



IBM TotalStorage®

## G25 - SAN Basics for Mainframers

Scott Drummond  
Program Director – Storage Networking  
[spd@us.ibm.com](mailto:spd@us.ibm.com)

## Permissions & Abstract

- **Permission to Copy:** Permission is granted to SHARE to publish this presentation paper in the SHARE proceedings. IBM retains the right to distribute copies of this presentation to whomever it chooses.
- **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 terminology as a "Rosetta Stone" for translation purposes

# Storage Management Challenge

## ■ Costs of Storage Management


- A Well Managed SAN in 2005 costs 1 Person per Terabyte (industry generic number)
  - This assumes the 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!
- z/OS w/Automated Storage Policies costs 1 Person per 3 to 17 Terabytes
  - Customer quotes




# 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 - a few customers today!
- ◆ Petabytes ( $10^{15}$ ) (Quadrillion) of tape storage - large customers
- ◆ Exabytes ( $10^{18}$ ) (Quintillion) of tape storage - a few accounts by 2004
- ◆ Zettabytes ( $10^{21}$ ) (Sextillion) WW digital data by 2004/5 - IDC
- ◆ Yottabytes ( $10^{24}$ ) (Septillion) - How Long will it take?

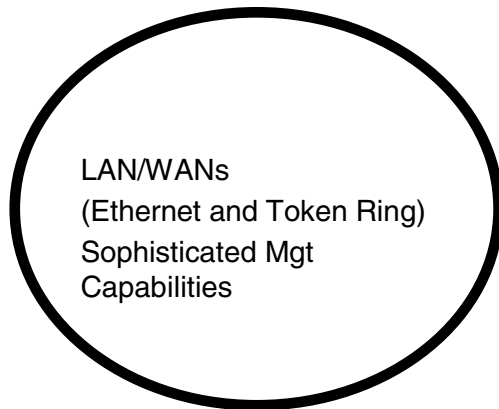
# IBM SAN Evolution



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



- 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

- Delivers immediate benefits:
  - ▶ Improved performance
    - ▶ Leads to more transaction or higher throughput
  - ▶ Improves availability
    - ▶ Less planned outages (90+% of all outages are planned)
  - ▶ Better consolidation & costs
    - ▶ Storage consolidation (15-50% of DAS storage is utilized) – SAN plus storage virtualization can be higher than 80%
    - ▶ Environmental costs can be improved with SANs rather than having many scattered “mini-IT shops”
    - ▶ The costs of more sophisticated equipment (like the DS8000/6000 or tape libraries) can be amortized across many servers – making them more affordable
    - ▶ People costs can be improved by having a dedicated SAN mgt group rather than forcing system administrators and end users to perform storage management tasks or storage recoveries
  - ▶ Better storage capacity sharing
  - ▶ Better storage attachment flexibility
  - ▶ Better scalability
  - ▶ Prerequisite for Storage Virtualization & Information Life Cycle Mgt (ILM...really HSM!)

## 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
  - ▶ 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

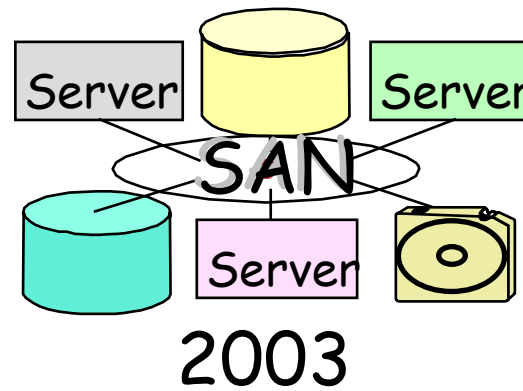
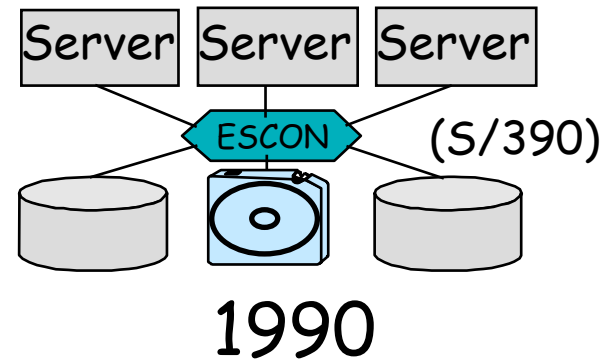
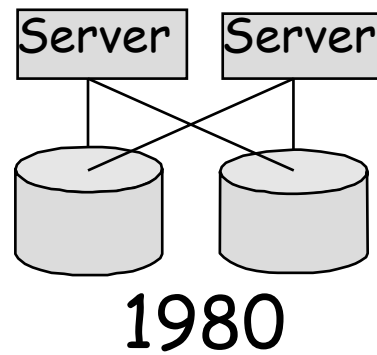



# z/OS ESCON/FICON SAN

- 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



# SAN Solution Building Blocks

## ■ Servers

- UNIX, Windows, Novell, LINUX, z/OS

## ■ SAN Fabric Components

- Switches/Directors, Hubs, Gateways/Routers

## ■ Storage Virtualization Engines

- Symmetric – LUN Management & Utilization
- Asymmetric Virtualization – Shared File Systems

## ■ Storage

- Disk Control Units, Tape Drives, Tape Libraries


## ■ Software

- SAN Management, SAN Exploitation (LAN Free and Server Free Backup)

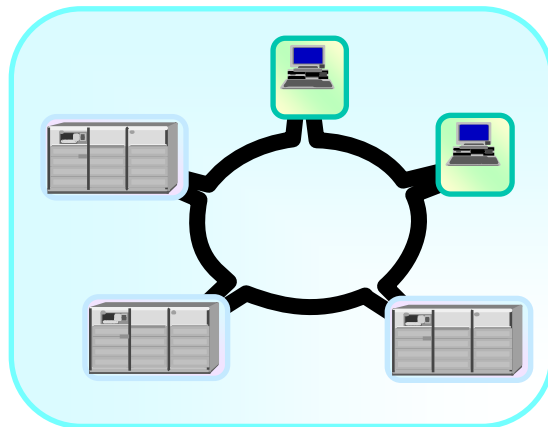
## ■ Services

- Planning, Testing and Implementation, Education


# SAN Topologies



FC Point to Point

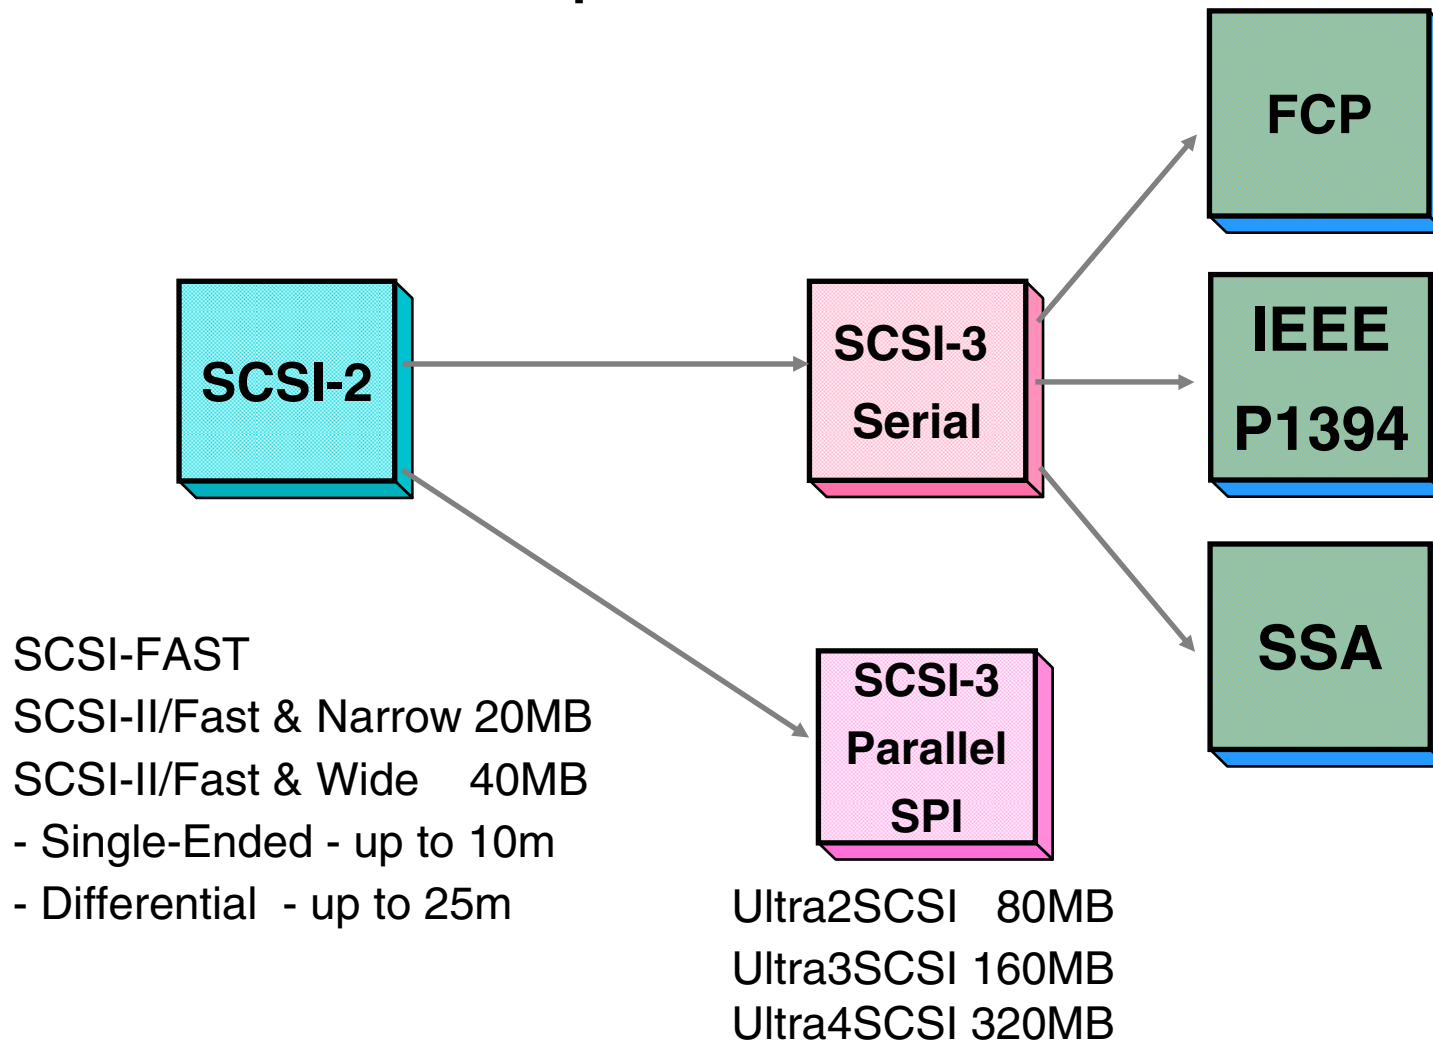


FC Arbitrated Loop



Switched Fabric

# SCSI Roadmap



# 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)

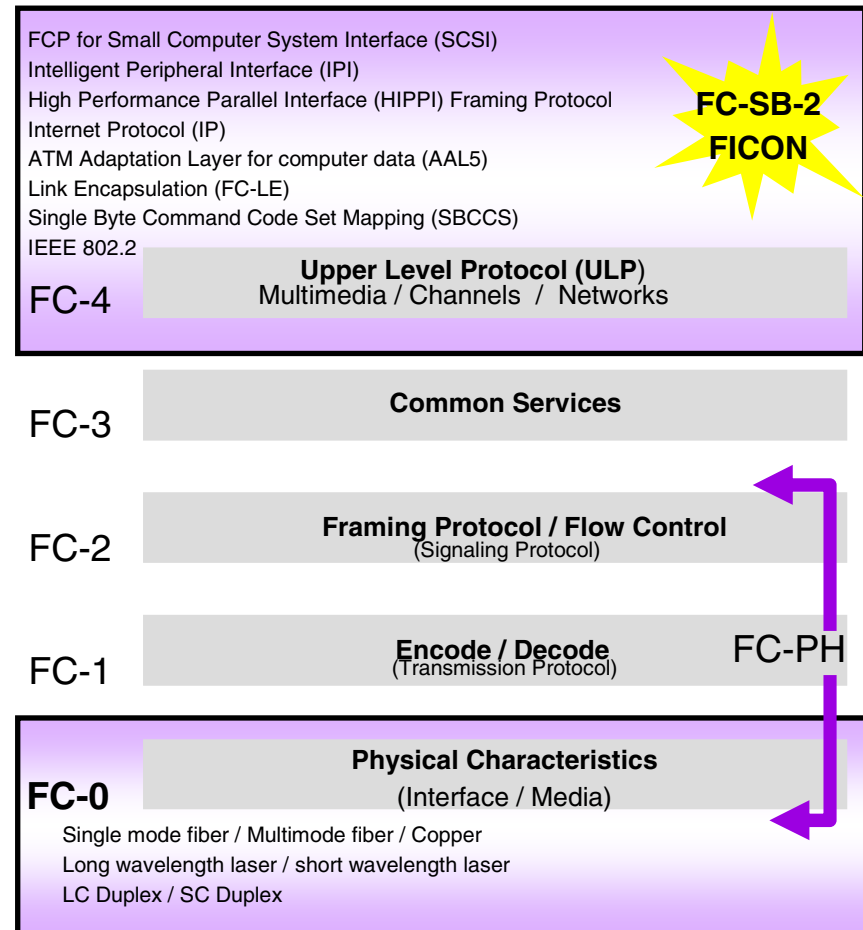
See also [www.t11.org](http://www.t11.org)

## ■ FICON

- IBM ^™ zSeries™ implementation of Fibre Channel Architecture
- An industry standard under the name FC-SB-2

## ■ FCP

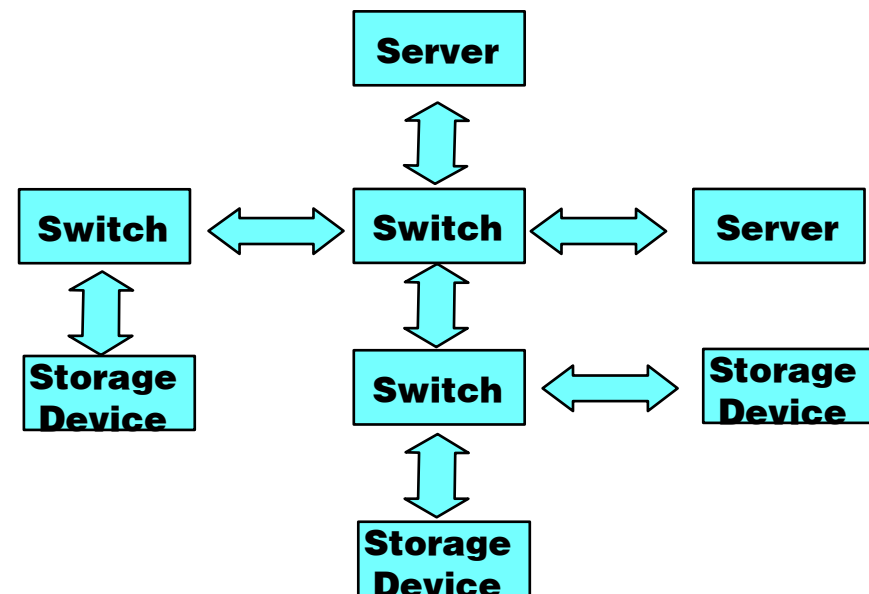
- Fibre Channel Protocol for SCSI
- Mapping of the SCSI command protocol onto the Fibre Channel Architecture



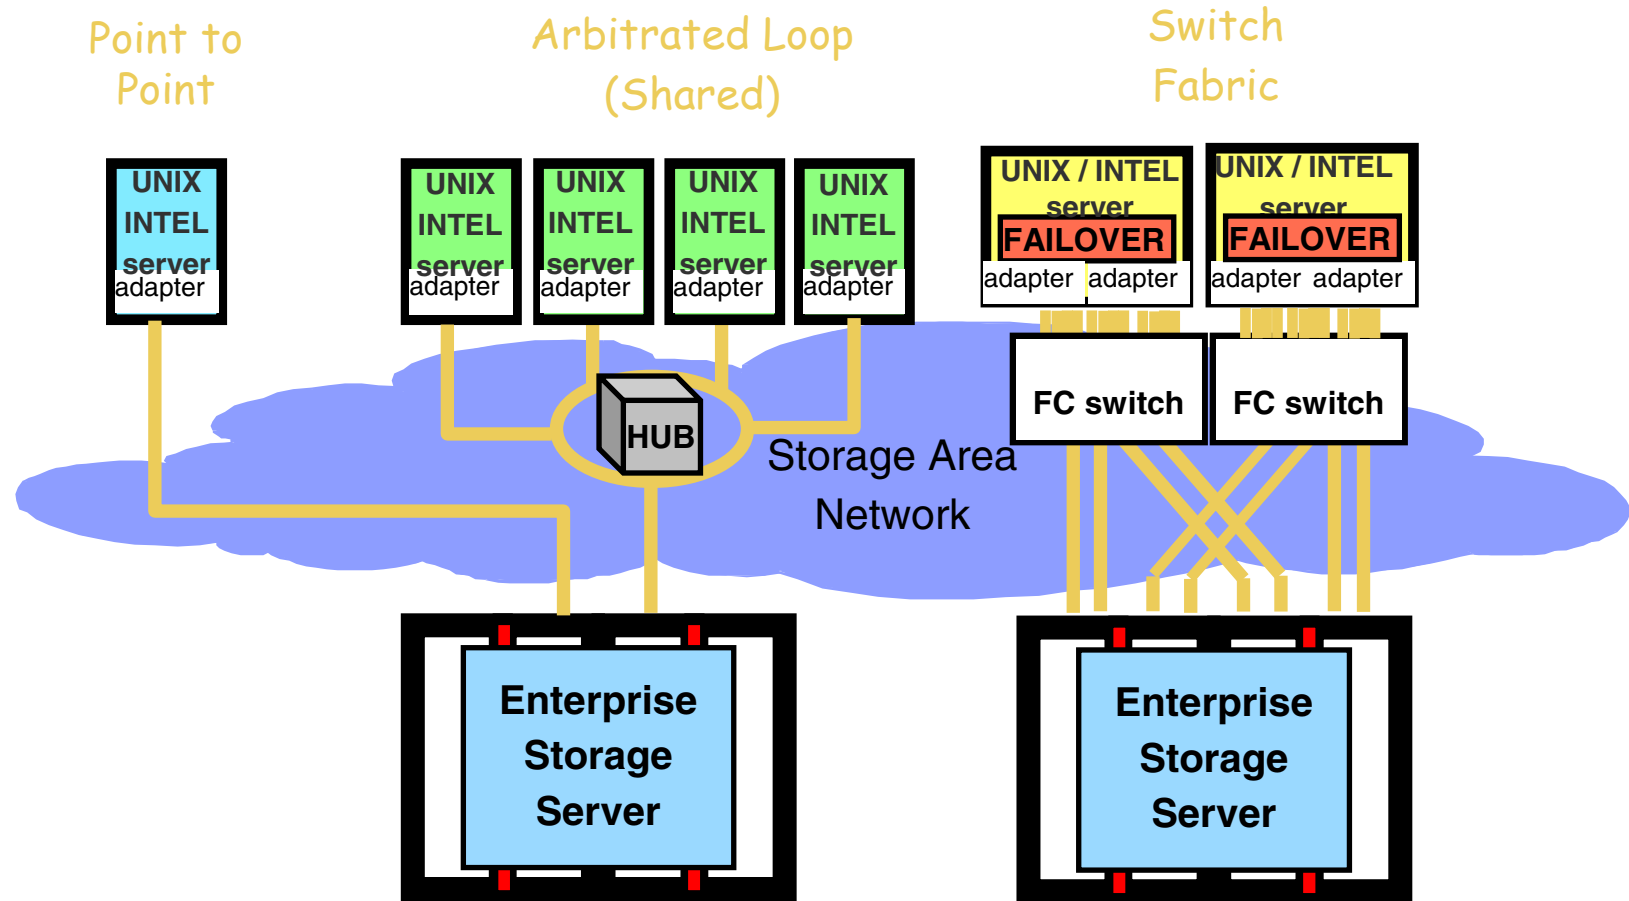
# Cascading Switches/Directors

## ■ Cascading Switches

- Currently like switches only – provides redundancy, ISL trunking, preferred routing
- Communicate Fabric Topology to Each Other
- Switch manufacturers
  - working on standards for better interoperability (OSFI)



# FCP Topologies



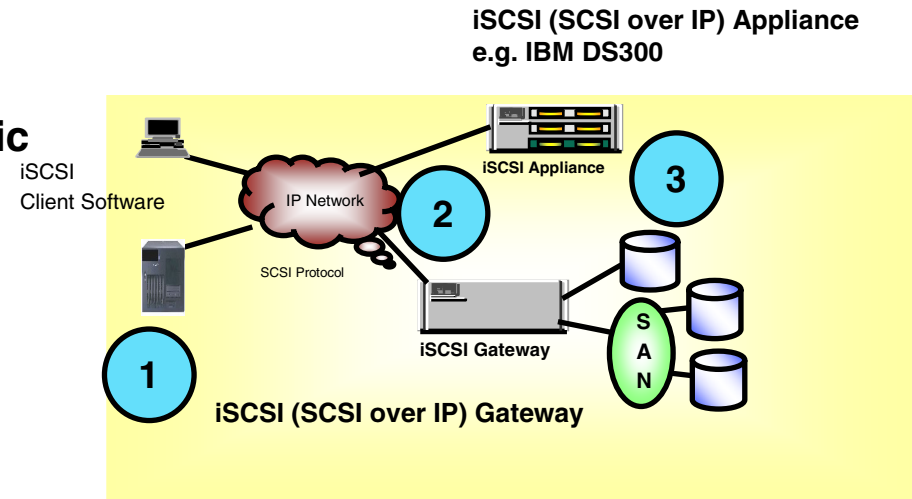
# Internet SCSI (iSCSI) A New and Emerging Technology

- 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--

iSCSI appliances (with embedded storage)  
Gateways (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)






# Who's Playing with iSCSI?

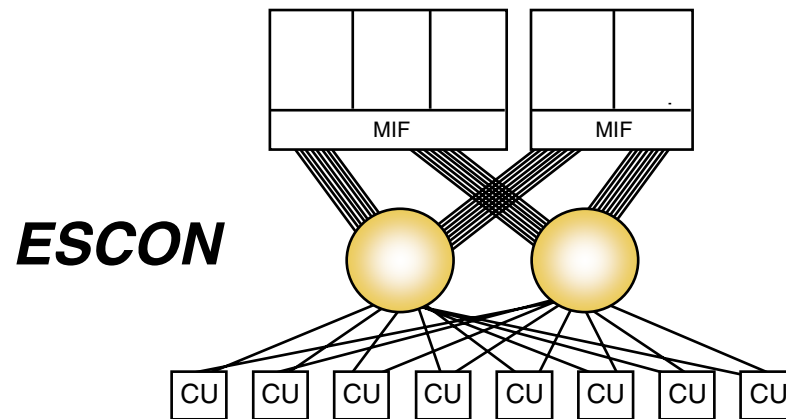
- Brocade, Cisco, McDATA (CNT), NetApps and several other companies offer an iSCSI gateway functionality that will provide the physical and logical connection between the ethernet attached servers and the fibre channel attached storage
- IBM offers the DS300 iSCSI Storage Array
  - 14-Drive, 1Gbe iSCSI to U320 SCSI subsystem
  - Single and dual RAID controller options
  - Scalable to 2 TB (14 drives)
  - Common software, RAID code and drive carrier
  - IBM eServer xSeries iSCSI Adapter (73P3601) available to provide higher performance, remote boot and failover protection
- Other companies offer iSCSI Storage Array solutions – Stonefly, Lefthand, SANRAD, Intransa, EqualLogic, Nimbus Data

# IBM Enterprise Systems CONnection = ESCON




## ■ ESCON using fiber optics

- Light weight
- Long distance
- Nondisruptive changes
- Up to 120 CUs
- Up to 1 K devices (512 devices for CTC)




# IBM Fibre CONnection = FICON




## ■ FICON and FICON Express

- Greater bandwidth
- Configuration flexibility
- Longer extended distance
- Up to 256 CUs
- Up to 16 K devices  
(240 devices for FCP)




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



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

With 2 Gbps Inter Switch Links (ISLs), fewer fiber links may be needed for cross-site connectivity



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

- 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

# Why Use Native FICON Channels?

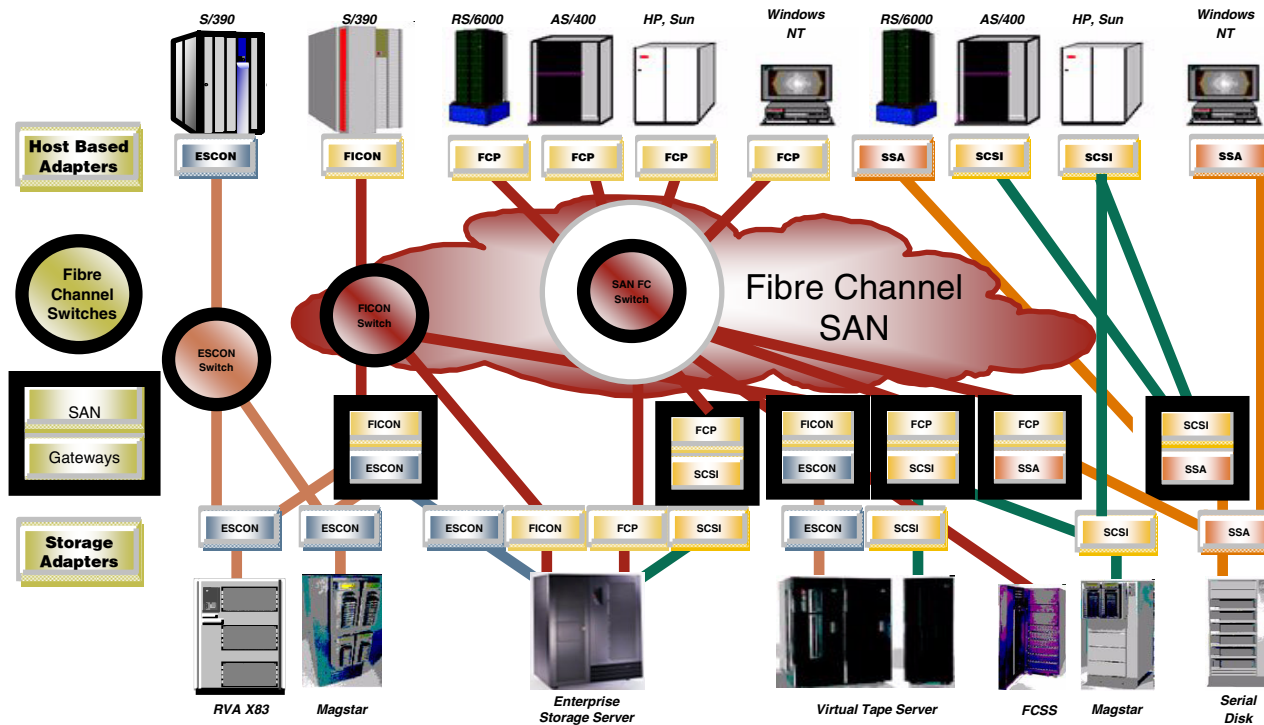
- FICON is more than just an I/O card or a channel type
  - FICON is at the heart of some very important solutions
- Performance
  - Faster backup/ recovery times, shorter batch windows, faster data access for large data queries
  - See the FICON and FICON Express Channel Performance white paper (February 2002) located at [ibm.com/servers/eserver/zseries/connectivity/](http://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

# 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"



# Entry SMB Switch Portfolio

 <p><b>SAN16B-2</b> (2005-B16, PN200516) 8, 12, 16 ports, 4 Gbps <a href="http://www.ibm.com/totalstorage/san/b-type">www.ibm.com/totalstorage/san/b-type</a></p>	 <p><b>IBM SAN16M-2</b> (2026-416, PN 202616E) 4/8/12 ports, 2 Gbps ports <a href="http://www.ibm.com/totalstorage/san/m-type">www.ibm.com/totalstorage/san/m-type</a></p>	<p>White Space Opportunity First-time SAN buyers</p>
 <p><b>SAN Switch H08</b> (2005-H08, PN2005H08) 4-8 ports, 2 Gbps <a href="http://www.ibm.com/totalstorage/san/b-type">www.ibm.com/totalstorage/san/b-type</a></p>	 <p><b>IBM SAN12M-1</b> (2026-E12, PN 202612E) 4/8/12 ports, 2 Gbps ports <a href="http://www.ibm.com/totalstorage/san/m-type">www.ibm.com/totalstorage/san/m-type</a></p>	 <p><b>Storage Switch L10</b> (PN2006L10) 4-10 ports, 1 or 2 Gbps <a href="http://www.ibm.com/storage/2006">www.ibm.com/storage/2006</a></p>

- Integrated, simple storage consolidation and data protection solutions
  - Homogenous Windows/Linux servers
  - IBM® eServer™ xSeries® Server sales channels
  - IBM TotalStorage DS4000 (FAStT) and LTO storage
  - Designed to provide high availability with dual fabric deployment



# Midrange SMB 4 Gbps Switch Portfolio

**SAN Upgrade Opportunity**  
4Gbps, SAN extension

## IBM TotalStorage b-type (Brocade) switches



**IBM SAN16B-2** (2005-B16 PN 200516B) **IBM SAN32B-2** (2005-B32 PN 200532B)  
8, 12, 16 ports, 4 Gbps **FICON**  
[ibm.com/totalstorage/san/b-type](http://ibm.com/totalstorage/san/b-type) [ibm.com/totalstorage/san/b-type](http://ibm.com/totalstorage/san/b-type)

## IBM TotalStorage m-type (McDATA) switches



**IBM SAN16M-2** (2026-416 PN 202616E) **IBM SAN32M-2** (2026-432 PN 202632E)  
8, 12, 16 ports, 4 Gbps **FICON**  
[ibm.com/totalstorage/san/m-type](http://ibm.com/totalstorage/san/m-type) [ibm.com/totalstorage/san/m-type](http://ibm.com/totalstorage/san/m-type)

## Cisco MDS 9000 switches



**Cisco 9120/40** (2061-020/040)  
4-20/40 FC ports, 2 Gbps  
(4/8 target/16; 32 host optimized)  
[www.ibm.com/storage/cisco](http://www.ibm.com/storage/cisco)

**Cisco 9216i** (2062-D1A/D1H)  
14/16/32/48 FC ports, 2 Gbps  
2/4/8 GbE IP storage ports  
[www.ibm.com/storage/cisco](http://www.ibm.com/storage/cisco)

- Designed to offer integrated, scalable, high availability *IBM Virtualization* family solutions
- Heterogeneous Windows, Linux, iSeries, UNIX servers
  - xSeries®, iSeries, pSeries® Server™ sales channels
  - IBM DS4000 (FAStT), DS6000, DS8000, LTO and ETS storage

# IBM Director Portfolio

Technology Upgrade Opportunity  
SAN Consolidation & extension

**NEW**




**IBM SAN256B**  
256 FC/FICON , 2, 4 Gb ports  
[ibm.com/totalstorage/san/b-type](http://ibm.com/totalstorage/san/b-type)




**IBM M14 (2109-M14)**  
32-128 FC/FICON 2 Gb ports  
[ibm.com/totalstorage/san/b-type](http://ibm.com/totalstorage/san/b-type)


**NEW**



**IBM SAN256M**  
256 FC/FICON 2 Gb ports  
10Gbps ISL ports  
[ibm.com/totalstorage/san/m-type](http://ibm.com/totalstorage/san/m-type)



**IBM SAN140M (2027-140)**  
16-140 FC/FICON 2 Gb ports  
[ibm.com/totalstorage/san/m-type](http://ibm.com/totalstorage/san/m-type)



**IBM SAN256N (2045-N16)**  
32-256 FC/FICON 2Gb ports  
10Gbps ISL ports  
[ibm.com/totalstorage/san/n-type](http://ibm.com/totalstorage/san/n-type)

**NEW** McDATA/CNT roadmap  
final enhancements in  
firmware R1.5



**Cisco 9606 (2062-D04/-T04)**  
14,16-128 FC/FICON 2 Gb ports  
Three 2+14, 4 or 8 IPS Modules  
[www.ibm.com/storage/cisco](http://www.ibm.com/storage/cisco)



**Cisco 9509 (2062-D07/-T07)**  
28, 32-224 FC/FICON 2 Gb ports  
Five 2+14, 4 or 8 IPS Modules  
[www.ibm.com/storage/cisco](http://www.ibm.com/storage/cisco)


- Designed to provide high availability and scalability and intelligent software to help simplify management of complex, integrated enterprise SANs
- Heterogeneous Windows, Linux, iSeries, UNIX and mainframe servers
  - xSeries, iSeries, pSeries, zSeries Server sales channels
  - IBM DS4000, DS6000, DS8000, LTO and ETS storage

# IBM WAN and SAN Routing and Intelligent Fabrics



**NEW**

**Cisco Multiprotocol Services**  
 Fourteen 2 FC, two GbE ports  
 iSCSI Gateway  
 FCIP Tunneling w/compression  
 Inter-VSAN routing  
[www.ibm.com/storage/cisco](http://www.ibm.com/storage/cisco)



**CNT UltraNet Edge Storage Router\***  
 Two 2Gbps FC ports, two GbE ports  
 Tape pipe lining,  
 automatic failover,  
 data compression  
[www.cnt.com](http://www.cnt.com)



**NEW**

**IBM SAN16M-R SAN Router**  
 Twelve 2Gbps FC ports, four GbE ports  
 iSCSI Gateway  
 iFCP routing w/compression  
 FC-FC SAN Routing  
[www.ibm.com/totalstorage/san/m-type](http://www.ibm.com/totalstorage/san/m-type)



**NEW**

**Cisco Storage Service Module**  
 32 FC ports and intelligent fabric  
 support for DS4000 metro &  
 SVC global mirroring  
[www.ibm.com/storage/cisco](http://www.ibm.com/storage/cisco)



**NEW**

**IBM SAN16B-R SAN router**  
 8/16 ports, 2 Gbps or GbE  
 iSCSI Gateway  
 FCIP Tunneling  
 FC-FC SAN Routing  
[www.ibm.com/totalstorage/san/b-type](http://www.ibm.com/totalstorage/san/b-type)



**NEW**


**IBM SAN04M-R SAN Router\*\***  
 Two 1Gbps FC ports, two GbE ports  
 iSCSI Gateway  
 iFCP routing w/compression  
 FC-FC SAN Routing  
[www.mcdata.com](http://www.mcdata.com)

\* IBM storage (DS8000, SVC) supported and \*\* IBM TSP McDATA 1620



# Storage Evolution

*Technology Shifting Towards the Network and Software*



*Value in Network-based Storage Infrastructure Software*

# Storage Virtualization

- Block Level Virtualization solutions w/wo storage - "in-band" (symmetric) or "out-of-band" (asymmetric)
- Block Level Virtualization solution wo/ storage allow multiple vendors storage to attach and look as one large pool for better investment protection (IBM SAN Volume Controller)
- SAN-Wide File System (IBM SAN File System) provide file system virtualization for better availability, manageability and data access - uses "out-of-band" technology

# 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

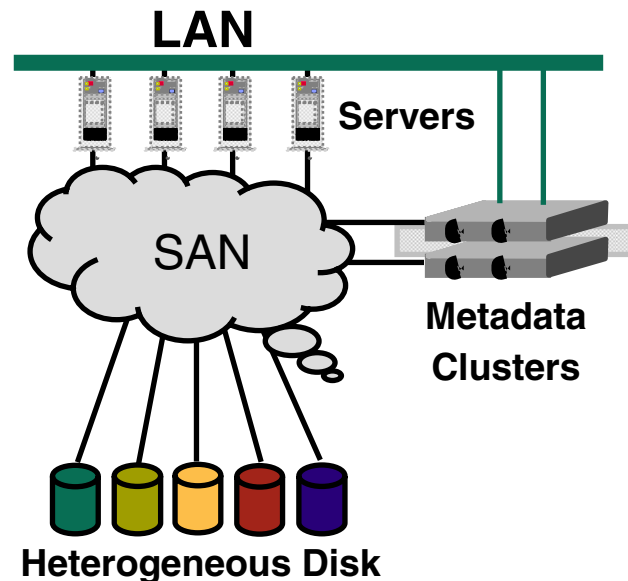
## ■ Symmetric (in-band)

- 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 ...



# 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 Volume Controller Architecture

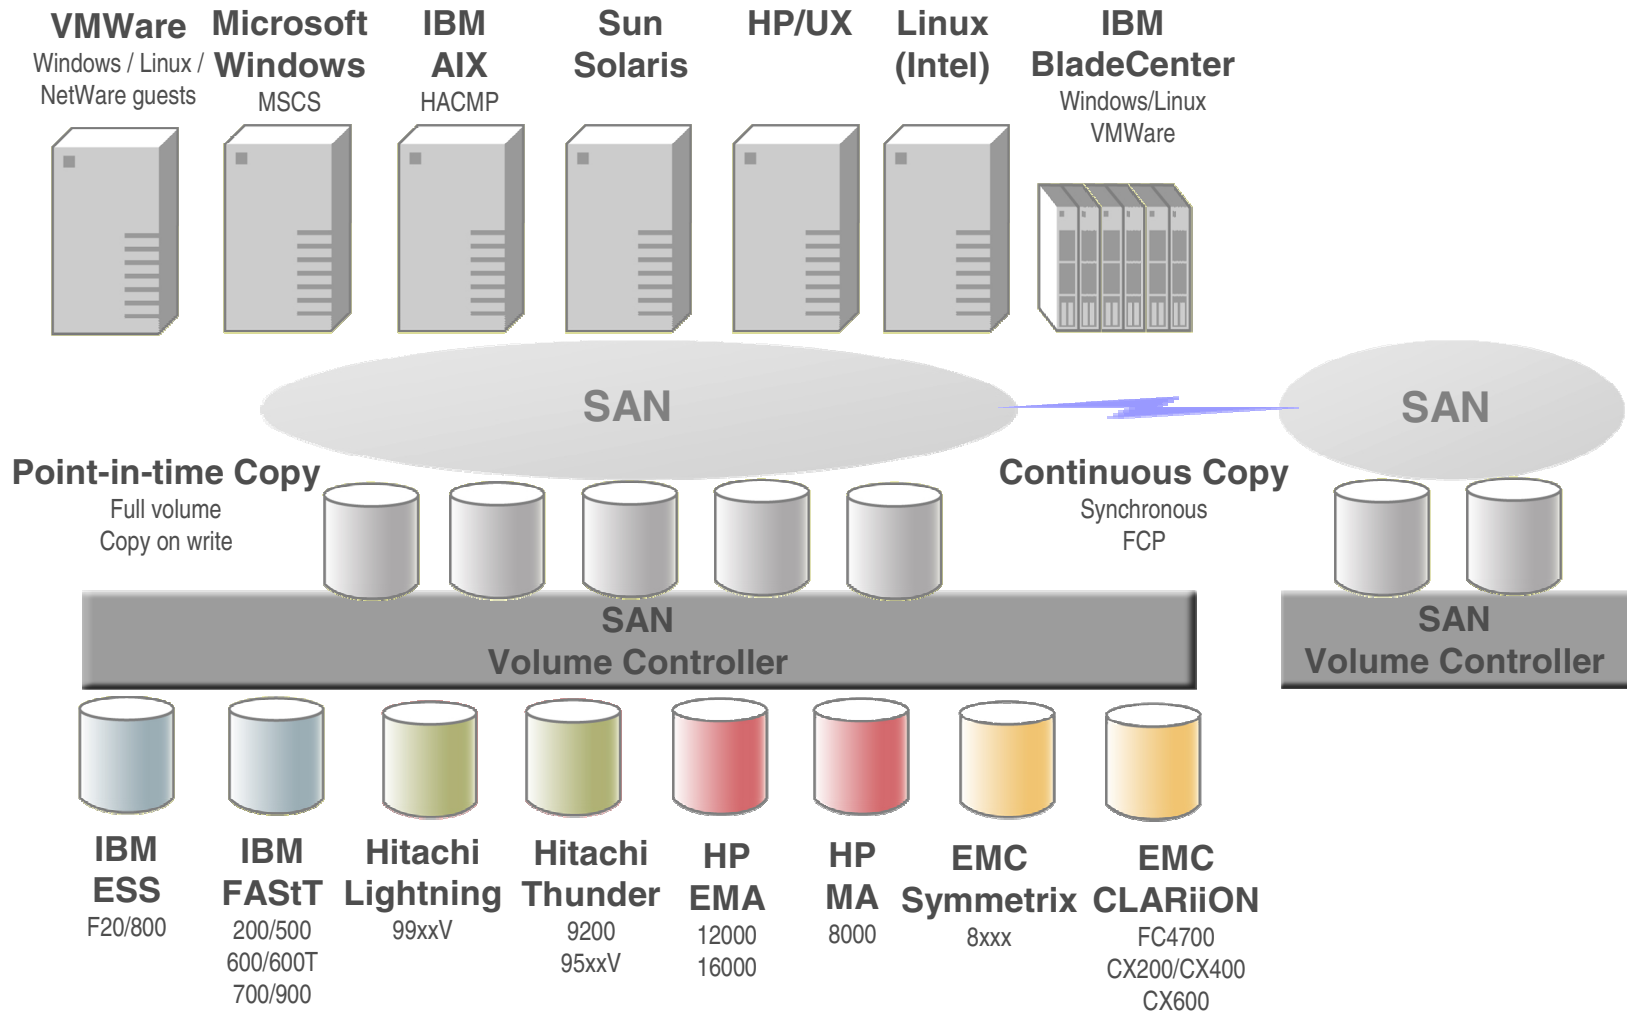


To ensure high availability and peak performance, each **SAN Volume Controller node** pair is connected to your SAN using 8 fibre channel ports – four ports per node.


The SAN Volume Controller provides in-band storage virtualization by creating a pool of **managed disks** from attached back-end disk storage subsystems. These managed disks are then mapped to a set of virtual disks for use by various host computer systems.

**vDisks** are created using managed disks within a managed disk group (MDG). Accordingly, all the managed disks in a single MDG should have the same (or very similar) performance characteristics.

# SAN Volume Controller Supported Environments



# IBM® TotalStorage® SAN File System



## Rules for File Provisioning

- /HR go into Storage Pool 1
- \*.bak go into Storage Pool 4
- DB2.\* go into Storage Pool 2
- \*.tmp go into Storage Pool3

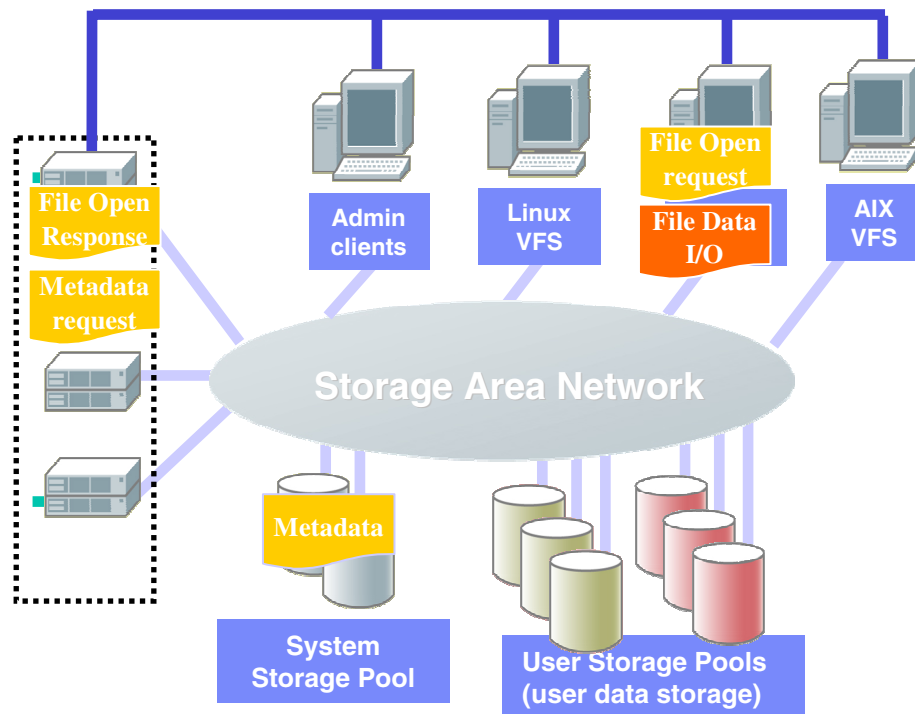
## BENEFIT

- Simplify and improve access to data
- Simplify storage management improving productivity
- Improve storage utilization
- Improve application availability
- Provide true heterogeneous sharing of data in a SAN environment
- Provide automated policy based data management

## FEATURES

- Policy management for file placement
- Supports non-disruptive migration of files
- Reduces downtime for server consolidation
- Improves access to files
- Reduces out of space conditions
- Supports heterogeneous servers and storage


# SAN File System Architecture



## Typical I/O Request:

1. Application makes “File Open” request. This is intercepted by SAN File System Client and directed to Metadata server.
2. Metadata server gets metadata (file location, permissions, locks, attributes) stored on system location on SAN
3. Metadata server verifies available locks and responds to “File Open” request
4. SAN File System Client verifies security permissions and provides access to file
5. Application conducts direct I/O with storage device on SAN. Achieves near local file performance

# SAN File System Policy-Based Automation



- **Storage and file management feature**
  - Designed to provide automatic file allocation
  - Can help ease management of storage growth
- **Policy-based management uses policy sets**
- **A policy is an ordered list of rules**
  - Multiple policies can be stored but only one is active
- **Rules and how are they used**
  - Rules determine which storage pool each file is automatically allocated to
    - Files matching a rule are placed in stg. pool specified by the rule
    - If no rule applies, then file is allocated to the default pool
  - File allocation to storage based on business requirements
    - Storage pools may have different service characteristics
  - Rules use SQL-like language
- **Placement only enforced at file creation time**

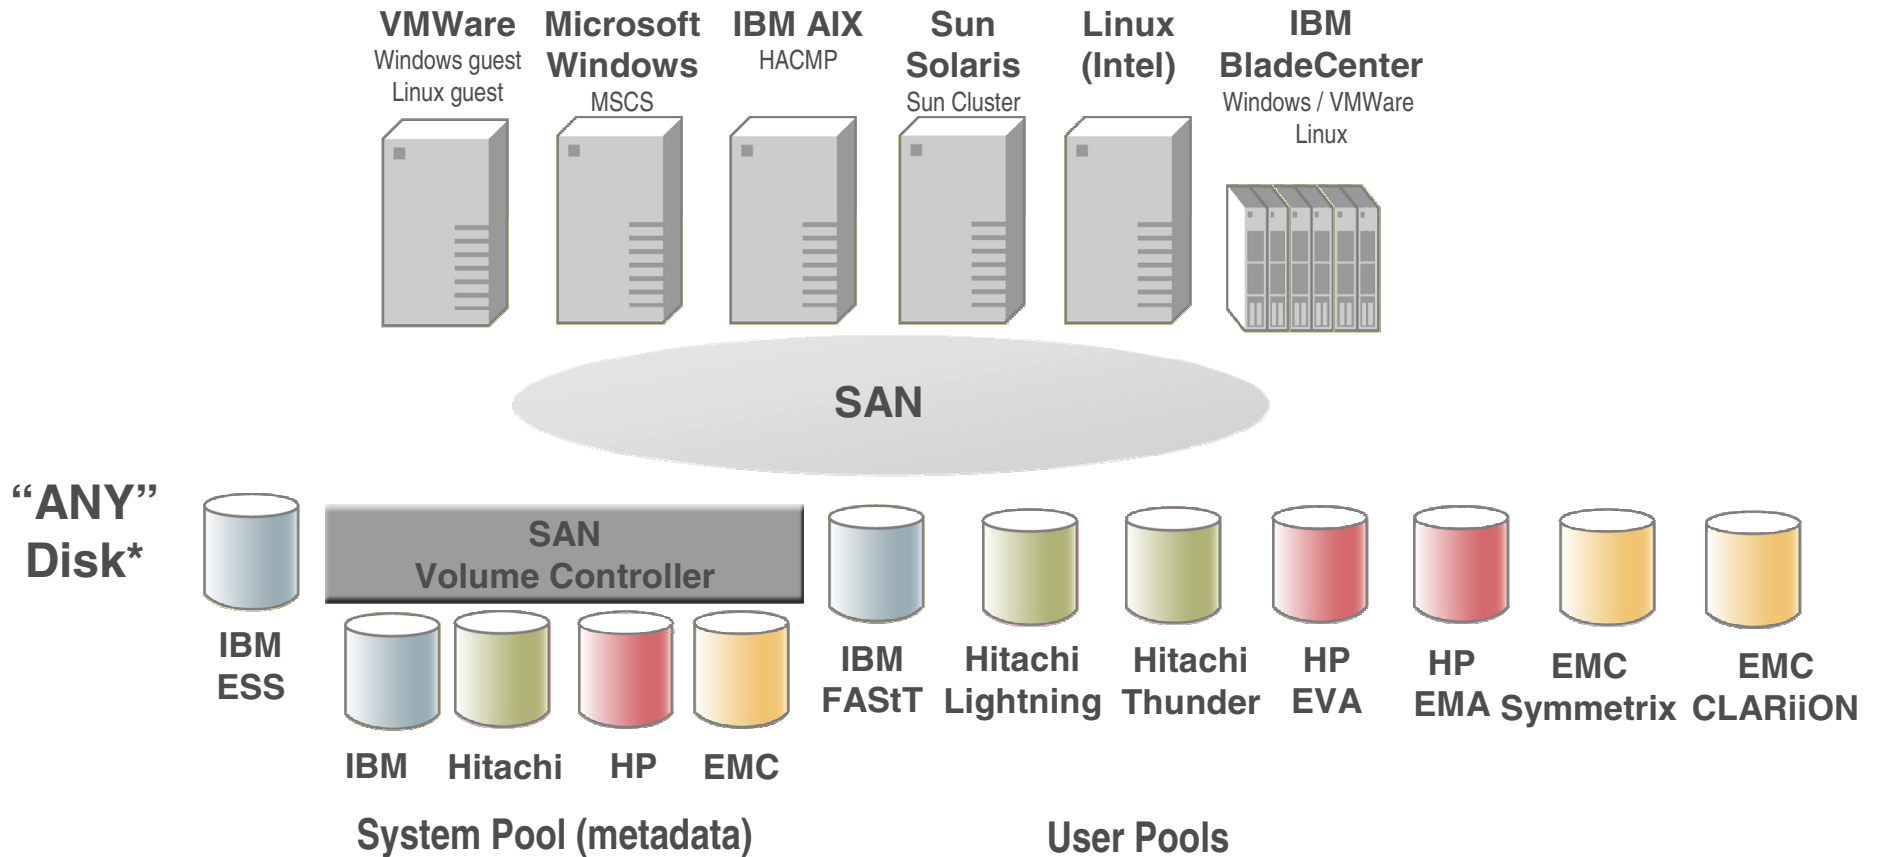


## System Managed Storage - Terminology

System-Managed Storage Concept	Mainframe Terminology	SAN File System Terminology
Information and Location of data	The Integrated Catalog Facility (ICF)	Metadata stored in System Pool
Concurrent access to data	Global Resource Serialization (GRS)	Read-only vs. Read/Write Shared vs. Exclusive
Pooled Storage Resources	DFSMS Storage Groups	User Pools
Policy-Based Management for Automation	Automatic Class Selection (ACS) Routines	Policy Sets containing Placement and Service Rules



# SAN File System Supported Environments

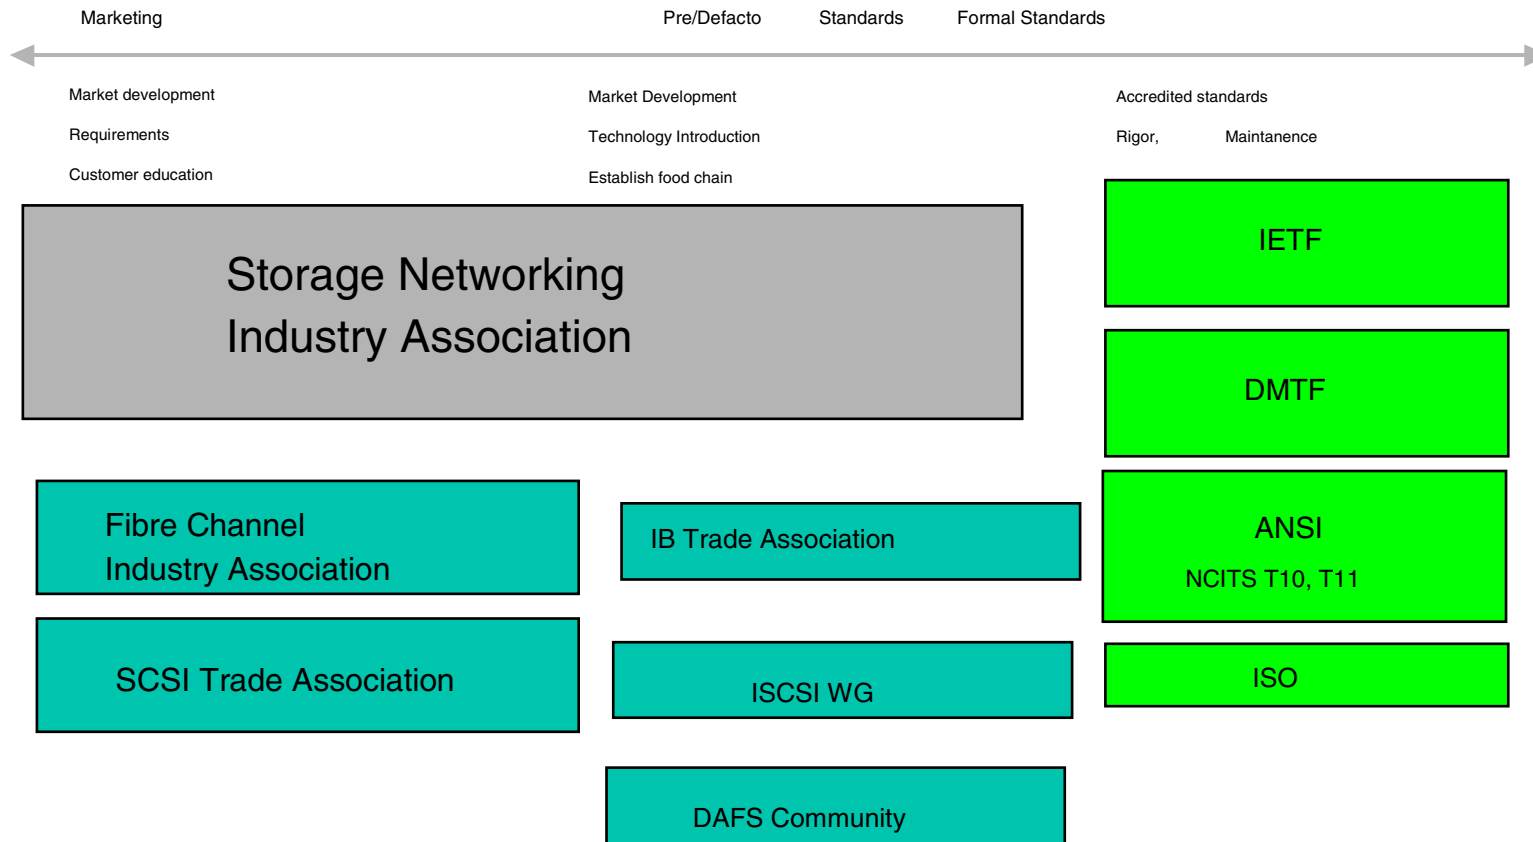


\* One IBM SAN Volume Controller, FASTt or ESS required for metadata information

# 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



# SAN Software Toolkit

## ■ Framework Managers

- Tivoli TEC, HP OpenView, CA TNG, BMC

## ■ Storage Resource Managers (SRM)

- IBM TotalStorage Productivity Center for Data (used to be TSRM)

## ■ SAN Fabric Managers

- IBM TotalStorage Productivity Center for Fabric (used to be TSANM)

## ■ Storage Element Managers

- IBM TotalStorage Productivity Center for Disk
- IBM TotalStorage Productivity Center for Replication

## ■ SAN Exploitation

- Tivoli Storage Manager (LAN-Free Backup/Server-Free Backup)

# 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 Murthy, 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."

Fibre Channel, SSA, ESGON

# 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 companies 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

<b>ESCON Channel</b>	<b>Native FICON Channel</b>	<b>FCP Channel</b>
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***

# Trademarks

The following terms are trademarks or registered trademarks of the IBM Corporation in either the United States, other countries or both.

IBM, DFSMS, DFSMSHsm, DFSMSdss, Enterprise Storage Server, SSA, ESCON, FICON, OS/390, S/390, RS/6000, AIX, Netfinity, z/OS, zSeries, pSeries, iSeries, xSeries and TotalStorage

Windows 2000, Windows NT, Windows 95 and Windows 98 are registered trademarks of Microsoft Corporation.

Novell and NetWare are registered trademarks of Novell, Inc.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

CA ArcServeIT is a registered trademark of Computer Associates International, Inc.

Legato and NetWorker are registered trademarks of Legato Systems, Inc.

Tivoli and Tivoli Storage Manager are registered trademarks of Tivoli.

Veritas is a registered trademark of Veritas Software Corporation.

Intel is a registered trademark of Intel Corporation.

HP and HP-UX are trademarks of Hewlett-Packard Company.

Sun and Solaris are trademarks of Sun Microsystems, Inc.

Brocade, SilkWorm, Extended Fabrics, Remote Switch, Fabric OS, Fabric Watch, QuickLoop, Zoning, Inter-Switch Link Trunking are trademarks or registered trademarks of Brocade Communications System, Inc.

INRANGE is a registered trademark of INRANGE Technologies Corporation. FC/9000 and IN-VSN are trademarks of INRANGE Technologies Corporation.

McDATA and Fibre Channel Director are a trademark or registered trademark of McDATA Corporation.

Other company, product, and service names may be trademarks or registered trademarks of their respective companies.



## Additional Resources

- **The Campus for more education:**

**BPs:** <http://www.ibm.com/partnerworld/sales/systems/education>

Includes link to IBM PartnerWorld University (Web lectures for key topics)  
<http://www.ibmweblectureservices.ihost.com/pwu>

**IBMers:** [w3.ibm.com/sales/systems/education](http://w3.ibm.com/sales/systems/education)

Includes links to the Online Universities for Cross-Brand and each Brand (Web lectures for key topics)

**Customers:** [www-1.ibm.com/servers/eserver/education](http://www-1.ibm.com/servers/eserver/education)

# Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml): CICS\*, DB2\*, DB2 Universal Database, e-business logo, IBM, IBM Logo\*, IMS, Intelligent Miner\*, Lotus\*, MQSeries\*, Multiprise\*, MVS, Net.Commerce, NetView\*, On demand business logo, OS/390, Parallel Sysplex\*, PassPort Advantage\*, PR/SM, Tivoli\*, Websphere, z/OS\*, zSeries\*, z/VM\*, zSeries Entry License Charge.

The following are trademarks or registered trademarks of other companies

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation  
Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries  
LINUX is a registered trademark of Linux Torvalds  
UNIX is a registered trademark of The Open Group in the United States and other countries.  
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.  
SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.  
Intel is a registered trademark of Intel Corporation  
\* All other products may be trademarks or registered trademarks of their respective companies.

## NOTES:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

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

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.