON is more than just an I/O card or a channel type
  - FICON is at the heart of some very important solutions
- Performance
  - Faster backup/ recovery times, shorter batch windows, faster data access for large data queries
  - See the FICON and FICON Express Channel Performance white paper (February 2002) located at **ibm.com**/servers/eserver/zseries/connectivity/
- Distance
  - Extended distance and remote vaulting (up to 100 km)
  - Reduced data rate performance droop at extended distances
- Channel consolidation
  - Fewer channels, CHPIDs, ports, fiber optic cabling
- FICON support of Cascaded Directors
  - Integrity features integral to FICON architecture
  - 2 Gbps Inter Switch Link capability

FICON Support of Cascaded Directors – Simplify Cross-Site Connectivity – Reduce Costs

- **Dynamic – channels, ports, fiber shared**
- **Fewer cross site connections - Repeaters, DWDM, Fibers, Channels, Director Ports**
- **Reduce implementation cost for disaster recovery applications; GDPS and Remote Copy**
- **Support: z/OS V1.3 and V1.4 plus PTFs**
- **Requires single-vendor high integrity fabric**

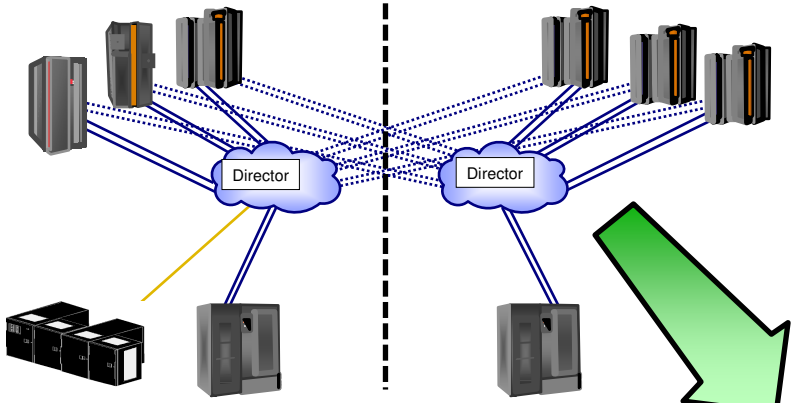

Two site non-cascaded director topology
Each CEC connects to directors in both sites

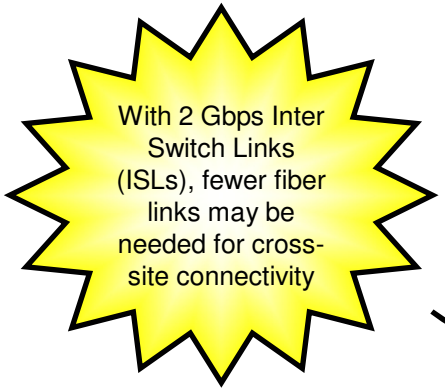

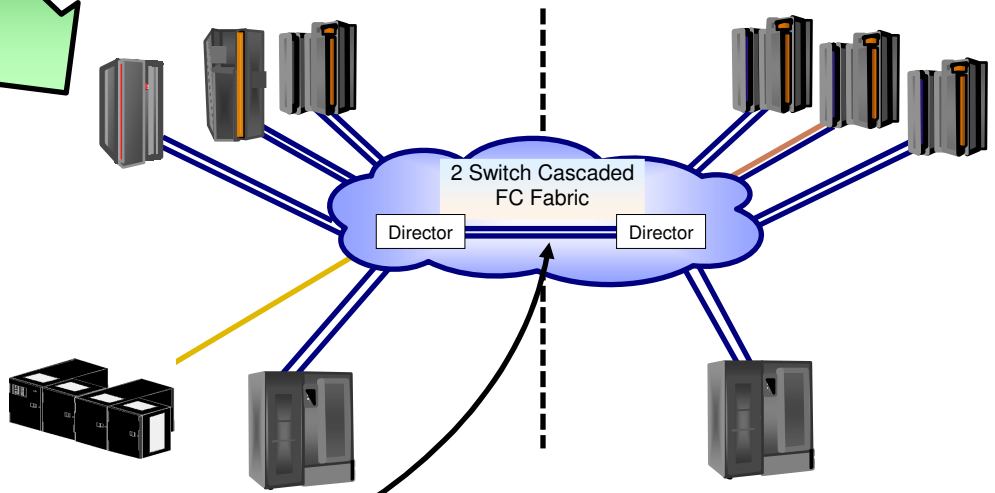

Two Site cascaded director topology
Each CEC connects to local directors only

FICON and FCP Intermix Considerations

- Reasons for keeping them in separate physical SANs:
  - If both environments are large enough to justify self-contained SANs anyway, it's Less Complex to just keep them separate.
  - Adding in FICON imposes design considerations on the SAN that would otherwise not apply
    - e.g. need for SANtegrity if cascaded
    - e.g. 1 hop ISL limit for FICON
    - e.g. Product selection decisions (eg FICON VTS Peer to Peer does not support all switch products)
  - Feature code costs: SANtegrity / FICON feature costs for a large consolidated switch/director may work out more expensive than using small switches with the s/w features only applied to those switches that need them. Consolidating is therefore not by default cheaper (at least for Feature Codes)
  - Expediency: It's the path of least resistance
    - IT Services' "Managed Storage Service" offering is familiar with managing shared Fibre-Channel-only SANs
    - In today's IT Services Service Delivery support model, Mainframe and Midrange environments are usually run by separate teams
    - The environment may have evolved with different SAN vendors chosen for FC and FICON SANs. If this is the case, standardizing to one vendor will be a big change for some teams
  - Different Connectivity standards: FICON is almost always Longwave; FC is almost always Shortwave.
  - Possible Conflicting code level requirements.
    - e.g.. A new level required for FICON may require new drivers on AIX.
    - e.g. A problem in AIX with pathing requires a newer version of switch microcode that is not yet supported by VTS.
  - There is a fear (unjustified) that mixing FC and FICON like this will allow some traffic to negatively impact other traffic

FICON and FCP Intermix Considerations

- Reasons for keeping them together in same physical SANs:
  - For smaller requirements, where a mix of FC and FICON is required, it is more cost-effective to share the SAN Fabric. Fewer boxes will need managing, so less effort (less cost) is required overall. The equipment will be from one vendor, again simplifying management.
  - Where physical Machine room space is a constrained resource, it should be more space-efficient to share the SAN Fabric
  - It allows better pooling of SAN resources.. spare ports can be pooled for either FC or FICON growth (as long as shortwave vs longwave connectivity standards are set)
  - One "SAN team" can easily manage a shared FC and FICON SAN (FICON is address-centric, definition oriented and control is set at the zSeries server level by the System Programmers, so little "additional" FICON work for the SAN team
  - For smaller requirements, consolidating workload may justify "director class" switches with higher availability characteristics
  - Use of the latest SAN Virtualization features may remove some previous inhibitors

Types of Fibre Channel Ports

- E_Port - Expansion Port - a port on a switch used to link multiple switches together into a Fibre Channel Switch Fabric; May also be used to bridge a SAN into a ATM Frame Relay Wide Area Network (WAN)
- F_Node - a Fabric Attached Node; e.g. a server or a storage control unit
- F_Port - Fabric Port - a port used to attach a Node Port (N_Port) to a switch fabric (switch side of F_Node)
- FL_Port - The access point of the fabric for physically connecting the user's Node Loop port (NL_Port); e.g. FCAL FCSS
- G_Port - Generic Port - a generic switch port that functions as either a E_Port or F_Port.
- L_Port - Loop Port - a node or fabric port capable of performing FCAL functions and protocols. NL_Ports (node side) and FL_Ports (switch side) are loop capable ports.
- N_Port - Host attachment side of a fabric port (F_Port) attached to a switch
- NL_Port - Node Loop Port - a node port that supports FC-AL devices e.g. FC-AL port on a Storage Control Unit
- Private NL_Port - An NL_Port which does not attempt login with the fabric and only communicates with other NL_Ports on the same loop (e.g. HP uses this technology)
- Public NL_Port - An NL_Port that attempts login with the fabric and can observe the rules of either public or private loop behavior. A Public_NL Port may communicate with both both private and public NL ports (Dell Servers use this technology; Dell Storage uses private loops)
- U_Port - Universal Port - allows the attachment/function of any other industry standard port technology

Storage Area Network "Fabric"

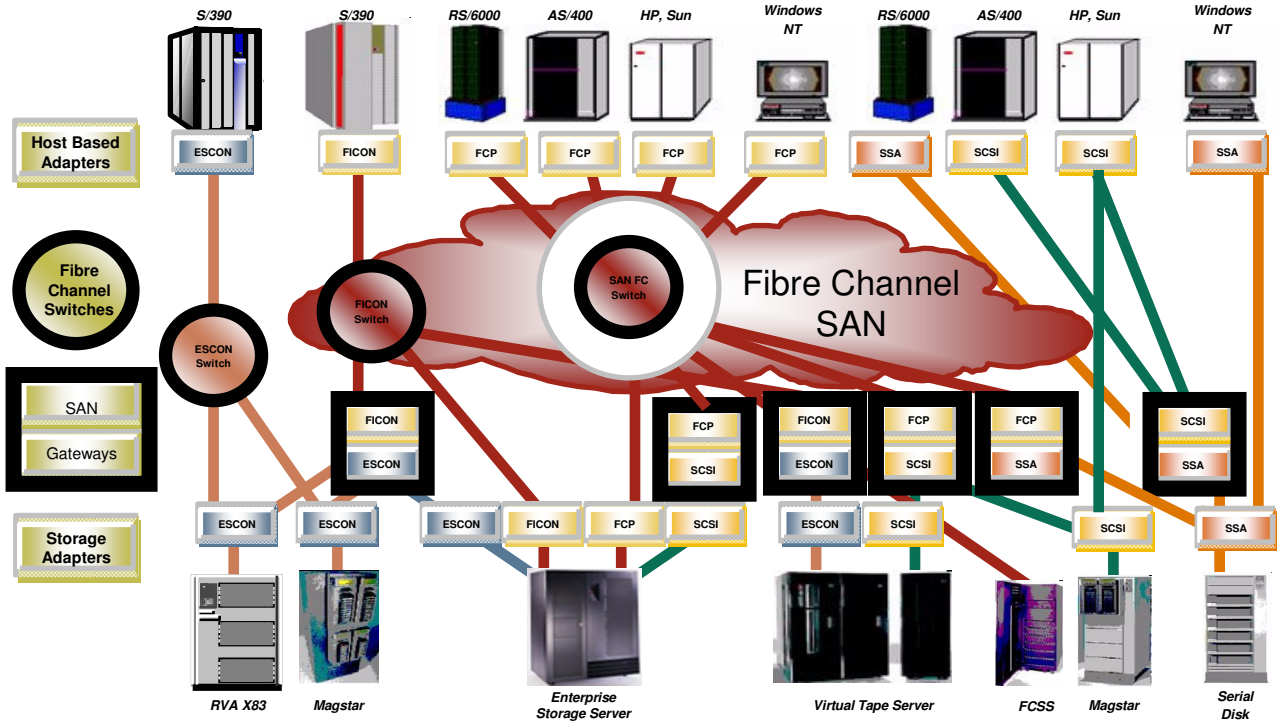

Two Major Manufacturers in IBM's Portfolio

- **Brocade (with McDATA) has largest market share**
  - Both Open and FICON
  - Also channel extension products used for FICON and mainframe XRC over long distances
  - Used in mission-critical applications such as TPF*

- **Cisco has comprehensive portfolio**
  - Emphasize synergy with corporate IP network
    - SAN fabric management uses familiar network interface
  - VSAN is most versatile / flexible

*TPF = Transaction Processing Facility (airline reservations)

IBM Entry Switch Portfolio

(2498-B24 / 249824E)

IBM System Storage SAN24B-4 Express
- 8-16-24 ports-on-demand; 8, 4, 2, 1 Gbps
- Shortwave, longwave, extended distance SFPs
- Full Fabric, Extended Distance and other options
- See SAN 24-B Express Hotlink*

(2053-424 / 241724C)

Cisco MDS 9124 Express for IBM System Storage
- 8-16-24 ports; 4, 2, 1 Gbps
- Shortwave, longwave SFPs
- Full member of Cisco SAN fabric
- See Cisco MDS 9124 Express Hotlink*

*** Note: Hotlinks work when slide is viewed in "Slide Show" mode**

IBM Midrange Switch Portfolio

(2498-B80)

IBM System Storage SAN80B-4
- 48-64-80 ports-on-demand; **8**, 4, 2, 1 Gbps
- Shortwave, longwave, extended distance SFPs
- Extended Distance and other options
- See SAN 80B-4 Hotlink*

(2498-B40 / 249840E)

IBM System Storage SAN40B-4
- 24-32-40 ports-on-demand; **8**, 4, 2, 1 Gbps
- Shortwave, longwave, extended distance SFPs
- Full Fabric, Extended Distance and other options
- See SAN 40B-4 Hotlink*

(2053-434 / 2053-S34)

Cisco MDS 9134 for IBM System Storage (2053-434/S34)
- 24-32 ports; 4, 2, 1 Gbps, also 2 10 Gbps ISL links
- Shortwave, longwave SFPs
- Can be stacked to create a 64-port switch
- See Cisco MDS 9134 Hotlink*

*** Note: Hotlinks work when slide is viewed in "Slide Show" mode**

IBM Enterprise Director Portfolio – b-type

IBM System Storage SAN768B (2499-384)

- Extremely high performance
- Up to 768 end-user ports in two chassis connected via Inter-Chassis Links
- 16, 32 and 48-port 8 Gbps switch blades; sixteen FC + two GbE ports MPR (Multi-protocol Router) blade; 10 Gbps FC blade
- Eight blades per chassis
- 8 & 4 Gbps SW, 4 Gbps LW and Extended Distance LW SFPs
- FICON w/CUP, Extended Distance and other options
- See SAN 768B Hotlink*

IBM TotalStorage SAN256B (2109-M48)

- 16, 32 and 48-port 4 Gbps switch blades; sixteen FC + two GbE ports MPR (Multi-protocol Router) blade; iSCSI blade; 10 Gbps FC blade
- Eight blades per chassis
- Shortwave, longwave, extended distance SFPs
- FICON w/CUP, Extended Distance and other options
- See SAN 256B Hotlink*

*** Note: Hotlinks work when slide is viewed in "Slide Show" mode**

IBM Enterprise Director Portfolio – Cisco

	**Cisco MDS 9506 for IBM System Storage (2054-E04*/2062-D04 WFM)** Blades: 12-24-48 port FC switching with 8/4/2/1 Gbps (192 ports maximum); 18/4 Multiservice Module (18+4); 32-port Storage Services; 8-port IP; 4-port 10 Gbps FC switching for ISLs and distance extension Shortwave, longwave, CWDM extended distance, GbE SFPs FICON w/CUP, Extended Distance and other options See Cisco MDS 9506 Hotlink* with IBM warranty service
	**Cisco MDS 9509 for IBM System Storage (2054-E07*/2026-D07 WFM)** Blades: 12-24-48 port FC switching with 8/4/2/1 Gbps (336 ports maximum); 18/4 Multiservice Module (18+4); 32-port Storage Services; 8-port IP; 4-port 10 Gbps FC switching for ISLs and distance extension Shortwave, longwave, CWDM extended distance, GbE SFPs FICON w/CUP, Extended Distance and other options See Cisco MDS 9509 Hotlink* with IBM warranty service
	**Cisco MDS 9513 for IBM System Storage (2054-E11*/2062-E11 WFM)** Blades: 12-24-48 port FC switching with 8/4/2/1 Gbps (528 ports maximum); 18/4 Multiservice Module (18+4); 32-port Storage Services; 8-port IP; 4-port 10 Gbps FC switching for ISLs and distance extension Shortwave, longwave, CWDM extended distance, GbE SFPs FICON w/CUP, Extended Distance and other options See Cisco MDS 9513 Hotlink* with IBM warranty service

*** Note: Hotlinks work when slide is viewed in "Slide Show" mode**

IBM Router Portfolio – b-type

(2005-R04)

IBM System Storage SAN04B-R (2005-R04)

- Two 4, 2, 1 Gbps FC plus 2x50 MbE ports (optical or copper)
- Performance Enhancement feature enables 14 additional FC ports and increases IP link speed to GbE to make functionally equivalent to SAN18B-R
- FICON Accelerator and FICON w/CUP optional
- See SAN04B-R Hotlink*

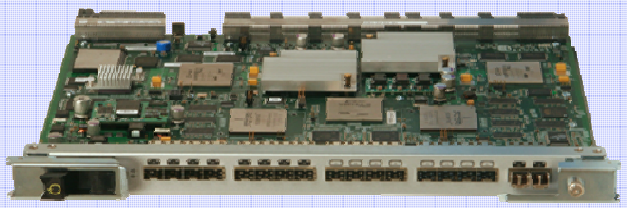

M48 Routing Blade for SAN256B and SAN768B

- Sixteen 4, 2, 1 Gbps FC plus 2 GbE ports (optical or copper)
- FC Routing + FCIP (optional)
- Shortwave, longwave SFPs
- See SAN256B Hotlink*

*** Note: Hotlinks work when slide is viewed in "Slide Show" mode**

IBM Switch and Blade Portfolio – Cisco

(2054-E01)

Cisco MDS 9222i for IBM System Storage (2054-E01)

- 18 FC ports + 4 GbE ports; 4, 2, 1 Gbps; 3U height
- Shortwave, longwave SFPs
- Slot for one additional blade, including 12, 24 or 48-port 4 and 8 Gbps switch blades or 14/2 Multiprotocol Services Module
- See Cisco MDS 9222i Hotlink*

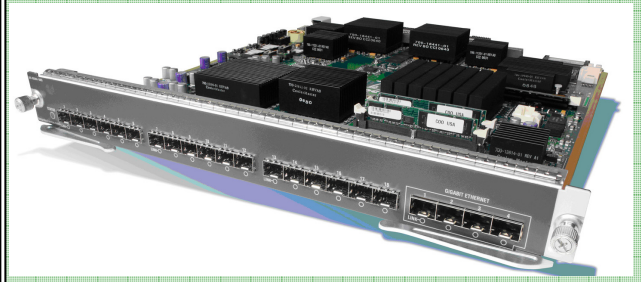

Cisco MDS 9000 Family 18/4 Multiservice Module (Feature Code 2250 for Cisco MDS 9500 directors)

- Eighteen 4, 2 or 1 Gbps FC ports + four GbE ports
- Shortwave, longwave, CWDM extended longwave SFPs
- Optional FCIP
- See Cisco 18/4 Multiservice Module Hotlink*

*** Note: Hotlinks work when slide is viewed in "Slide Show" mode**

Converged Networks: FCoE and FCoCEE Overview

Converged Networks

- **Fibre Channel over Ethernet (FCoE) is designed to provide a way of mapping the Fibre Channel Protocol (FCP) over a Converged Enhanced Ethernet network.**

- **FCoE allows converging a server's Ethernet and Fibre Channel adapters into a single Converged Network Adapter (CNA).**
  - This enables a cost savings by avoiding the use of two adapters (one for Fibre Channel (FC) and another for Ethernet) inside the data center (less than 1 km)
- **As FCoE matures, clients can also converge their Ethernet and Fibre Channel networks.**

- **Additionally, by keeping Fibre Channel's upper layers the same, FCoE preserves the investments made in FC infrastructure (e.g. management and operating system stacks).**

Storage Connectivity Models

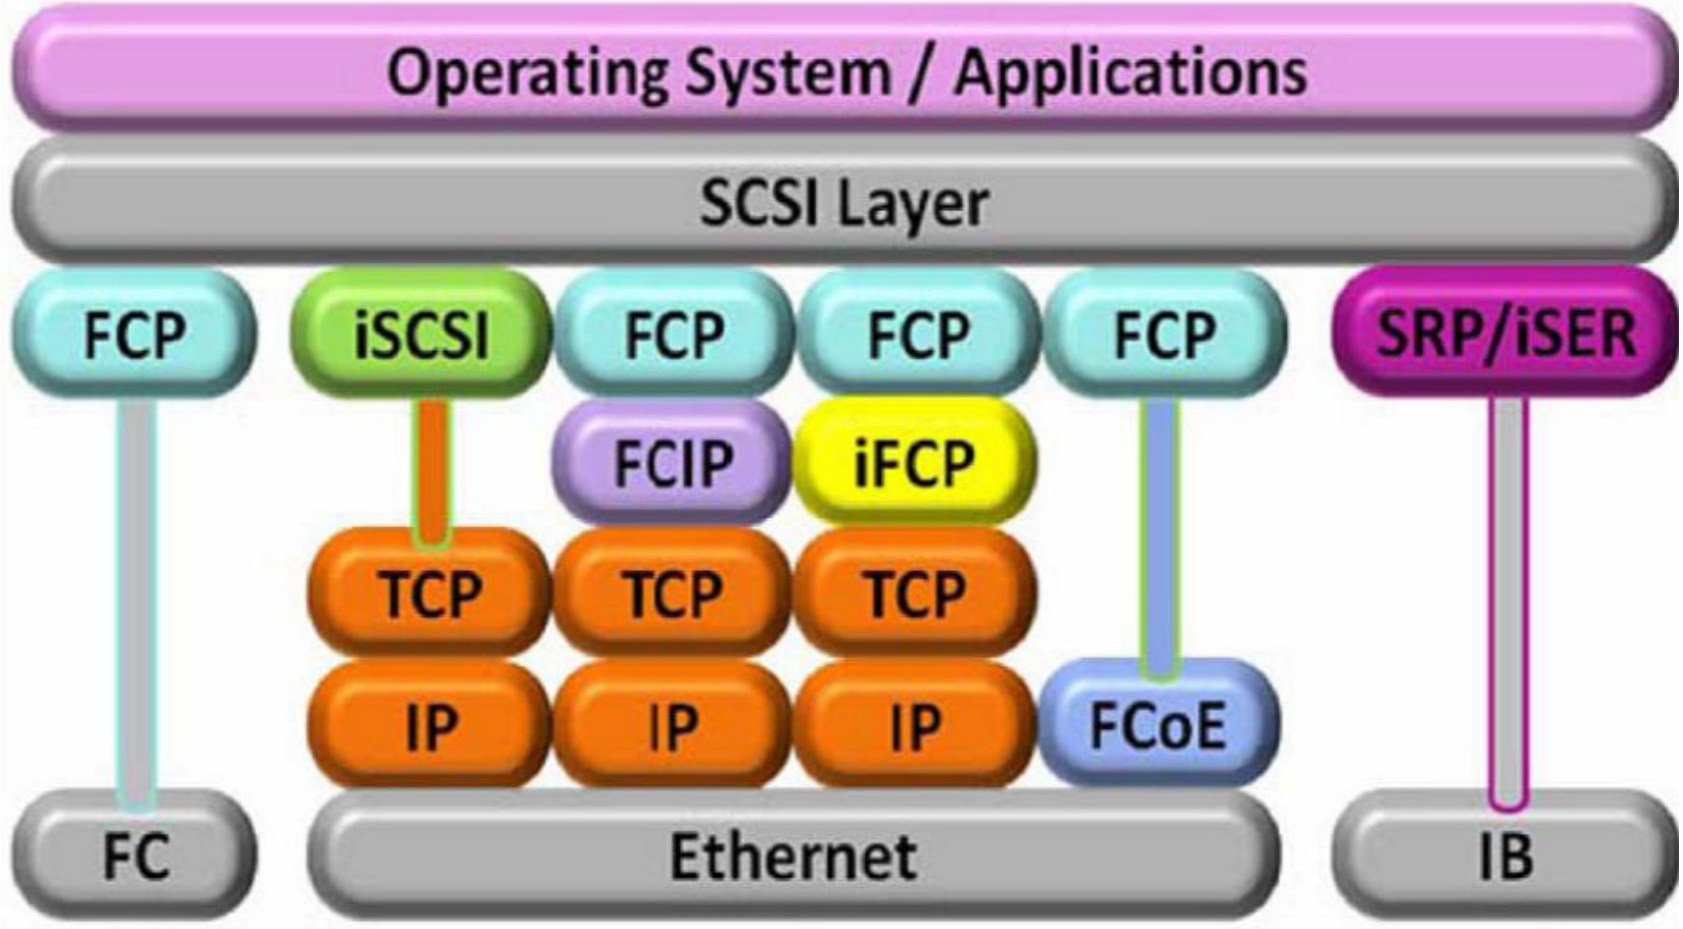

Fibre Channel over Ethernet (FCoE) and Fibre Channel over Convergence Enhanced Ethernet (FCoCEE) Part 1

- **FCoE is the standard that is driving convergence and the emergence of FCP over the Ethernet.**
- **FCoE is encapsulates FCP frames over the Ethernet.**

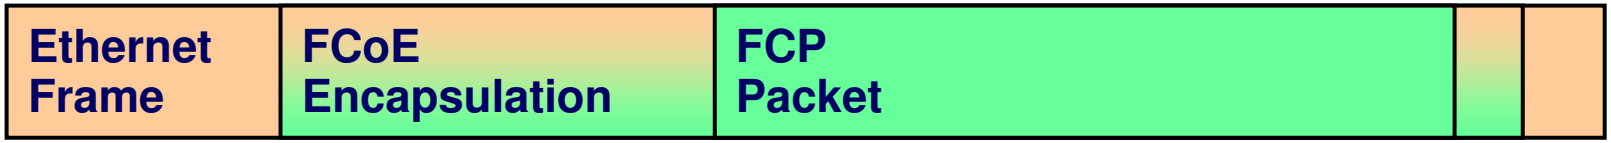

- **Ethernet provides the physical interface.**
- **However, Ethernet needs enhancements, to:**
  - Provide no-drop behavior in face of fabric congestion.
  - Manage traffic interferences.

Fibre Channel over Ethernet (FCoE) and Fibre Channel over Convergence Enhanced Ethernet (FCoCEE) Part 2

The INCITS T11 technical committee that started the impetus for the creation of the FC-BB-5 standard – FCOE:

- "This project proposal recommends the development of a set of additional and enhanced mechanisms, services, and protocols to connect Fibre Channel entities over selected non-Fibre Channel protocol infrastructures. Enhancements to Ethernet protocols, such as the Pause mechanism defined in IEEE 802.3-2005, make it possible to define a direct mapping of Fibre Channel over Ethernet (FCoE). This mapping provides several technological benefits over the currently defined Fibre Channel over IP (FCIP) mapping and gives a significant business advantage to Fibre Channel over competing technologies, such as iSCSI, because Fibre Channel provides seamless compatibility with existing storage, drivers, and management tools. The FCoE mapping allows Fibre Channel to be used in Ethernet-based I/O consolidated environments and will be especially useful in both the Data Center and Metro Ethernet environments.

Fibre Channel Frame Encapsulation into Ethernet

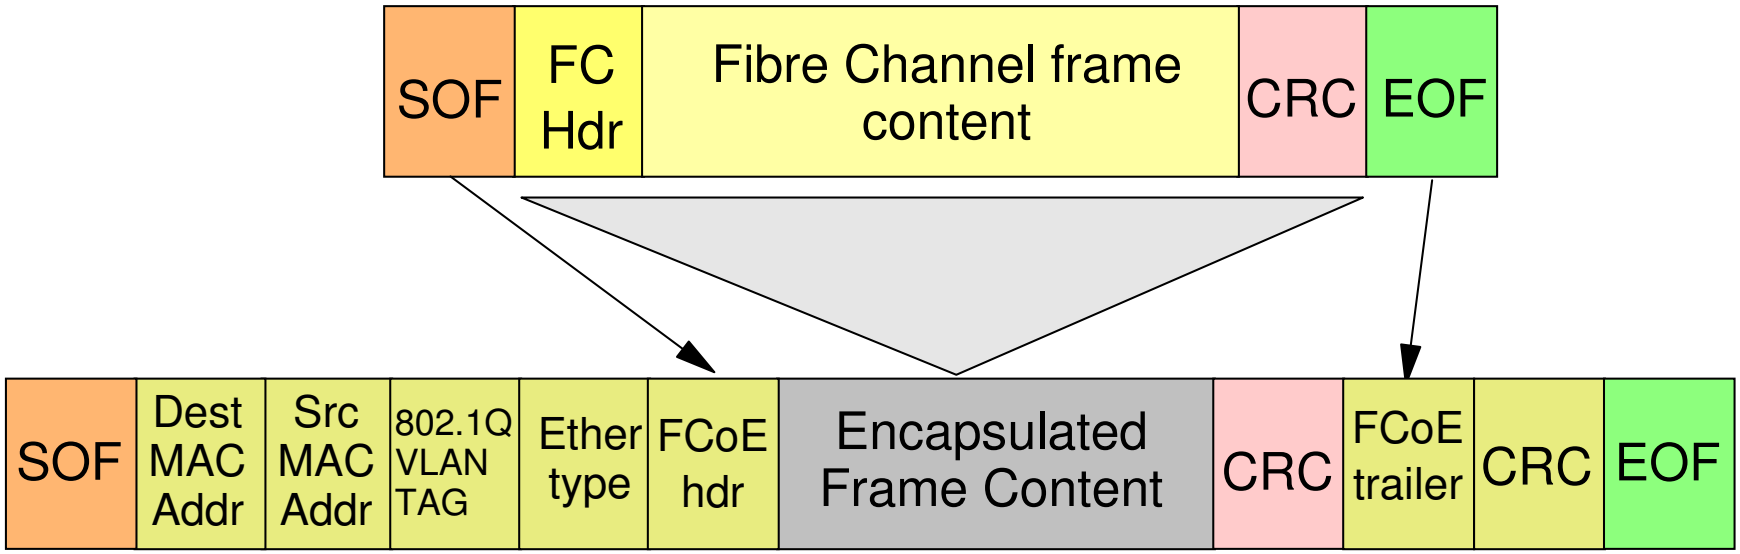

Standard Ethernet Frame = 1518 Bytes Maximum

Frame size considerations

- **Please note that the standard Ethernet frame is a maximum of 1518 bytes as compared with the maximum FCP frame size of 2148 bytes.**

- **As you can see, if we were to simply wrap the frames into an Ethernet frame, some form of fragmentation, or segmentation, must occur due to the difference in maximum sizes (1518 compared to 2148). Although this segmentation is possible, it is inefficient.**

- **The answer to this difference in frame sizes is to increase the size of the Ethernet frame.**
  - Conventionally, jumbo frames can carry up to 9000 bytes of payload but this is not standard. Many, but not all, Gigabit Ethernet switches and Gigabit Ethernet network interface cards support jumbo frames, but all Fast Ethernet switches and Fast Ethernet network interface cards support only standard-sized frames.
  - Although jumbo frames are not a standard, the quickest and simplest option was to require jumbo frame support for every device in the FCoE/FCoCEE network.
  - “Baby” jumbo frames of approximately 2500 bytes are desirable for the future.

Fibre Channel over Ethernet (FCoE) and Fibre Channel over Convergence Enhanced Ethernet (FCoCEE) Part 3

- **FCoE is the standard that is driving convergence and the emergence of FC over the Ethernet.**
- **However, in and of itself it (FCoE) will not be enough to allow for fabric convergence. Without any enhancements Ethernet itself does not meet the requirements for data center convergence (FCP networks are lossless. Ethernet networks are not lossless.)**
- **FCoCEE is the new base transport that starts to support FCoE - Enhanced Ethernet will include new extensions to the existing Ethernet standard that will eliminate the lossy nature of the Ethernet and make it a viable storage networking transport.**
- **With the emergence of 10 Gbps Ethernet (10 GbE), we now have a base on which all FCoE/FCoCEE solutions will be built. And it only gets better, and faster, with 40 GbE and 100 GbE in plan for the future.**

What is Convergence Enhanced Ethernet?

Today's Dual Fabrics (Fibre Channel and Ethernet)

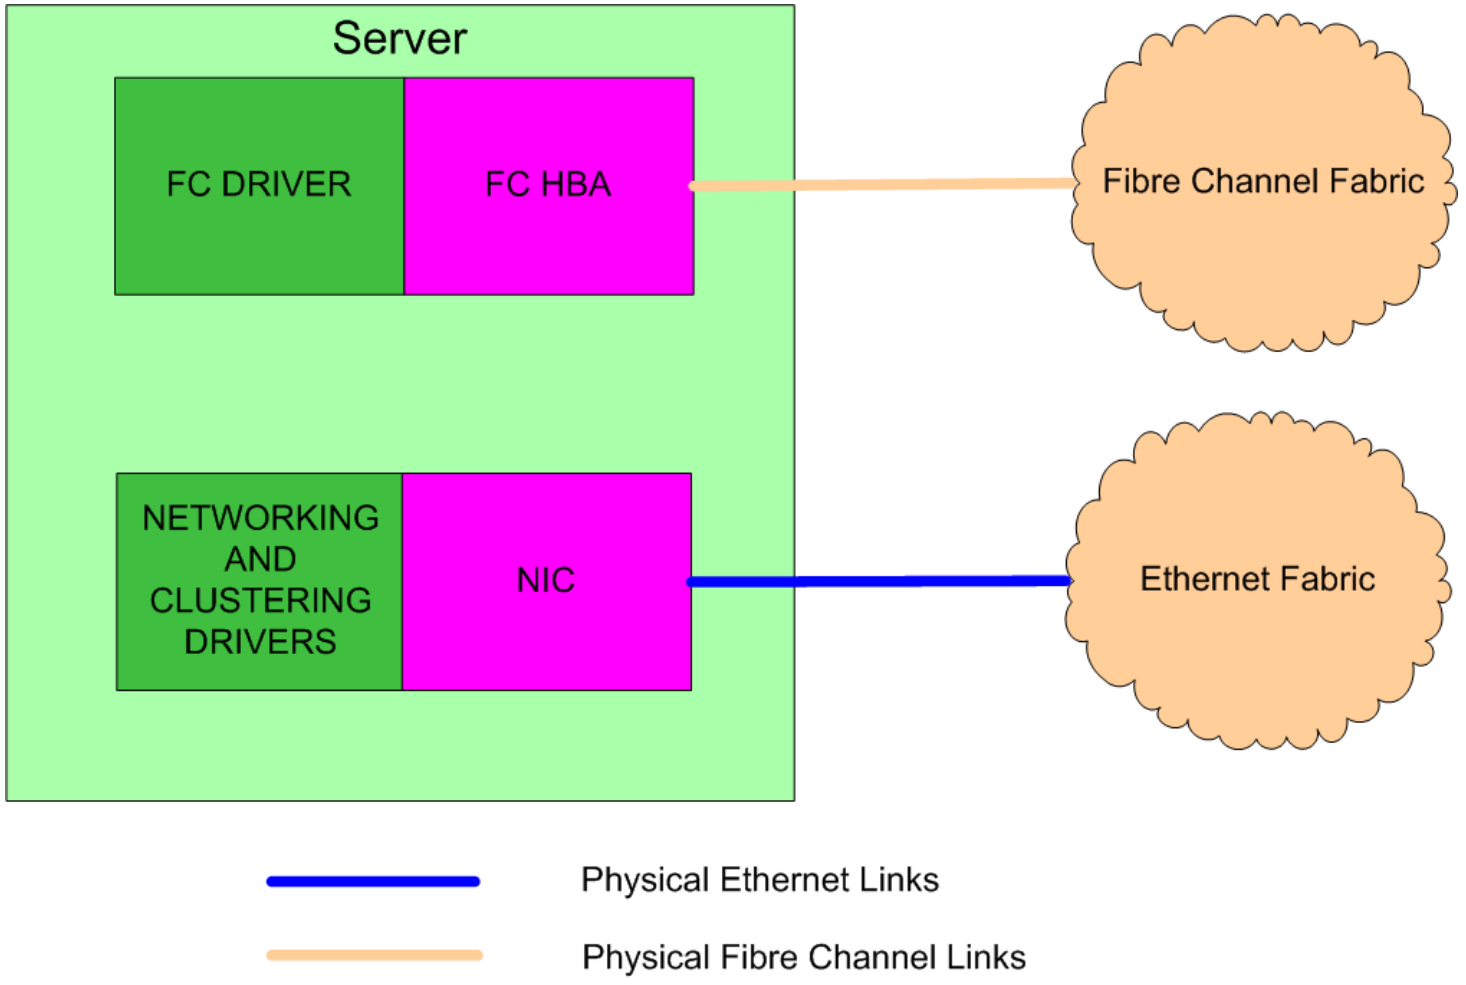

Converged Networks with Enhanced Ethernet

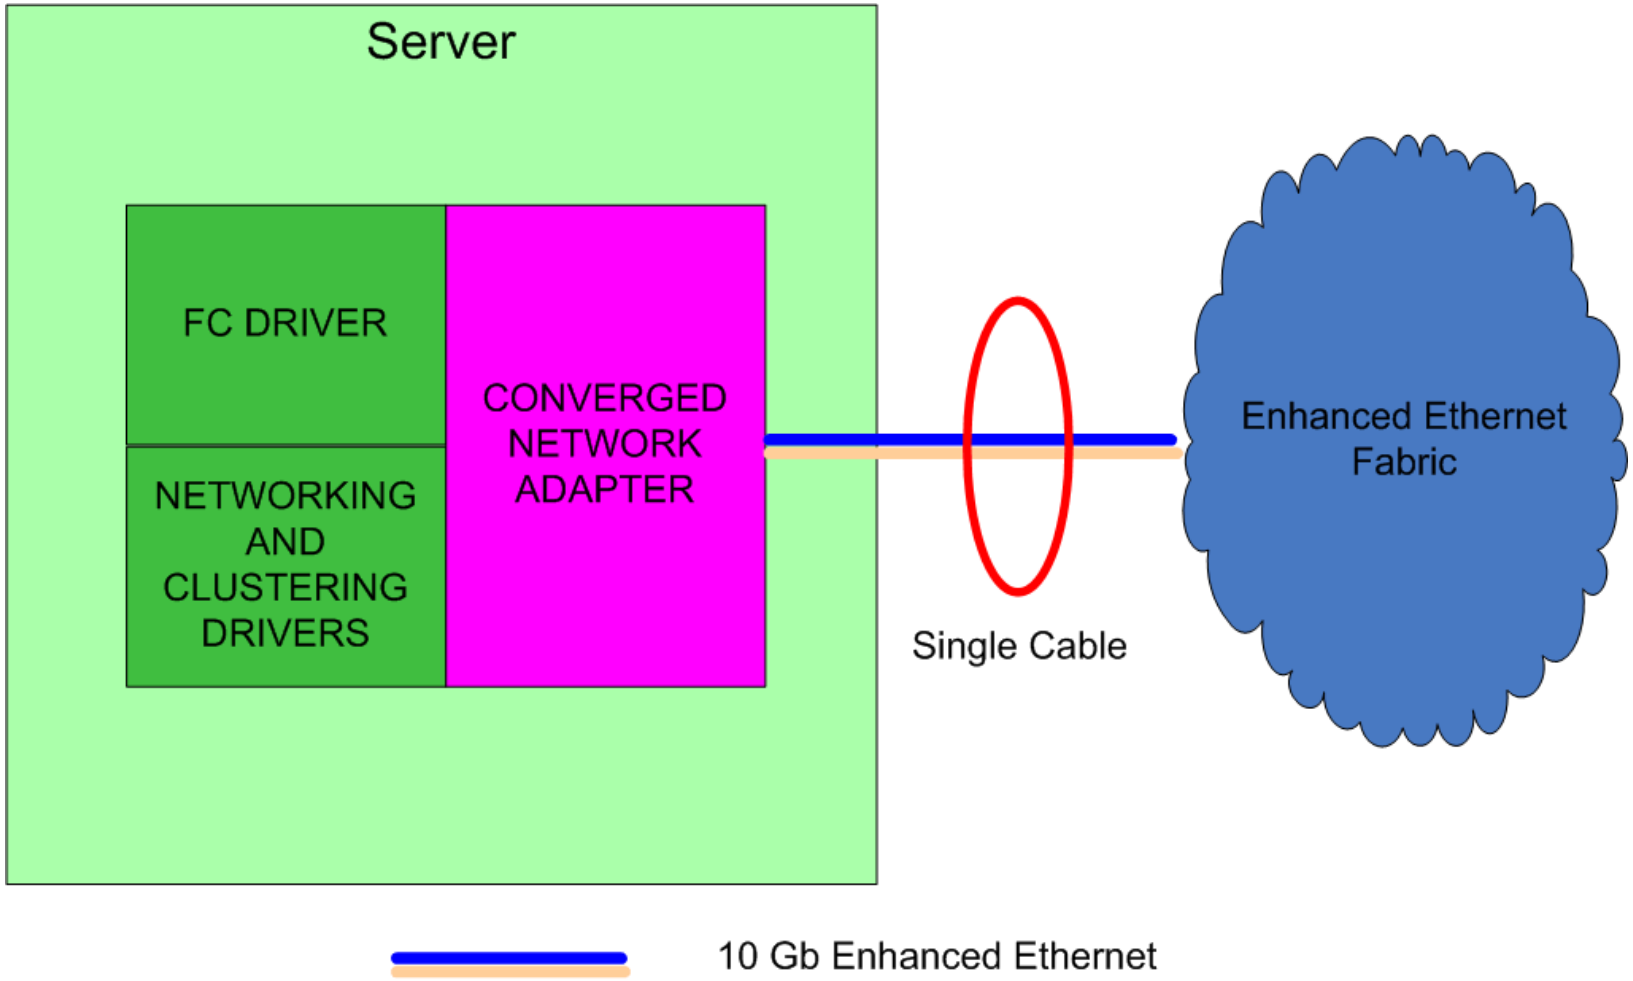

FC vs FCoCEE Protocol Stack Comparison

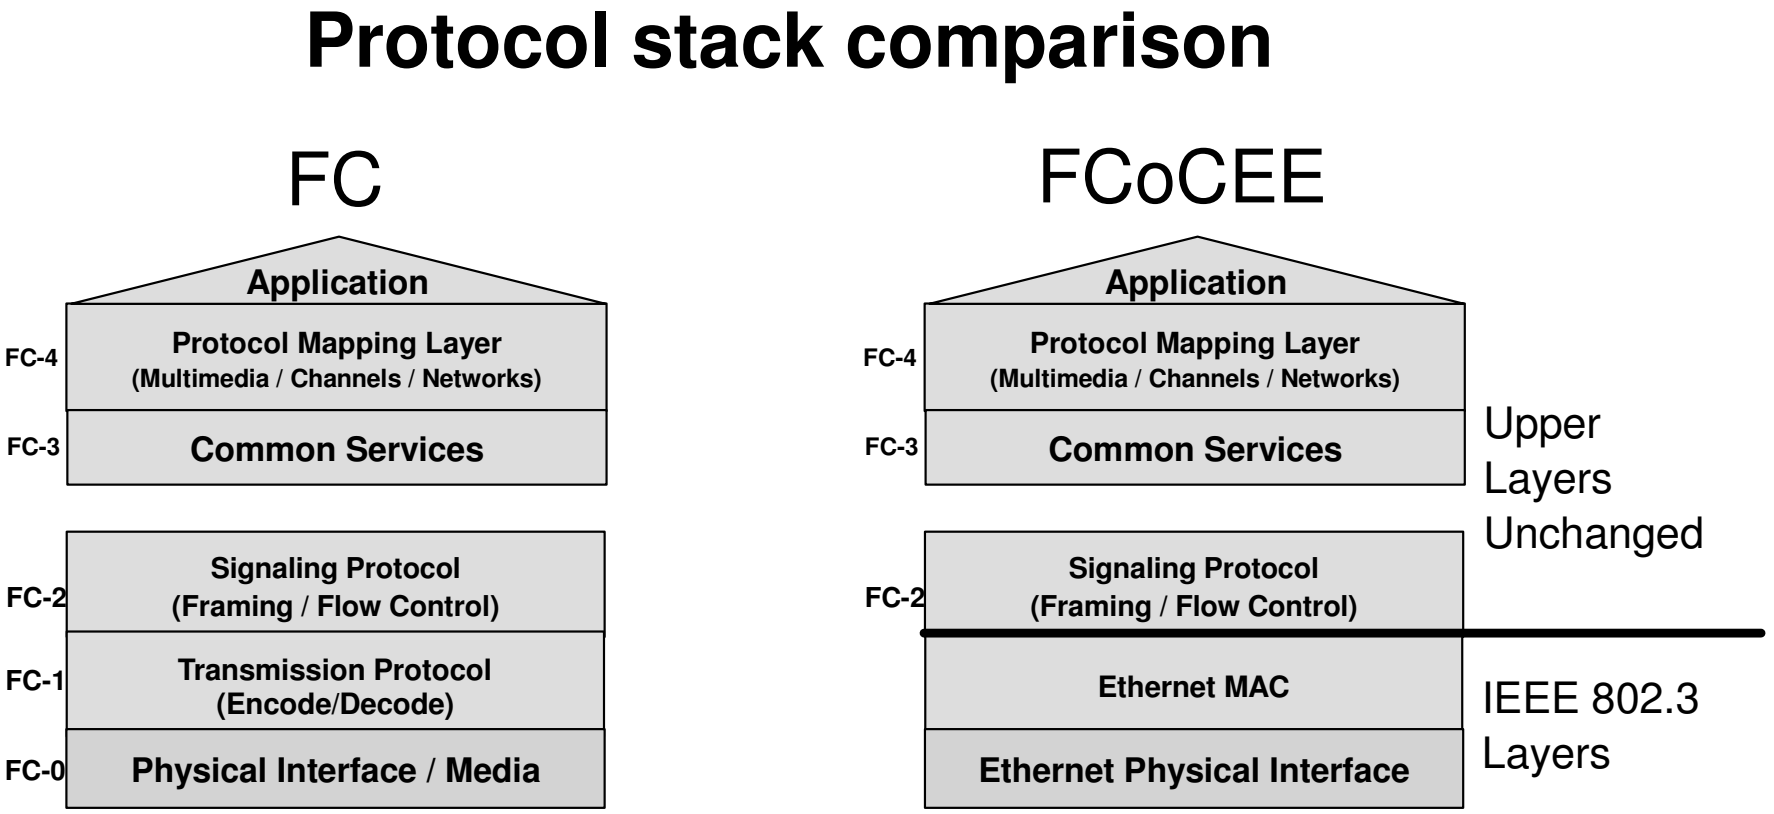

Addressing

- **In an FCP network, the links are based on a "point-to-point" topology. The Ethernet network does not create a point-to-point connection in the same way that FCP does. FCoE will use the destination and source MAC addresses to forward a frame to its intended destination.**

- **Addressing Schemes**
  - Server-Provided MAC Addresses
    - As its name suggests, an SPMA is a MAC address that is issued in accordance with Ethernet standards and set by the manufacturer at installation.
  - Fabric-Provided MAC Addresses (WWPN)
    - The FPMA is a fabric-unique address that is assigned by the fabric. The low-order 24 bits are equivalent to the N_Port ID (FC-ID) assigned during fabric login, and the high-order 24 bits are equal to the FCoE MAC address prefix (FC-MAP) associated with the fabric.

Converged Ethernet Standardization

- **Data Center Bridging (DCB) refers to a set of Ethernet enhancements currently being pursued in IEEE 802.1**
  - Per-priority Flow Control
  - Enhanced Transmission Selection
  - Discovery and Capability Exchange
  - Congestion Notification

- **CEE is the name folks (EMC, HP, IBM and trade press) gave to the "version 0" specifications the CEE Authors proposed to the IEEE 802.1 DCB workgroups, for the following 3 functions:**
  - Per-priority Flow Control
  - Enhanced Transmission Selection
  - Discovery and Capability Exchange

- **CEE Authors was formed by IBM to facilitate/accelerate definition of the above.**

- **Note: the version 0 specifications serve as a de-facto standard for initial deployments.**

Storage Convergence Protocols

- **FCoE/FCoEE is an enhancement that expands FCP into the Ethernet by combining two leading-edge technologies (FCP and Ethernet)**
  - We see three deployment usage cases for the initial generation of FCoCEE:
    - Usage case 1 is a rack upgrade scenario. In this case, an FCoCEE rack is deployed into an existing data center (DC), without changing the data center's Ethernet or FC infrastructure.
    - Usage case 2 is a new Dual-Fabric, data center scenario, where FCoCEE is used within each rack, but at the DC level, there are still two separate fabrics: Ethernet and FC.
    - Usage case 3 is a new Converged Fabric, data center scenario, where FC is used at the perimeter of the converged FCoCEE fabric to attach storage, but CEE and FCoCEE switches are used throughout the DC.
  - All of these usage cases will also enable a lower operational expense by integrating FC and Ethernet fabric management.
- **FCoE/FCoEE could supplant iSCSI depending on adoption and cost factors**
- **FCoE/FCoEE will not supplant FCIP/iFCP at this time**
- **While Infiniband is used in many server interconnect applications – it has not caught on in the industry as a Storage connectivity protocol.**

Potential FCOE/FCoCEE rollout

Marty Lans, Brocade's senior director of product marketing, believes that network convergence is a multiyear process consisting baby steps:

1. **General availability of FCoE fabric and host products**
2. **I/O consolidation using FCoE top-of-rack switches on the edge of the data center**
3. **General availability of FCoE storage**
4. **Corporations smooth over storage vs. networking IT politics**
5. **Actual total cost of ownership savings are defined by early adopters**
6. **Build out the CEE/FCoE fabric in the core of the data center**
7. **SSDs drive the bottleneck back to the network**
8. **40/100-Gb/s CEE/FCoE networks**

Source – http://www.byteandswitch.com/blog.asp?blog_sectionid=750&doc_id=176057&WT.svl=blogger1_2

References

- **An Introduction to Fibre Channel over Ethernet, and Fibre Channel over Convergence Enhanced Ethernet – Jon Tate**
  (http://www.redbooks.ibm.com/abstracts/redp4493.html?Open **)**
- **INCITS – Technical Committee T11, which is the committee responsible for Fibre Channel Interfaces – (** http://www.t11.org/index.html **)**
- **Fibre Channel over Ethernet in the Data Center – An Introduction**
  (http://www.fibrechannel.org/OVERVIEW/FCIA_SNW_FCoE_WP_Final.pdf **)**
- **FCoE development tracker – (** http://www.fcoe.com/ **)**
- **FCoE standard – (** http://www.t11.org/ftp/t11/pub/fc/bb-5/09-056v5.pdf **)**
- **FCoE Wikipedia Entry – (** http://en.wikipedia.org/wiki/FCoE **)**
- **CEE Authors** http://tech.groups.yahoo.com/group/cee-authors/
  - IEEE 802.1 Data Center Bridging www.ieee802.org/1/pages/dcbridges.html
  - Congestion Notification www.ieee802.org/1/pages/802.1au.html
  - Enhanced Transmission Selection www.ieee802.org/1/pages/802.1az.html

Storage Evolution

Technology Shifting Towards the Network and Software

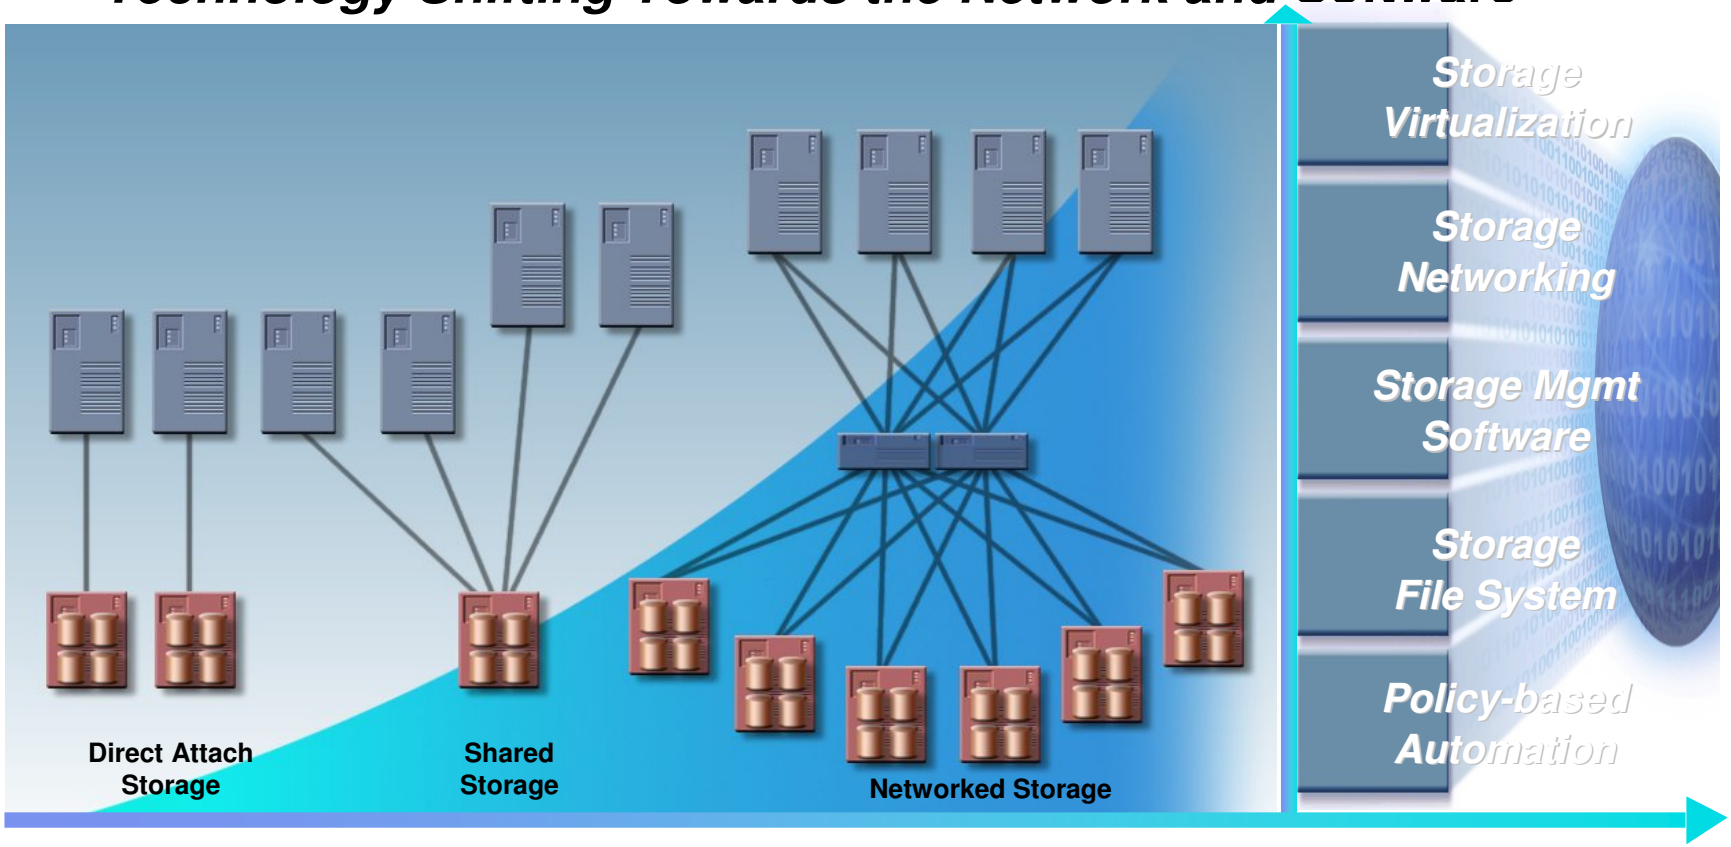

Value in Network-based Storage Infrastructure Software

Storage Virtualization is . . .

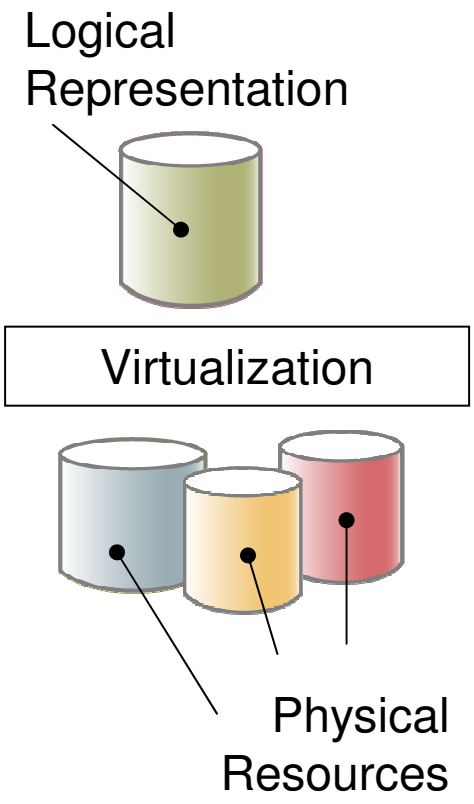

Technology that makes one set of resources look and feel like another set of resources, preferably with more desirable characteristics…

A logical representation of resources not constrained by physical limitations

- Hides some of the complexity
- Adds or integrates new function with existing services
- Can be nested or applied to multiple layers of a system

IT Historical Precedent

- *Operating Systems*
  - MVS, Windows, UNIX ...
- *Mainframe*
  - Virtual memory, Virtual Channels ...
- *Logical Volume Managers*
  - Veritas, Tivoli ...
- *Tape technology*
  - Virtual Tape Server ...
- *Disk systems*
  - Cache, RAID, Iceberg, 'Shark' ...

Hardware/Software Virtualization

- Symmetric (in-band)
  - Virtualization Appliances
  - Data/control stream through appliance
  - Favored by IBM, DataCore ...

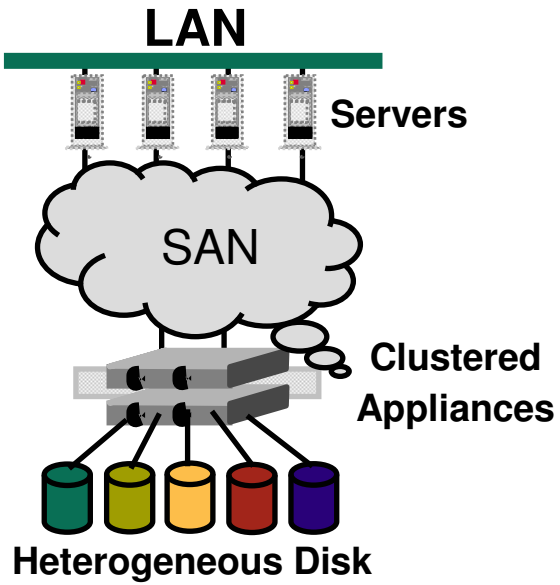

- Asymmetric (out-of-band)
  - Metadata Servers (clustered)
  - Server data stream, Metadata control stream through appliance
  - Favored by IBM, Compaq ...

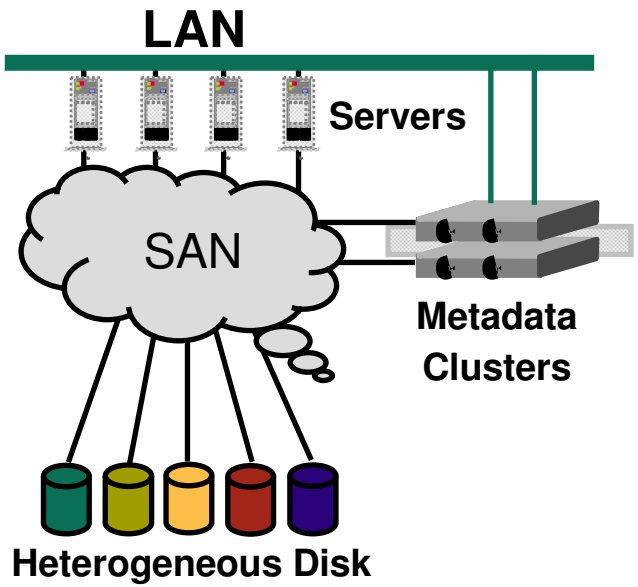

IBM Storage Management within the Infrastructure

Storage Virtualization

- *Common Class of Service* across uncom vices

Netcool Omnibus
TADDM
ITM

ers

Data Management

rs, including heterogeneous environments

Storage Infrastructure Analytics

- *Storage management* with a single managem t layer
- Storage configuration management across t work
- Centralizes management of all storage devi
- Disk and SAN topology views and manage
- Host file systems, data bases, and storage
- Wide range of IBM and non-IBM device su

ver and
elopment,

r resources
roves

ger,
IIA (4Q)

Tape

eve

- *Au* storage technology
- *Automatic file expiration* fo
- *Disk Pooling* to effectively enterprise tape
- *Information-on-Demand aw* Records Management, Wareho Management functions

Application Integration

- *Data offload* for performance, scaling, operational simplification and compliance
- *Archive aware* functions in DB2 Commonstore, Content Manager, Filenet, Optim and key ISVs
- *Continuous data protection*, local and remote, through TSM Fastback

Offsite,
ecure

VTL
icaton

Tivoli Provisioning Manager

SAN Volume Controller

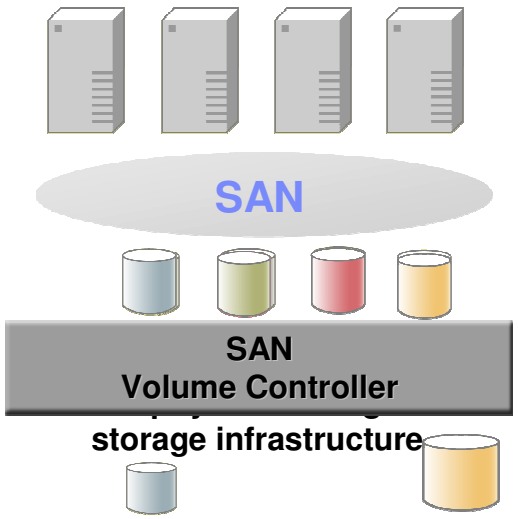

Virtual disks, however, can remain constant while physical changes in the infrastructure are carried out.

Value

- Centralized point of control for volume management
- Reduce or eliminate downtime for planned outages, maintenance and backup
- Improved resource utilization
- Single, cost effective set of advanced copy services

Functional Summary

- Single storage pool grouped into disk groups
- Dynamic Data Migration
- Image Mode to transition from existing SANs
- Copy Services with consistency groups
- SAN-wide FlashCopy
- SAN-wide Synchronous PRC

Storage Engine

- Modular, HW/SW integrated solution 1-2 pairs of xSeries storage engines

SAN Software Toolkit

- Framework Managers
  - Tivoli TEC, HP OpenView, CA TNG, BMC
- Storage Resource Managers (SRM)
  - IBM Tivoli Storage Productivity Center (TPC) Standard Edition, Data or Disk
- Storage Element Managers
  - SSPC with TPC Basic Edition
- SAN Exploitation
  - Tivoli Storage Manager (LAN-Free Backup/Server-Free Backup)

SAN Standards

- z/OS - Sure, it's open - anyone can code to IBM's standards and have an excellent chance of compatibility
- Windows/NT - Sure, it's open - anyone can code to Microsoft's standards and have some chance of compatibility
- Storage Area Networks - Sure, it's open - It's multi-vendor! Easy - Right?

Storage Networking Standards

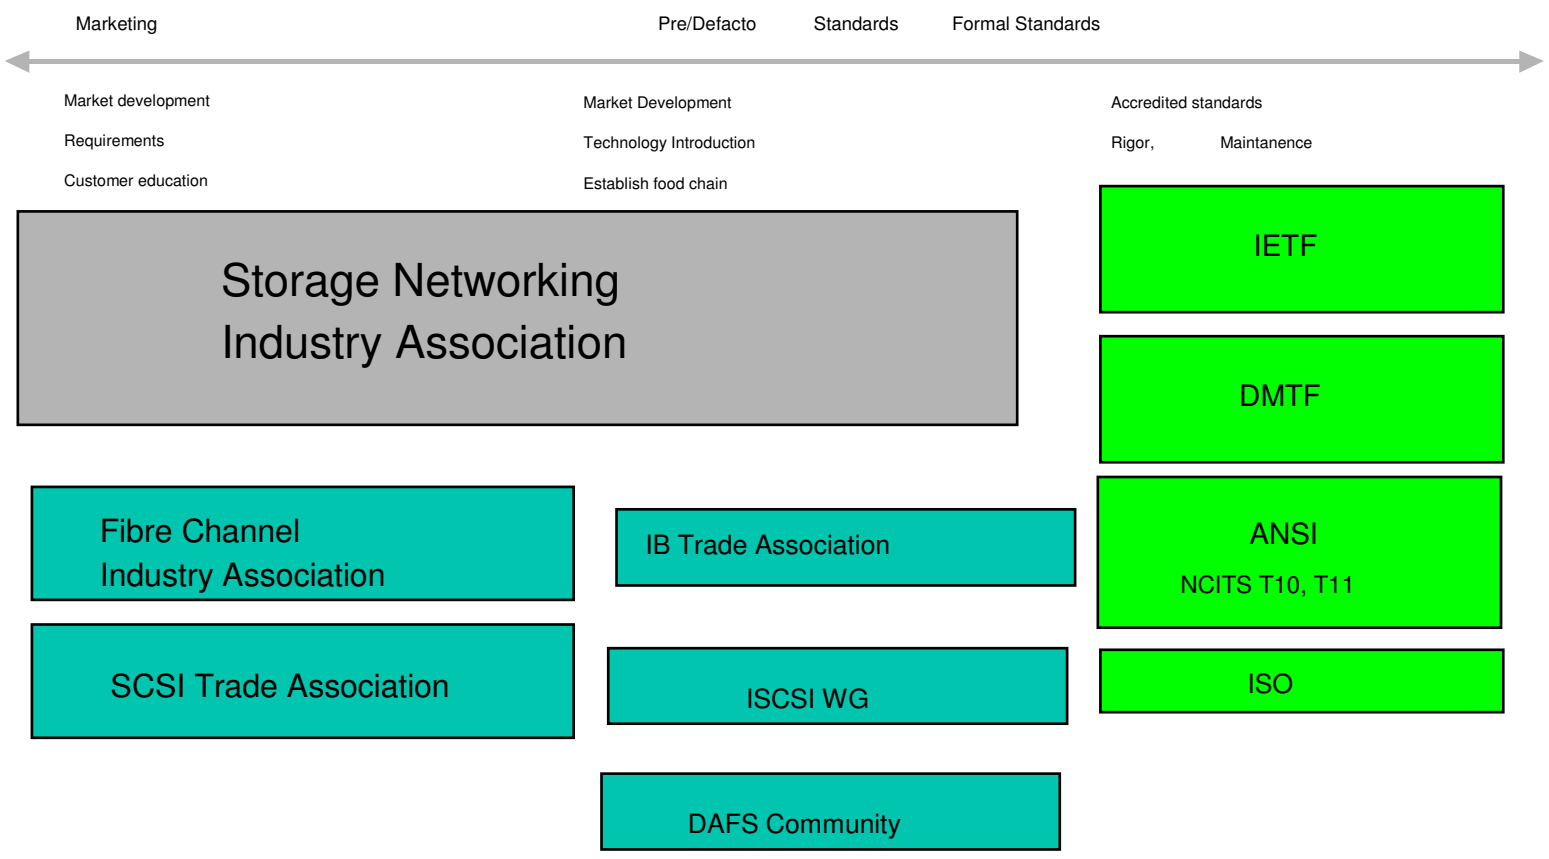

OS/390 - UNIX/NT Interoperability

- Various Ethernet/GbE and Token Ring Implementations (including NFS clients/servers, TSM, DFM, DRDA, etc.) for Server to Server data transfer/access
- InfoSpeed (now owned by CNT) - Server to Server high speed access
- EMC/Innovation Data Processing - Symmetrix/FDRSOS - Direct read of UNIX/NT volumes to MVS tape for full volume backup
- HDS/Harbor - 7700E/Harbor - Direct read of UNIX/NT volumes to MVS Tape for full volume backup (IDP also does this now for HDS)
- EMC Infomover and HDS HMDE - FTP file transfer between UNIX/NT and OS/390 using control unit cache
- Encore/SUN - CD-ROM view (ISO 9660 Standard) read of OS/390 data from UNIX/NT systems (not in a current product)
- "Twin-Tailed" attachment of Storage Control units like the 3494 to OS/390 and AIX for storage hardware sharing (partitioning)
- UNIX Systems Services - HFS and OS/390 flat file transfers (IEBGENER + OGET/OPUT TSO Cmds + XSAM access method)
- EMC/BMC's DataReach (now called Data Extractor from EMC)
- Ex-XPE (9399) function

OS/390 SAN interoperability with FCP SANs - Generic Requirements

- Most customers want to gain painless and quick access to their data regardless of the server on which it resides
- Customers want to be able to manage all SAN elements from a single point of control across heterogeneous operating systems and multiple vendor storage control units
  - Consolidated Multi-Network Management
  - Automated Policy Management for Storage regardless of platform
- There are some storage software exploitation items of the ESCON SAN infrastructure like PPRC, but customers want more - like server free movement of data
- Some people confuse storage control unit functions with SAN functions - i.e. direct read of UNIX/NT volumes from OS/390 inside the ESS/Shark to do volume backup - but call this a SAN requirement
- Have OS/390 provide integrated Security coverage for itself and UNIX/NT
- Have a single unified Tape Management System for OS/390 and UNIX/NT
- Have easier ways of replicating data between OS/390 and UNIX/NT
  - Automated, works with scheduling software on both sides
  - Controllable from either OS/390 or UNIX/NT
  - "Conceptually Simple" - e.g. single step transfer (hiding data extract, data transformation, data loading processes)

SAN Summary

"Information is the currency of an e-business and must be managed as a valuable corporate asset. SANs are the digital bank for that information. This makes SANs one of the most important IT advancements since the emergence of Network Computing."

"Today's business environment is being driven, in large part, by the data explosion fueled by e-business, the commercialization of the Internet, the emergence of data-intensive technologies such as multimedia and data warehousing, and the focus on server and storage consolidation. IBM's Storage Area Network initiative is the next step towards providing centrally managed, open software and hardware solutions designed to help companies get the most value out of their entire business information and IT infrastructures."

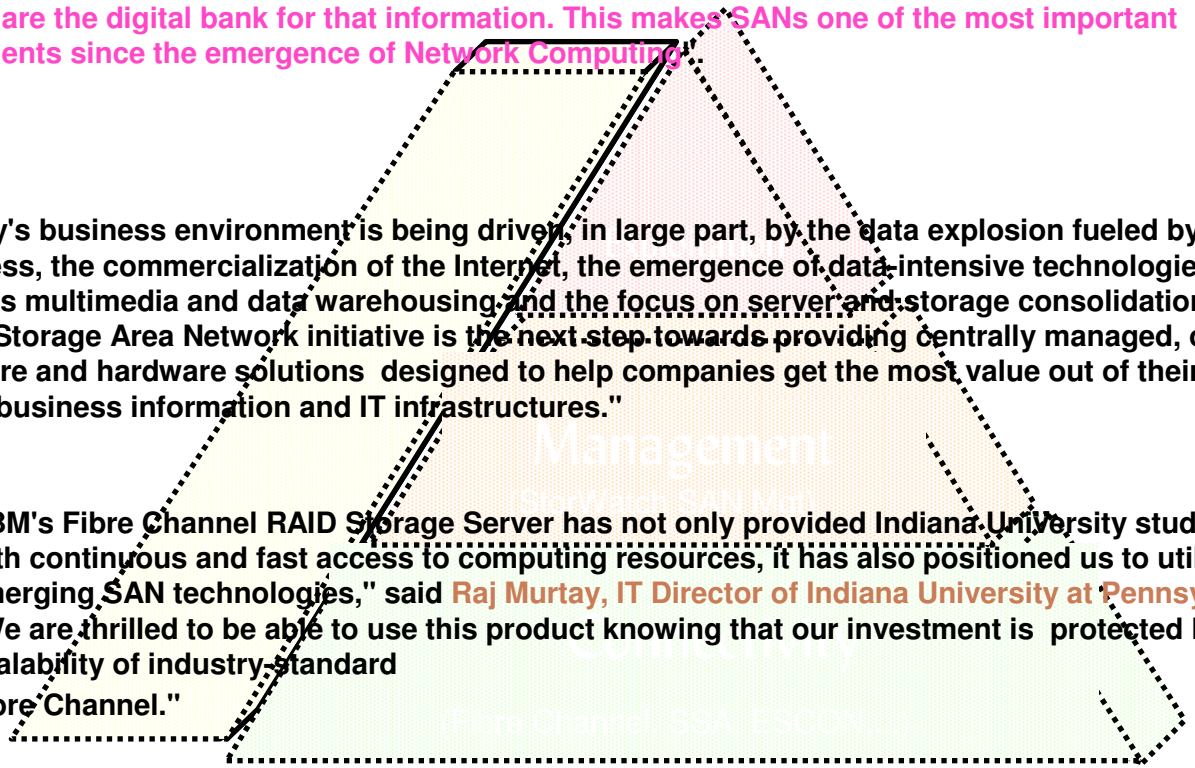

"IBM's Fibre Channel RAID Storage Server has not only provided Indiana University students with continuous and fast access to computing resources, it has also positioned us to utilize emerging SAN technologies," said Raj Murtay, IT Director of Indiana University at Pennsylvania. "We are thrilled to be able to use this product knowing that our investment is protected by the scalability of industry-standard
Fibre Channel."

IBM System Storage & Services

Intelligent Management. Protected Information. Smarter Insights.

For more information: **http://www.ibm.com/systems/storage**

Disk Systems

- SAN Volume Controller
- DS family
- N series

Infrastructure Management

- TotalStorage Productivity Center
- SAN Fabric Management software
- Tivoli Provisioning Manager
- Tivoli Storage Process Manager
- IBM Systems Director family

Tape Systems

- TS family of drives, libraries and virtualization

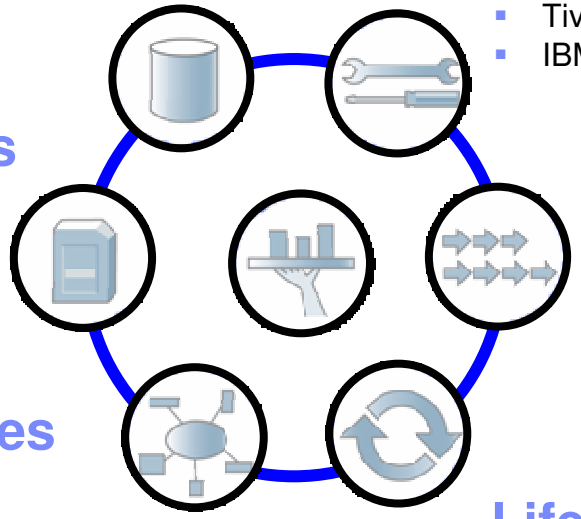

Business Continuity

- Productivity Center for Replication
- Basic Hyperswap for System z
- Advanced copy services
- Tivoli Storage Manager (TSM) family
- Tivoli Continuous Data Protection (CDP)
- Tape cluster grids and Peer-to-Peer
- GDOC, GDPS

Storage Services

- Consulting
- Assessments
- Design
- Deployment
- Outsourcing
- Hosting

Storage Networking

- Switches
- Directors
- Routers

Lifecycle and Retention

- DR550 archive and compliance storage
- Grid Access Manager, GMAS
- TSM Space Management for Unix/Windows
- GPFS, DFSMS, SOFS
- N series with SnapLock™
- WORM tape support

61

Definitions

- CIFS - Common Internet File System; commonly used as a replacement for Netbios on PC systems
- Fabric - A Fibre Channel network consisting of multiple devices interconnected by one or more switches that use Fibre Channel methodology to link nodes and route frames
- FC-AL - Fibre Channel Arbitrated Loop - a shared gigabit media for up to 127 nodes, one of which may be attached to a switch fabric. Uses arbitration to decide which node is using the loop at any one time
- FCP - Fibre Channel Protocol - the mapping of serial SCSI-3 commands to the physical Fibre Channel carrier
- FICON - Mapping an enhanced ESCON protocol onto the physical Fibre Channel carrier
- Gateway - a node of a network that interconnects two otherwise incompatible networks (e.g.. IBM SAN Data Gateway (2108-G07) SCSI - FCAL + LUN Masking function)
- GBIC - Gigabit Interface Converter - allows the optical fibre to be attached to the Host Bus Adapter (HBA) in a server (optical laser to copper)
- HBA - Host Bus Adapter - the I/O card that fits in a server and allows attachment of fibre channel media (optical cable or copper cable)
- Hub - A Fibre Channel device that connects nodes into a logical loop by using a physical star topology. Facilitates FC-AL loops and allows the FC-AL loops to be extended to 10 KM by connecting longwave GBICs.
- LAN/WAN - Local Area Network; Wide Area Network - usually using TCP/IP or CIFS protocols to carry end-user interactive traffic.

Definitions

- NAS - Network Attached Storage - a term used to describe technology where an integrated storage systems is usually attached to a LAN or WAN using TCP/IP or CIFS. There are two general types of NAS - File servers (e.g. NFS) or Backup/Archive servers (IBM 3466 Network Storage Manager)
- Node - An entity with one or more N_Ports or NL_Ports (e.g. Servers or Storage Devices)
- Point to Point - A Fibre Channel topology in which each point has physical links to only one neighbor resulting in a closed circuit. The available bandwidth is dedicated solely to this connection.
- Port - The hardware entity within a node that performs data communications over the Fibre Channel
- Router - A dedicated hardware and/or software package which manages the connection between two or more networks (e.g. IBM SAN Data Gateway Router(2108-R03) - SCSI to FCAL protocol conversion only)
- SAN - Centrally managed high speed networks of multivendor storage subsystems, applications servers, clients and networking hardware that allow compagnies to exploit the value of their business informations via universal access and sharing of resources.
- Switch - A hardware entity with multiple entry/exit points (ports) that provides dynamic connection between any two of these ports.

A Comparison of Channels

ESCON Channel	Native FICON Channel	FCP Channel
Circuit Switching	Packet Switching	Packet Switching
Read or write Half-duplex data transfer	Simultaneous read and write full-duplex data transfers	Simultaneous read and write full-duplex data transfers
Connection-oriented	Connectionless	Connectionless
Dedicated path pre-established	Packets individually routed	Packets individually routed
When packet (frame) is sent, link is locked	When packet is sent, link is released	When packet is sent, link is released
Synchronous data transfer	Asynchronous data transfer	Asynchronous data transfer
Uses CCW Architecture	Uses CCW Architecture	Uses SCSI architecture
Traditional IODEVICE to access a logical volume	Traditional IODEVICE to access a logical volume	QDIO IODEVICE to access ALL logical units

IBM Networking Services fiber cabling solutions

- Introducing
  - Scalable solutions
    - From single product to entire enterprise
  - Fiber optic connectivity expertise deploying a proven methodology
  - Personalized services to effectively plan and install the fiber optic cabling needed for your zSeries with the future in mind
- Addressing the requirements of
  - The Data Center
  - Open Systems Environment (Storage Area Network/ Local Area Network)
  - Parallel Sysplex (Coupling Links)
- Considering
  - Current fiber optic cabling, connectors, transceivers
  - New industry-standard Small Form Factor (SFF) connectors and transceivers
- Delivering
  - A custom or contracted service for your products as well as small, medium, or large enterprise
  - Analysis of your current fiber optic cabling and the zSeries configuration
  - Cabling options customized for your system environment

Flexible, cost-effective, tailored fiber cabling solutions

Legal Disclaimer

NOTICES AND DISCLAIMERS

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This information could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) at any time without notice.

Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectually property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of anynon-IBM product, program or service.

Legal Disclaimer

The information provided in this document is distributed "AS IS" without any warranty, either express or implied. IBM EXPRESSLY DISCLAIMS any warranties of merchantability, fitness for a particular purpose OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (*e.g.*, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.