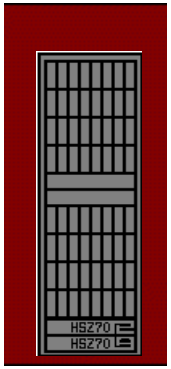

Overview of Fibre Channel Storage Area Networks

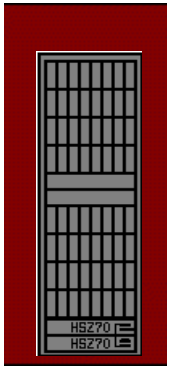

Challenges for SCSI

- Distance
  - 3 Meters SE
  - 25 Meters FWD
- Bandwidth
  - 40MB/sec UltraSCSI
  - 100MB/sec FCAL
- Throughput
  - 12000 IOs/sec UltraSCSI
  - 20000+ IOs/sec FCAL
- Scalability
  - 8 or 16 SCSI Ids
  - 128 nodes/FCAL
  - Virtually unlimited Switched Fibre

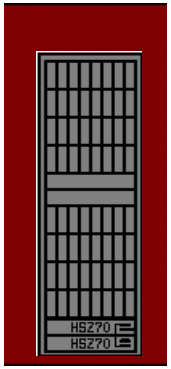

Fibre Channel - What is it?

- Integrated Set of Standards
  - Being developed by committees accredited by ANSI T11 committee (http://www.t11.org)
- New Protocol for Information transfer
  - Industry Standard interconnect
  - high performance serial I/O protocol
  - media independent
  - supports transfer of multiple protocols
- Information on technology and industry
  - http://www.fibrechannel.com

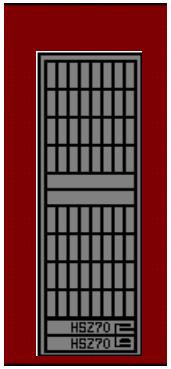

SANs - What are they?

- Storage Area Networks
  - dedicated networks that connect one or more systems to storage devices and subsystems
  - synonymous with fibre channel, however, fibre channel not a required component
    - any networking or serial SCSI technology can be used to create a SAN
  - extends the capabilities of server I/O channels by using standard networking technologies as a transport for high bandwidth data transfers.

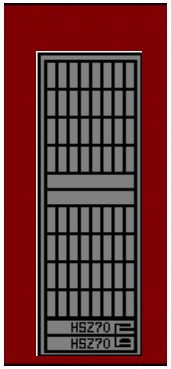

Distinguishing Fibre Channel Networks from SANs

- Fibre channel, a networking technology, was initially developed to be a high speed backbone technology for several purposes, including ip data networks.
- SANs are more an implementation of storage I/O methods over network transports (like fibre channel).
- SANs can be architected to provide the flexibility to implement the technology that fits best.
- Fibre channel is merely one of the transports chosen for a SAN implementation.

5

SANs provide flexible architecture

- ✓Enterprise Backbone
  - •Multiple switch fabric
- ✓Small clusters
  - •1 to 2 switches, 2 servers < 1TB
- ✓Departmental fabric
  - •3+ switches, > 2 servers, > 1TB
  - •Multi-Vendor

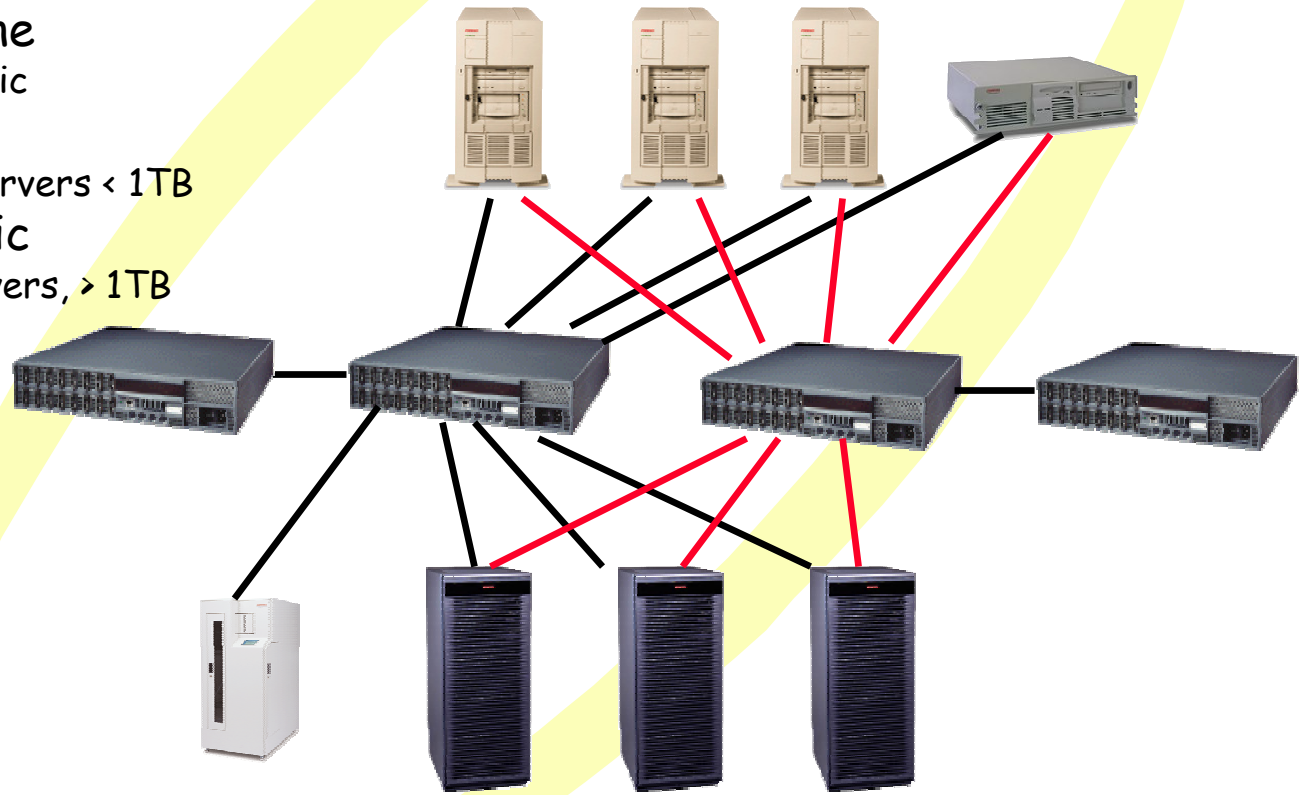

6

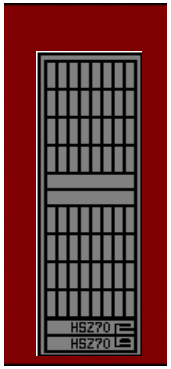

The Physical Fibre Channel Network

- Consists of ports, networking devices, and cabling
- Ports on systems and storage nodes are called N-ports or L-ports, depending on whether they're connected to switched or loop networks.
- Ports that work on both are called NL-ports
- Network devices are switches, hubs, bridges, and routers.
- Ports in networking devices are:
  - F-ports for N- port initiators
  - E-ports for switch to switch connections
  - FL-ports for connecting to loops

7

The Physical Fibre Channel Network

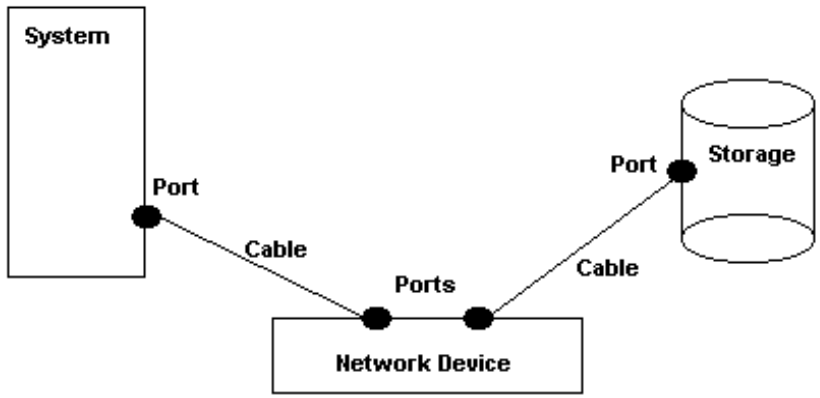

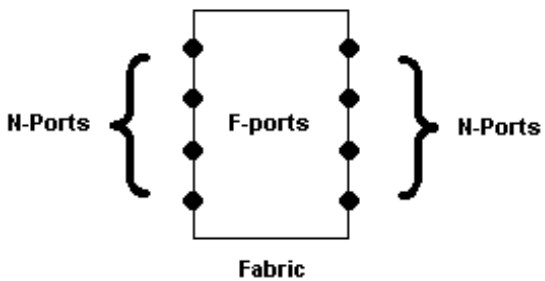

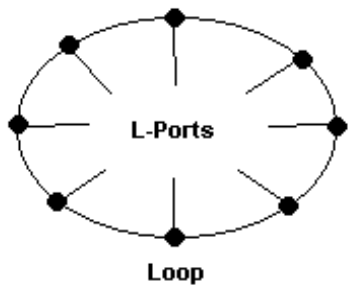

8

Fibre Channel Protocol Stack

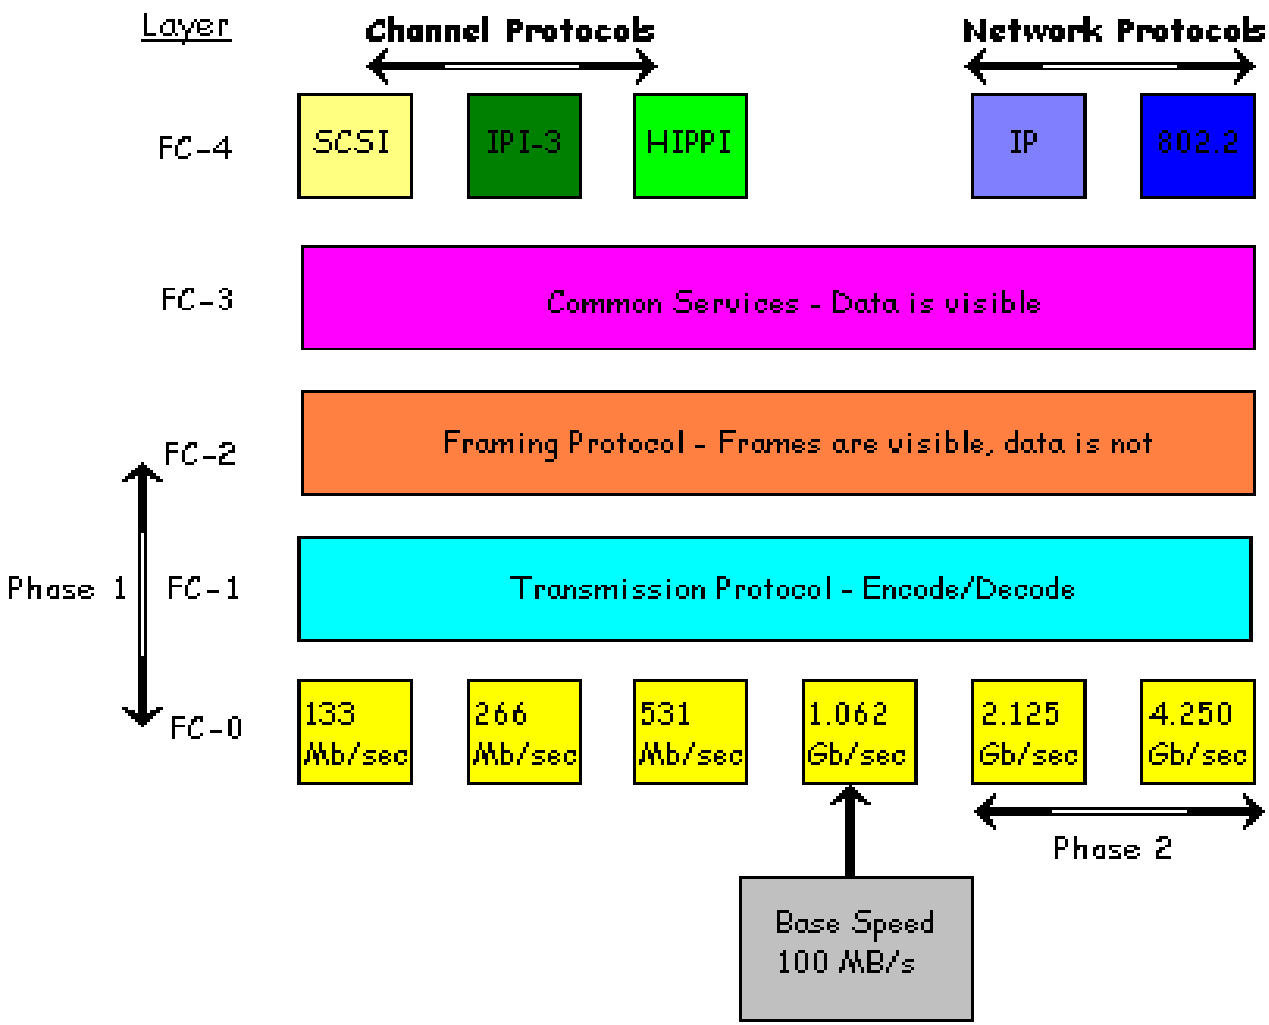

9

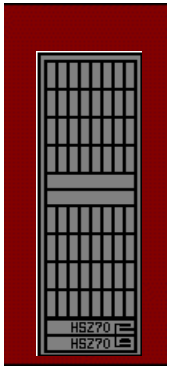

Protocol Stack

- FC-4 - Mapping
- FC-3 - Common Services
- FC-2 - Framing Protocol
- FC-1 - Encode / Decode
- FC-0 - Physical

10

Protocol Stack FC-0 (Physical)

Distance (m) at 100MB/s	Signal Source	Transmission Media	Source Wavelength	Fibre Core Diameter
10,000	Laser	SM Fiber	1300nm	9.0um
500	Laser/VCSEL	MM Fiber	780 to 850nm	50um
300	Laser/VCSEL	MM Fiber	780 to 850nm	62.5um
500	VCSEL	MM Fiber	980nm	62.5um
NA	LED	MM Fiber	1300nm	62.5um
25	ECL	Video Coax		
10	ECL	Mini Coax		
NA	ECL	Twisted Pair		

VCSEL = Vertical Cavity Surface Emitting Laser

NA - Cannot Achieve 100MB/s data rate

11

Fiber Optic Cable

50 / 62.5 micron Dia.

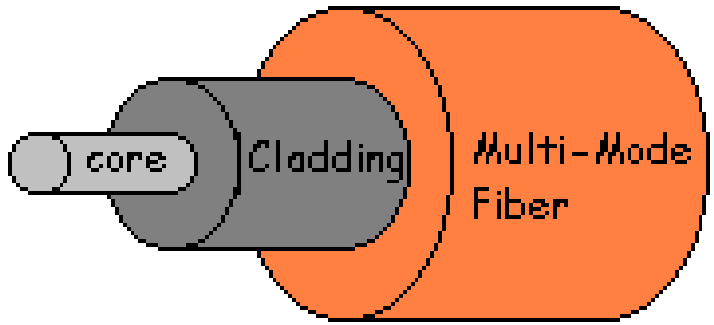

9 micron Dia.

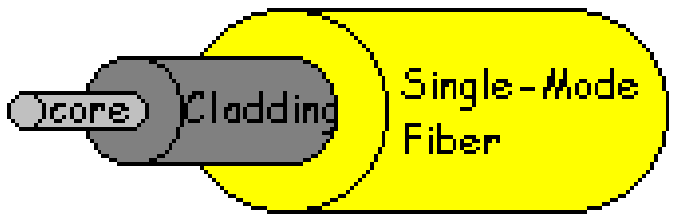

Connectors

Copper Connectors

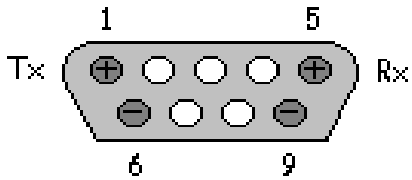

9 Pin STP Connector

BNC

Twisted Pair

Optical Connectors

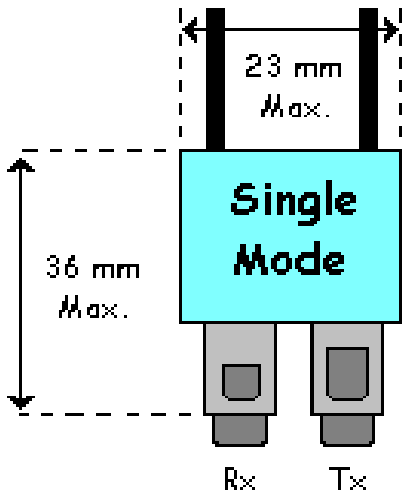

Single Mode SC Connector

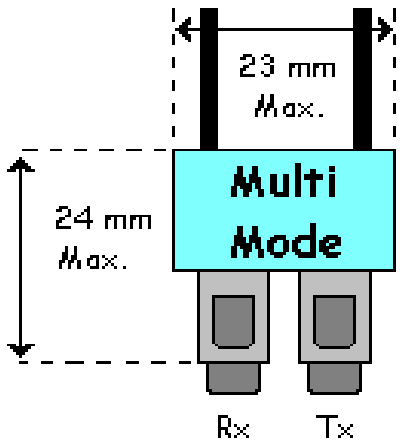

Multimode SC Connector

13

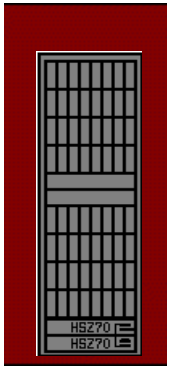

Transceivers

- Electrical
  - EL: ECL Level
- Optical
  - LL: longwave laser (1300nm)
  - SL: shortwave laser (780 to 850 nm)
  - LE: LED (1300nm)

The electrical signal is sent through a parallel-to-serial converter, converted into an optical signal, transmitted over the fiber media, changed back into an electrical signal and sent through a serial-to-parallel converter.

14

FC-1 Level -Encoding/Decoding

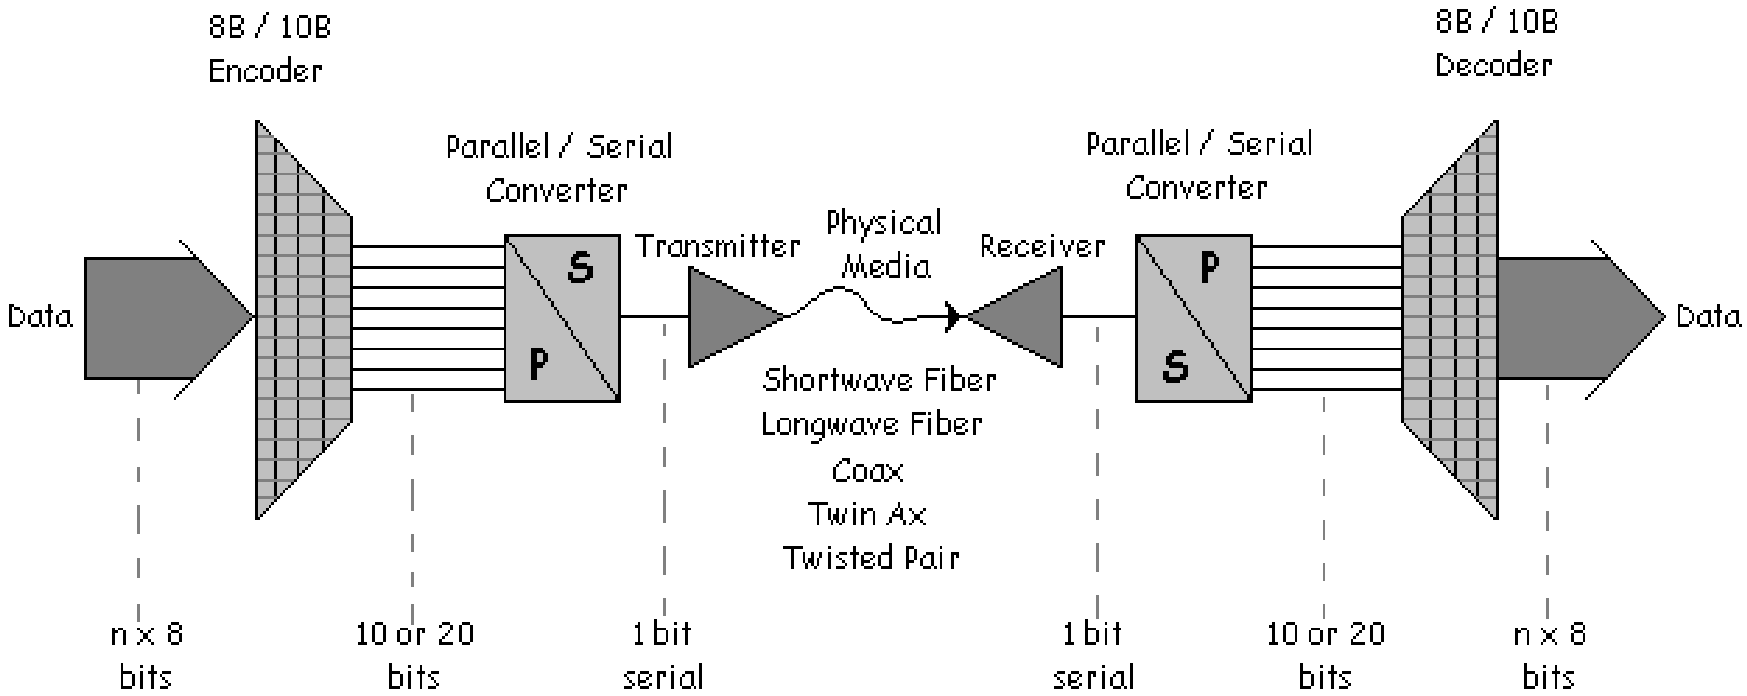

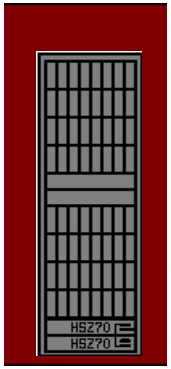

FC-2 Level - Framing

- A frame is a string of transmission words containing only data bytes,prefixed by a SOF (start of frame) delimiter and followed by an EOF (end of frame) delimiter.
- Several frames make a sequence, and several related sequences make an exchange. The following is the structure of each:
  - A *frame* is the smallest unit of information transfer
  - A *sequence* has at least one frame
  - An *exchange* has at least one sequence

SOF	Frame Header	Data Field	CRC	EOF
1TW	6TW	0-528 TW or 0-2112 Bytes	1 TW	1 TW

537 Transmission Words
or
2148 Bytes

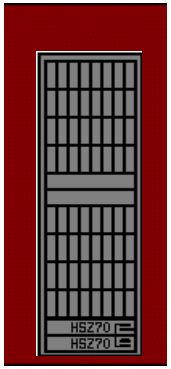

FC-3 Level - Common Services

- FC-3 is in the process of being defined. It is intended to provide the common services required for advanced features such as striping (to multiply bandwidth) and hunt groups (the ability for more than one port to respond to the same alias address). A hunt group can be referenced to a business that has 10 telephone lines, but requires only a single number to be dialed.

17

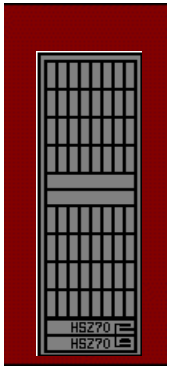

FC-4 Level - Mapping

- FC-4 defines the mapping of protocols between the lower levels of Fibre Channel and the command sets that use Fibre Channel.
- Here you will find separate standards for SCSI-3, IPI-3, HIPPI, FDDI, IP, and more.
- SCSI will retain its command set and a big portion of the device drivers and peripheral device codes, but instead of using 50-wire or 68-wire cables, connectors, and SCSI protocol chips, this lower level will be handled by Fibre Channel

18

SCSI-3 Architecture Structure

(Each box represents a separate SCSI-3 specification)

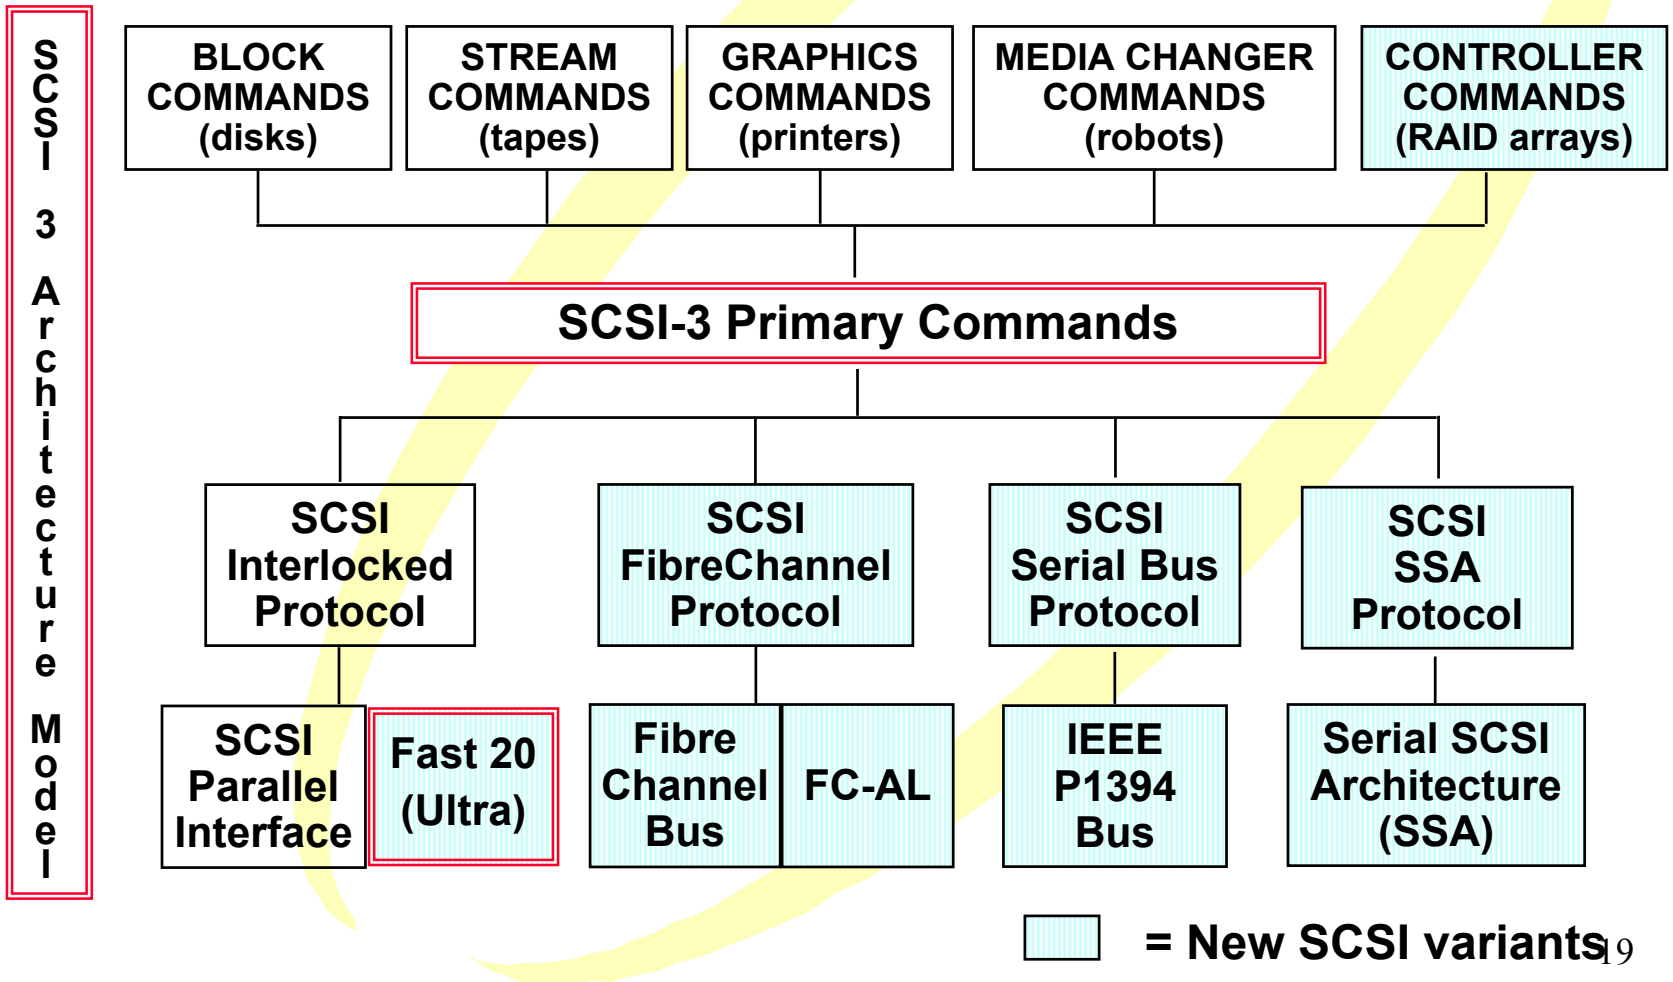

19

Optical Cables

- Multi-Mode Optical Cable Interconnect
  - Used with short wave optical transceivers, for distances up to 500 meters (300m for 2Gb)
  - Recommended as interconnect from hubs or switches to hosts and array controllers
  - In 2,5,15,30 and 50m lengths (longer cables must be custom made)
- Single-Mode Optical Cable Interconnect
  - Used with long wave optical transceivers, for distances of up to 10,000 meters
  - Fibre typically leased from telco supplier

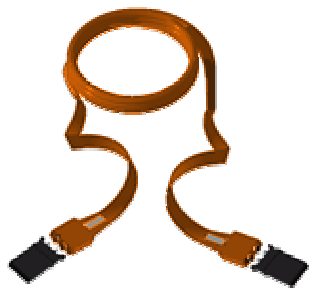

20

Adapters

- Enhanced PCI-to-FC adapter for selected systems
- Provides optical cable support with FC-AL or fabric switches

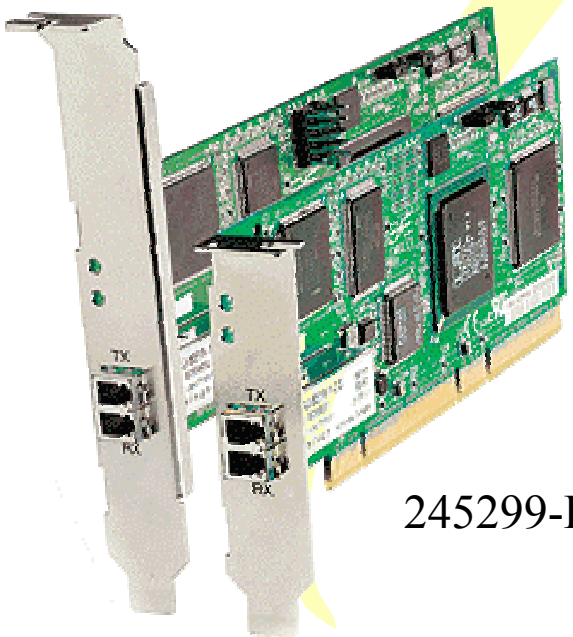

245299-B21/FCA2101

168794-B21/KGPSA-CA

21

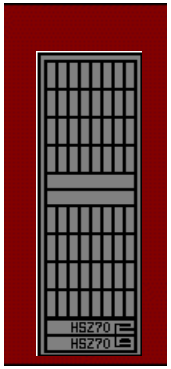

Transceivers (1Gb)

- Gigabit Interface Converter (GBIC)
  - Short-wave-based GBIC with SC cable connectors (380561-B21)
  - Long-wave-based GBIC with SC connector (127508-B21)
  - Hot-plug-able in optical hubs and switches.
- Gigabit Link Module (GLM)
  - modular media connectors
  - mount on HBA daughter card and HSG80 controller module
  - can be used to work with copper, long-wave and short-wave optical cables (HSG80 GLMs work with short-wave only)
- Connection Kits
  - FC Connection Kit
    3 short-wave GBIC's and 2 2m cables (380579-B21)
  - FC Connection Kit
    2 short-wave GBIC's and 2 2m cables (380596-B21)

Transceivers (2Gb)

- SFP Optical Transceivers
  - 221470-B21 (Brocade shortwave)
  - 300834-B21 (McData shortwave)
  - 300835-B21 (McData 10km transceiver)
  - 300836-B21 (McData 35km transceiver)
  - 332312-B21 (Cisco shortwave)
  - 332313-B21 (Cisco longwave)
- Optical cables (longer lengths available)
  - 221691-B21 (2m LC to SC)
  - 221691-B22 (5m LC to SC)
  - 221692-B21 (2m LC to LC)
  - 221692-B21 (5m LC to LC

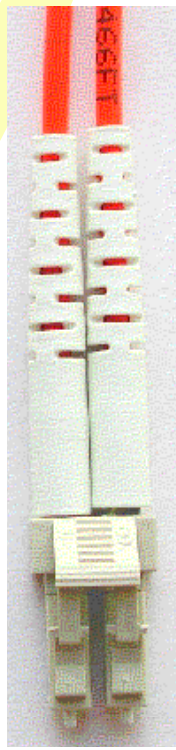

23

Fibre Channel Network Topologies

- Point-to-point topology

HSG80
N-Port GLM

KGPSA LP7000
N-Port GLM

* HP does not support using point-to-point on HSG80 controllers at this time

24

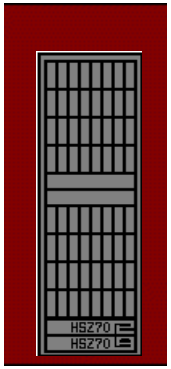

Fibre Channel Arbitrated Loop Topology

- FC-AL was developed to fill the gap between the limited access of point to point topology and the expensive switched fabric topology.
- Allows 126 node ports and 1 fabric port to be connected
- If more ports are connected to the loop - they are unable to be addressed
- The topology allows a single connection between one pair of ports at any point in time.
- Once a session (2 devices communicating) is started, the other devices on the loop must wait until the connection ends.
- Since all active ports on the loop share the loop's bandwidth, an individual port may only be able to recognize a small portion of the rated bandwidth.
- As a result, the total bandwidth available is limited to the bandwidth of the loop itself

25

Fibre Channel Arbitrated Loop Logical View

- In an arbitrated loop topology, the routing function is distributed to each loop port.

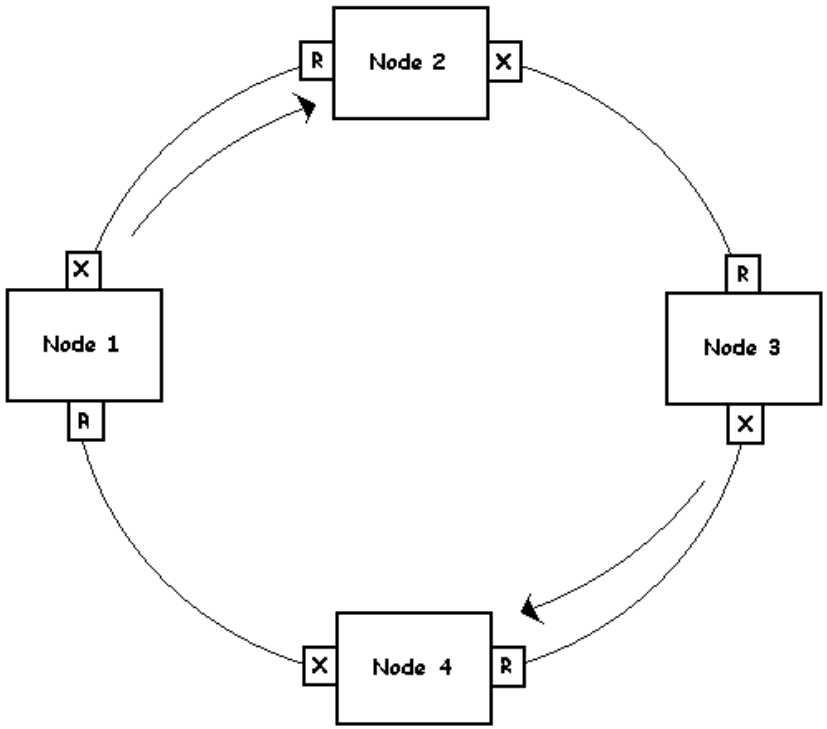

26

Fibre Channel Arbitrated Loop Physical View

- The Fibre Channel Storage Hub is the central device that routes transmissions around the arbitrated loop.

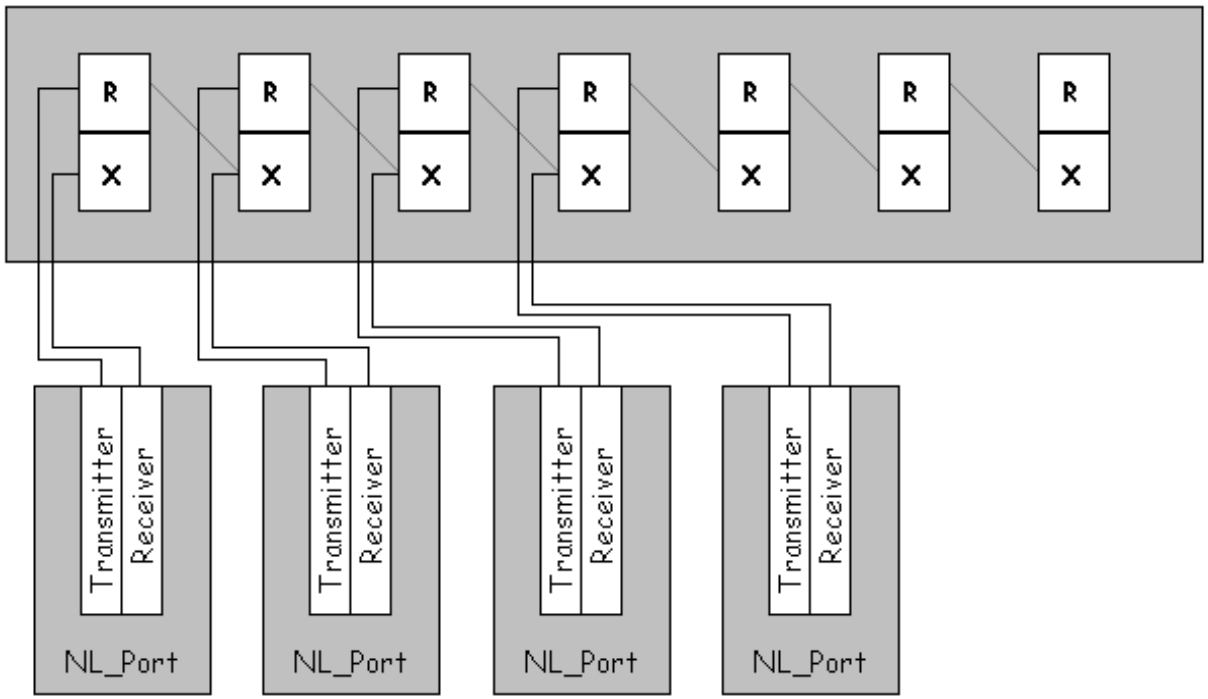

27

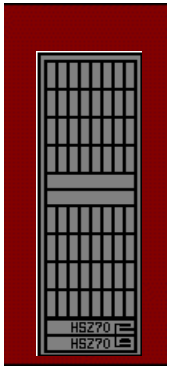

Fibre Channel Switched Topology

- Basis for Fibre Channel Fabric
- Provides high system bandwidth in multi-host systems
- Provides Name Server for node registration
- Requires Operating System support (more than FC-AL)
- Requires adapter, controller support (more than FC-AL)

28

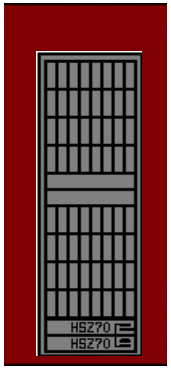

Switched Fibre Channel Fabrics

- Fibre channel networks built around switches are referred to as *fabrics*.
- The term fabric is used to distinguish the operations of the whole network as opposed to the operations of a single switch
- Ports on fabric networks connect nodes to switches on low-latency, point-to-point connections
- Each port has 2 complete 100MB/sec uni-directional connections that follow the same path on the network
- A node with multiple ports can have 200MB/s, 400MB/sec, 800MB/sec of aggregate bandwidth.
- To support this kind of bandwidth, fibre channel switches utilize high-speed integrated circuit and backplane technology that has the capability to support gigabit transmissions for multiple concurrent sessions

29

Switched Fibre Channel

- Each node on the Fibre Channel switch gets full bandwidth

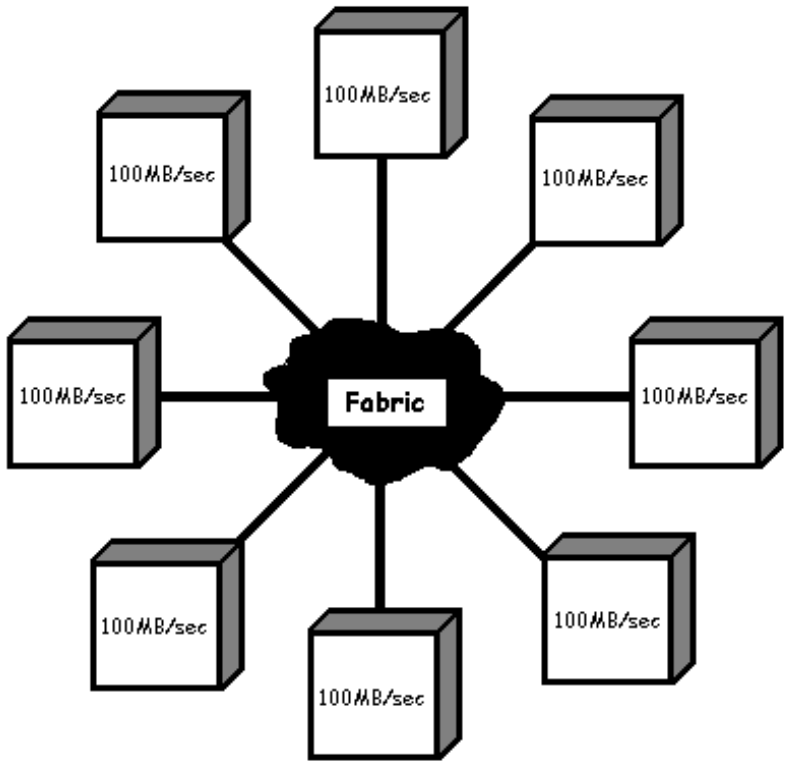

30

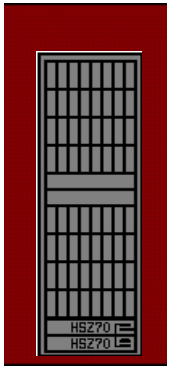

Addressing and Naming

- World Wide Name
  - 128 bit unique identifier (subsets implemented)
  - address stored in NVRAM
  - determined by IEEE
  - used to log into fabric
    - process called "registration" occurs
    - address is placed in simple name server
- Port Address
  - unique address for each port in network
  - 3 bytes long (24 bits)
  - defined by switch at fabric login
- Simple Name Server
  - directory service implemented in the switches
  - used by nodes, fabrics and applications for accessing port information

31

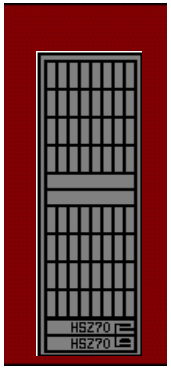

Zoning

- A technique used to segregate storage I/O traffic between groups of servers and their storage subsystems
- Implemented in the switch firmware
- Soft Zoning
  - filtering used to mask ports belonging to one zone from ports that don't belong
  - Can be assigned by WWN, port name, or alias name
- Hard Zoning
  - also known as port zoning
  - implemented in hardware
  - ports are physically blocked by port number

32

Zoning

- Hardware or Software restricted communications paths
- Can include/exclude any nodes in SAN

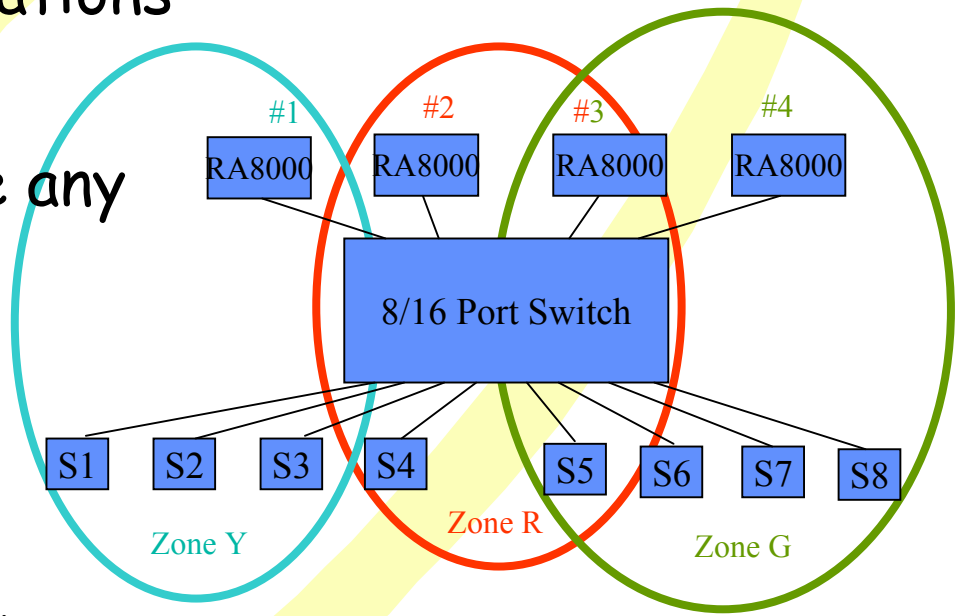

Zoning Illustration

Servers in a zone are restricted to accessing RA8000 in their zone.
As illustrated:

- S1,S2,S3 can access RA8000 #1
- S4,S5 can access RA8000 #2 and #3
- S5,S6,S7,S8 Can access #3 and #4

33

Interconnect product lines (June 2003)

markets

enterprise -class

Mid-range

entry level


core 2/64


director 2/64


director 2/140


MDS-9509


SAN switch 2/32


edge 2/32


MDS-9216


switch 2/16


switch 2/8 EL


blade 2/8


edge 2/24

B-Series M-Series C-Series

34

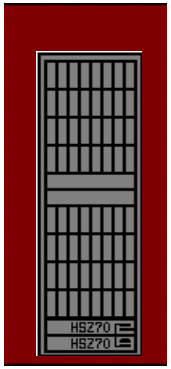

HP Interconnect product lines (SAN support)

	HP StorageWorks branded		Re-seller branded
O/S	HP-UX, Windows 2000, Windows DataCenter, Solaris, Netware, Linux, AIX, OpenVMS, Tru64 UNIX,	HP-UX, Windows 2000, Windows DataCenter, Solaris, Netware, Linux, AIX, OpenVMS, Tru64 UNIX	HP-UX, Windows 2000
Clustering	HP TruCluster, OVMS cluster, HP MC/ServiceGuard, HACMP lifekeeper, MSCS netware clusters, SUN clusters, Veritas clusters	HP TruCluster, OVMS cluster, HP MC/ServiceGuard, HACMP lifekeeper, MSCS netware clusters, SUN clusters, Veritas clusters	
Arrays	MSA 1000, EVA, EVA v3, XP, VA, MA8000/EMA12000,	MSA 1000, EVA, EVA v3, XP, VA, MA8000/EMA12000,	EVA v2, XP, MA8000/EMA12000
SW & Applications	EBS, DRM, Secure Path, Autopath, OVSAM, Continuous access, CASA	EBS, DRM, Secure Path, Autopath, OVSAM, Continuous access	
Scaling	28 switches, 1280 total ports	24 switches, 1632 total ports	Single Switch 192 ports
IP support	iSCSI bridging, FCIP SAN extension, NAS/SAN fusion	NAS/SAN fusion (by June)	Coming in Q3 2003

B-Series	M-Series	C-Series
• Market share leader • Mature technology • Full integration	• Market leader HA director product • Preferred solution in EMC environments •Full integration	• Evolutionary integration path with HP products • Newest technology entry in the switch market

B-Series product line

- Common 2Gbit/sec Fibre Channel ASIC
- Common Advanced Fabric Services
- Common Open Management
- Compatible with entire SAN Switch family
- >1000 total port SAN topologies & Growing
- **Software Power Pack** options on all switches (except embedded blade)

SAN core 2/64

BROCADE

SAN switch 2/32

SAN switch 2/16
SAN switch 2/16 EL

SAN switch 2/8
SAN switch 2/8 EL

msa SAN switch 2/8
(Embedded in
HP MSA1000 only)

36

M-Series product line

hp StorageWorks edge switch 2/24

- departmental /edge level connectivity
- capacity on demand (Flexport)
- 8/16/24 port configurations
- call-home features
- high performance
- firmware v4.01.02
- Optional software: HA Fabric Manager & SANtegrity Binding security software

hp StorageWorks edge switch 2/32

- departmental /edge level connectivity
- capacity on demand (Flexport)
- 16/24/32 port configurations
- call-home features
- high performance
- firmware v4.01.02
- Optional software: HA Fabric Manager & SANtegrity Binding security software

hp StorageWorks director 2/64

- full HA in a single box
- call-home features
- online upgrades – no impact to applications
- soft zoning
- firmware v4.01.02
- Optional software: HA Fabric Manager & SANtegrity Binding security software

hp StorageWorks director 2/140

37

C-Series product line (June 2003)

Future Entry Model

MDS 9216

- Includes 16 Port Fabric Switch with single expansion slot
- Expandable up to 48 FC ports with no disruptions
- Supports one,16 or 32 port FC expansion module
- Supports FC/IP 8 Port expansion Module Available Q3 2003
- Management software included

MDS 9509

- Director-Class Switch with 9 expansion slots
- Fully Redundant hot swappable hardware
- Expandable up to 224 FC ports
- Supports up seven,16 or 32 port FC expansion modules
- Supports FC/IP 8 Port expansion Module Available Q3 2003
- Management software included

Future Director Model

38

StorageWorks SAN Switch 2/8-EL

- Key Features/Benefits
  - Eight High Performance auto-sensing 1 and 2 Gb ports.
  - May be optionally programmed to a fixed 1 Gb or 2 Gb speed if desired.
  - Universal, self-configuring ports (F, FL, E) with new industry-standard removable Small Form Factor Pluggable optical transceivers (SFP) (ordered separately).
  - Four Switch SAN fabric support.
  - Fully non-blocking for all 8 ports Full 16 Gb switching capacity for uncongested fully sustained, 2 Gb full duplex throughput.
  - One fixed power supply.
  - N+ 1 redundant fan.
  - Rack optimized design: 1U (1.75") tall for high density rack installations.
  - SAN Switch 2/8-EL includes Web Tools, Advanced Zoning, and Quickloop.
  - Part number: 322120-B21
  - SAN Switch 2/8 PowerPak (322121-B21) includes: SAN Switch 2/8-EL, Fabric Watch, ISL Trunking, Advanced Performance Monitor, Remote Switch, Extended Fabric, full fabric, and Quickloop.

39

StorageWorks SAN Switch 2/16 (EL also available)

- Key Features/Benefits
  - Sixteen High Performance auto-sensing 1 and 2 Gb ports.
  - May be optionally programmed to a fixed 1Gb or 2 Gb speed if desired
  - Universal, self-configuring ports (F, FL, E) with new industry-standard removable Small Form Factor Pluggable optical transceivers (SFP) (ordered separately)
  - Fully non-blocking for all 16-ports Full 32 Gb switching capacity for uncongested fully sustained, 2 Gb full duplex throughput
  - High availability design
  - Redundant, hot swappable power supplies
  - Hot-swappable cooling unit
  - Refer to the SAN Design Guide Addendum for 2 Gb switch configuration information
  - Rack optimized design: 1U (1.75") tall for high density rack installations
  - Part number: 322118-B21
- PowerPack bundle (322119-B21) includes:
  - QuickLoop, Remote Switch, Fabric Manager, Extended Fabric, ZBun

StorageWorks SAN Switch 2/32

- Key Features/Benefits
  - Delivers 32-ports in a 1.5U enclosure and up to 384-ports in a single 42U cabinet, facilitating manageable SAN fabrics composed of thousands of ports
  - Meets enterprise level availability requirements with redundant, hot pluggable components, no-single-points-of-failure within the switch
  - Provides 1 Gb and 2 Gb operation today
  - Employs optional Inter-Switch Link (ISL) Trunking to provide a high-speed data path between switches
  - Redundant hot swappable N+ 1 cooling Fans (three units Standard)
  - Redundant, hot swappable power supplies (two units Standard)
  - Part Number: 240603-B21
- PowerPack bundle (333764-B21) includes:
  - QuickLoop, Remote Switch, Fabric Manager, Extended Fabric, ZBun

StorageWorks Core Switch 2/64

- Key Features/Benefits
  - New Core Switch base model switch which ships with Advanced Zoning and Web Tools
  - 2 Gb up to 64-port connectivity
  - Investment protection for previous SAN Switch installations
  - Full suite of software tools offers the most comprehensive software bundle for maximum performance and complete SAN management
  - Easily integrated into core-to-edge SAN Switch configurations
  - Investment protection with backward compatibility with presently installed 1 Gb and 2 Gb SAN
- PowerPack bundle (332178-B21) includes:
  - Fabric Watch, Advanced Performance Monitor, Remote Switch, Extended Fabric, ISL Trunking

42

StorageWorks edge switch 2/16

- Outstanding support for mission critical applications by providing a rich set of high availability features
- Key Features/Benefits
  - Hot plug redundant power supply
  - Hot plug redundant fans
  - Hot plug optics
  - On-line diagnostics
  - On-line, non-disruptive firmware load and activation
  - Fault isolation tools for network-wide activity
  - Call-home and e-mail automatic notification
- Optional Software
  - Product Manager License (300659-B21), Fabric Manager (287406-B21)
- Part number: 286811-B21

43

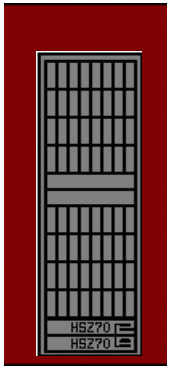

StorageWorks edge switch 2/32

- Outstanding support for mission critical applications by providing a rich set of high availability features
- Key Features/Benefits
  - Hot plug redundant power supply
  - Hot plug redundant fans
  - Hot plug optics
  - On-line diagnostics
  - On-line, non-disruptive firmware load and activation
  - Fault isolation tools for network-wide activity
  - Call-home and e-mail automatic notification
- Optional Software
  - Product Manager (300658-B21), 8-flexport upgrade (302660-B21), Fabric Manager (302660-B21)
- Part number: 286810-B21

44

StorageWorks Director 2/64

- Ensures Business Continuance through 99.999% availability
- Key Features/Benefits
  - Redundancy of all major hardware components
  - Automatic failover of active components
  - Fully hot swappable components
  - Non-disruptive code load/code activation
  - Configurations start at 32-ports increasing up to 64-ports, in four port increments.
  - Call home capabilities
- Optional Software and components
  - Fabric Manager (302660-B21), 8-port module kit (300833-B21)
- Part number: 286809-B21 (32 port configuration)

45

StorageWorks Director 2/140

- Ensures Business Continuance through 99.999% availability
- Key Features/Benefits
  - Redundancy of all major hardware components
  - Automatic failover of active components
  - Fully hot swappable components
  - Non-disruptive code load/code activation
  - Configurations start at 64-ports increasing up to 140-ports, in four port increments.
  - Call home capabilities
- Optional Software and components
  - Fabric Manager (287406-B23), 4-port module kit (316094-B21)
- Part number: 316093-B21 (64 port configuration)

46

Cisco MDS 9216 Multilayer Fabric Switch

- Key Features/Benefits
  - Modular, scalable Fabric Switch – 2 slots; 1 optional line card
  - High port density – 16-48 Auto-Sensing 2/1 Gb Fibre Channel ports in a single chassis
  - Multi-protocol Support – Fibre Channel, iSCSI, FCIP
  - Virtual SANs (VSANs)
  - QoS
  - Embedded Diagnostics – Fibre Channel ping and trace route, protocol analysis and decoding, SPAN, Call Home capability
  - Security – RBACs, ACLs, Hardware-Enforced Zoning, FC-SP, SNMPv3, RADIUS
  - Open Platform for Enabling Network-based Intelligent Storage Applications
  - Integrated Management – embedded Fabric Manager, integration with HP OpenView -SAM and CiscoWorks RME
  - Part Number: 332315-B21

The Cisco MDS 9216 Multilayer Fabric Switch shares a consistent architecture with the MDS 9509 Director and offers the same multilayer intelligence in a modular fabric switch.

47

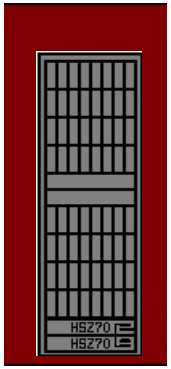

Cisco MDS 9509 Multilayer Director

- Key Features/Benefits
  - Modular, scalable Director – 9 slots
  - High port density – 16-224 Auto-Sensing 2/1 Gb Fibre Channel ports in a single chassis and 672-ports in a single rack
  - High Availability
  - Multi-protocol Support – Fibre Channel, iSCSI, FCIP
  - High Performance – 1.44 Terabit/sec internal system throughput ensure 10-Gbps readiness
  - Virtual SANs (VSANs)
  - PortChannels
  - QoS
  - Embedded Diagnositcs – Fibre Channel ping and trace route, protocol analysis and decoding, SPAN, Call Home capability
  - Security – RBACs, ACLs, hardware, enforced Zoning, FC-SP, SNMPv3, RADIUS
  - Open Platform for Enabling Network-based Intelligent Storage Applications
  - Integrated Management – embedded Fabric Manager, integration with HP OpenView -SAM and CiscoWorks RME
  - Part Number: 332306-B21 (base with 0 ports)

The MDS 9509 delivers industry-leading performance (1.44 Terabits internal system throughput), port density (up to 224 ports) and high availability, to lower TCO and enable integrated SAN infrastructures.

48

Summary of Fibre Channel Components