Powering Your Windows SAN With Emulex HBAs

Joe Teolis

Sr. Director, Product Marketing

HP Solution Symposium – Windows Track

EMULEX
We network storage

Emulex: Building on a Foundation of Strength

- 520 Employees
  - Costa Mesa, CA
  - Bothell, WA
  - Bolton, MA
  - Longmont, CO
- Worldwide manufacturing partners

- **Strong track record of growth**
  - ⇨ 5 Consecutive years of revenue growth
    - ⇨ FY03 Revenue of $308M
  - ⇨ #1 HBA Market Share
  - ⇨ 1.4M units shipped
- **Strategic investments fuel continued growth**
  - ⇨ Vixel acquisition
  - ⇨ Software technologies
  - ⇨ Intel/Emulex joint development

EMULEX
We network storage

Expanding Market Opportunities

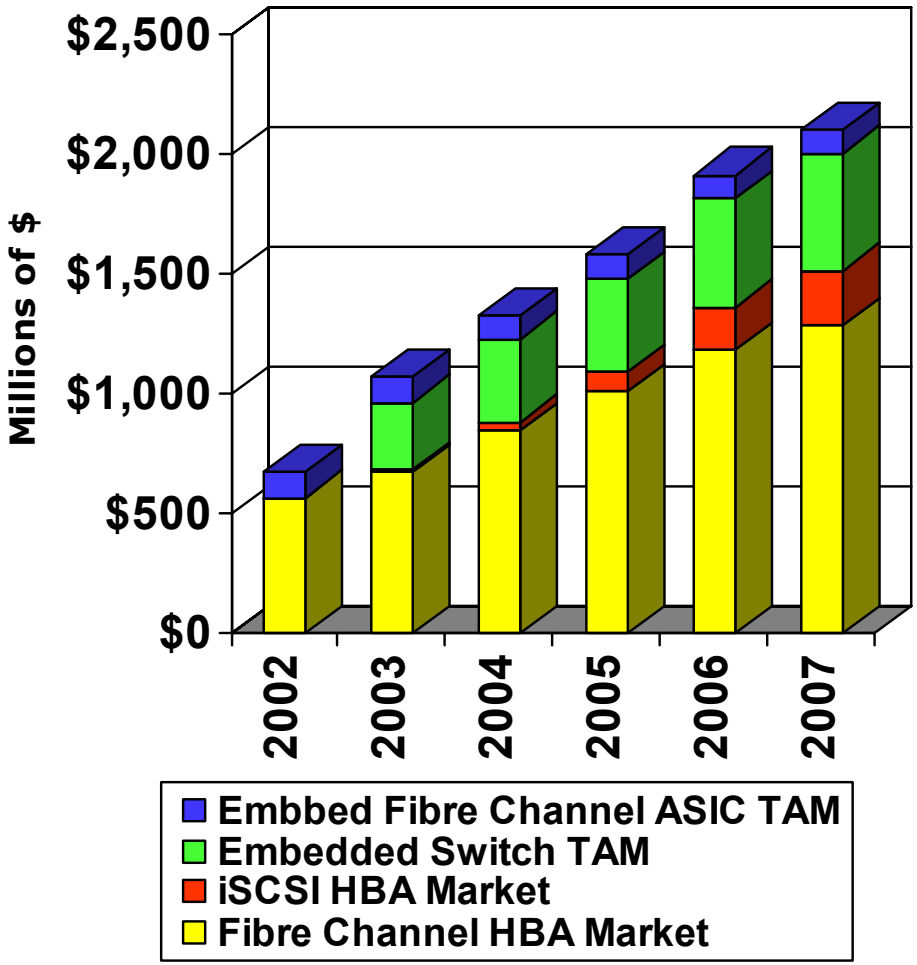

- Vixel acquisition adds strength in embedded switching market
  - ⇨ InSpeed technology ideally suited for non-fabric switching applications
    - ⇨ SBOD and Root Switches
    - ⇨ Server blades
    - ⇨ SMB SANs
- Growing opportunity in embedded Fibre Channel
  - ⇨ I/O Controllers and switch ASICs provide key synergies with OEM customers
- Driving Fibre Channel into the SMB Market
  - ⇨ Significant opportunity for easy-to-use SAN products

Source: IDC (July 2003), Gartner (July 2003), Vixel and Emulex

EMULEX
We network storage

EMULEX™ & Microsoft

We network storage

Providing World-Class Storage Network Connectivity

WHQL certified HBAs across Microsoft's server operating systems

Development, testing, and marketing

- ✓ **Storport, WMI, MPIO**
- ✓ **Storage, cluster, performance, SQL Server, SQL Exchange labs**
- ✓ **Exclusive HBA for Microsoft Systems Architecture (MSA)**
- ✓ **Executive Briefing Center (EBC), Microsoft Technology Centers (MTC)**
- ✓ **Microsoft Partner Solutions Center (MPSC)**

EMULEX

We network storage

Microsoft Systems Architecture (MSA)

Emulex is the exclusive MSA Fibre Channel HBA for Microsoft Systems Architecture – Internet Data Center

- Rigorously tested and validated technology architecture providing planning and implementation blueprints covering storage, networking, security, management, and application infrastructure.
- Designed to be used 2 ways:
  - ⇨ Proactive strategy to standardize IT architecture.
  - ⇨ Tactical approach to roll out specific projects.

Windows 2000 partners

Microsoft · EMULEX · hp invent · EMC² · UNISYS · BROCADE · Cisco Systems · CommVault Systems

Windows Server 2003 partners

Microsoft · EMULEX · hp invent · EMC² · DELL · HITACHI DATA SYSTEMS · UNISYS · NORTEL NETWORKS · Cisco Systems · BROCADE · McDATA

EMULEX
We network storage

Windows Overview

- Three production released drivers
  - FC Port driver (available for Windows Server 2003/2000/NT)
  - SCSIport Miniport driver (available for Windows Server 2003/2000/NT)
  - Storport Miniport driver (available for Windows Server 2003 only)
- Microsoft Windows Server 2003 In-Box drivers
  - Initial release - SCSIport Miniport driver v4.82a8 (32- and 64-bit versions)
  - SP1 Refresh - SCSIport Miniport v5.00a212 (Emulex v5.00a12 equivalent)
  - Future - switch to Storport Miniport driver
- All Emulex drivers WHQL qualified regularly
  - Windows Server 2003 and Windows 2000 adapter qualification
  - Windows Server 2003 Data Center qualification

Windows Driver Model Comparison
Miniport and Port Drivers

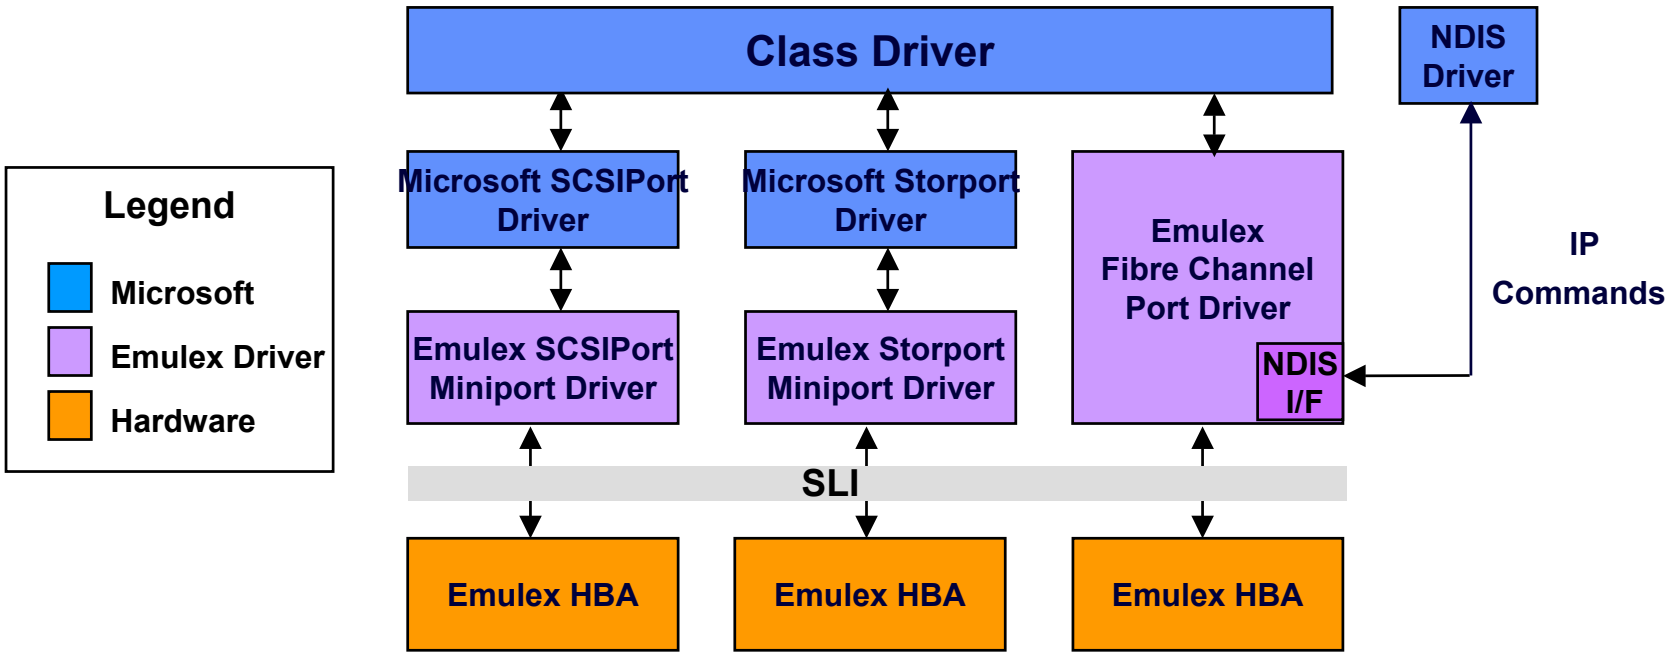

- Standard Microsoft Miniport models are portable across Microsoft operating systems
- Windows Server 2003 adds the Storport driver
  - ⇨ Fibre Channel and adapter-based RAID
- Emulex Fibre Channel Port driver coexists with Microsoft SCSI Port system driver
- Provides an NDIS socket

EMULEX
We network storage

Microsoft Storport

Emulex Storport Miniport Driver is written to Microsoft Storport standard

- **Microsoft Storport is a new port driver for Windows Server 2003 that delivers higher I/O throughput performance, enhanced manageability, and an improved miniport interface**
  - **Recommended for use with hardware RAID storage arrays and high performance Fibre Channel interconnects**
  - **Storport overcomes the limitations of the legacy SCSIport design, while preserving enough of the SCSIport framework that porting to the Storport device is straightforward for most developers.**
  - **Storport enables bidirectional (full duplex) transport of I/O requests, more effective interactions with vendor miniport drivers, and improved management capabilities.**
  - **Storport is the port driver of choice when deploying SAN or hardware RAID storage arrays in a Windows Server 2003 environment.**

EMULEX
We network storage

Microsoft Storport vs. SCSIport

	Storport	SCSIport
Function	Storport has been architected as a replacement for SCSIport, designed to enable the high performance capabilities of hardware RAID and Fibre Channel adapters to be realized.	Designed for managing SCSI block transport on parallel SCSI interconnects.
Adapter I/O Limit	> 254 outstanding I/O • 126 devices FCAL. • Up to 16 Million in FC Switched Fabric.	Max. 254 outstanding I/O request per adapter. Max. 15 attached devices in parallel SCSI.
I/O Function	Synchronous execution of I/O requests.	Sequential execution of I/O requests.
Effective Miniport Functioning	More effective miniport functioning • Circumvent the problem of delayed interrupt processing during miniport processing of the I/O • Eliminate the need for stepwise building of scatter/gather lists.	designed to do the majority of the work necessary to translate each I/O request packet into corresponding SCSI request blocks (SRBs).
Data Buffer Processing Overhead	Provides the flexibility to select the setting necessary to maximize the performance of the storage miniport driver. A Storport miniport can map all data buffers, no data buffers, or only those buffers that are not actual data intended for the application (such as discovery information).	Requires that the miniport driver repeatedly call the port driver to access the list of physical addresses that correspond to the IRP's data buffer information—one element at a time.
Queue Management	No limitation to the number of outstanding I/O requests.	Max. of 254 outstanding I/O request can be queued to an adapter.
Error and Reset Handling	Ability to instruct the HBA to only reset the afflicted LUN; no other device on that bus is impacted.	Respond using a SCSI bus reset. the code path will always disrupt I/O to *all* devices and LUNs connected to the adapter, even if the problem was related to only a single device.
Registry Access	Storport model allows registry access from the miniport in a much less restricted fashion. Writing back to the registry is also supported by Storport. This allows solving the problem of adapter specific parameters, such as persistent binding information or queue depth limits.	SCSIport does not allow free access to the registry from a miniport driver. SCSIport also cannot guarantee that multiple adapters using the same miniport will be able to use different sets of parameters. The total length of the parameter string passed is limited to 255 characters.

Emulex Advantages

EMULEX
We network storage

#1 HBA Supplier to HPQ

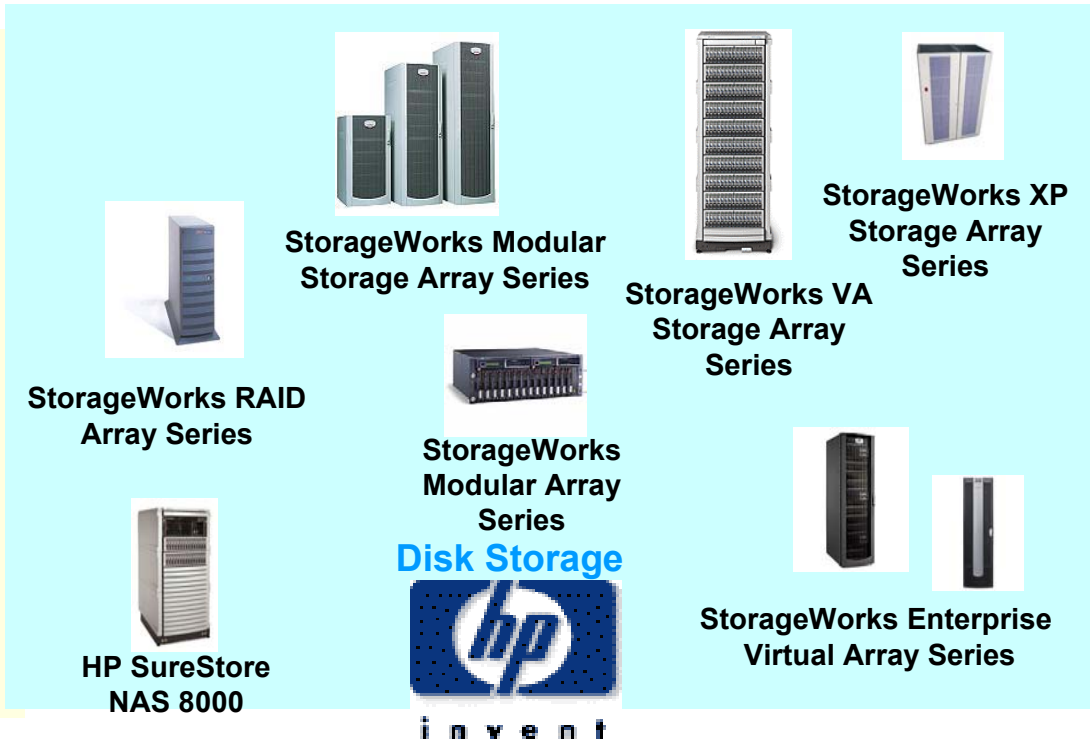

- ✓ Over 350,000 HBAs shipped to HPQ
- ✓ Support for strategic OSs: HP-UX, Windows, Linux
- ✓ HBAs and drivers for HP-UX: HP logo certified and hazard tested

EMULEX
We network storage

"Future-Proof Your Investment

Unique Emulex Capability:

- **Independent Firmware upgrade**
- **Upgrade firmware via management utilities**

- **Enables installed base to utilize "today's" features**
  - ⇨ **More than 1.4 M Emulex HBAs in service**
- **Upgrade without taking server offline**
- **Eliminates "fork-lift" upgrades**

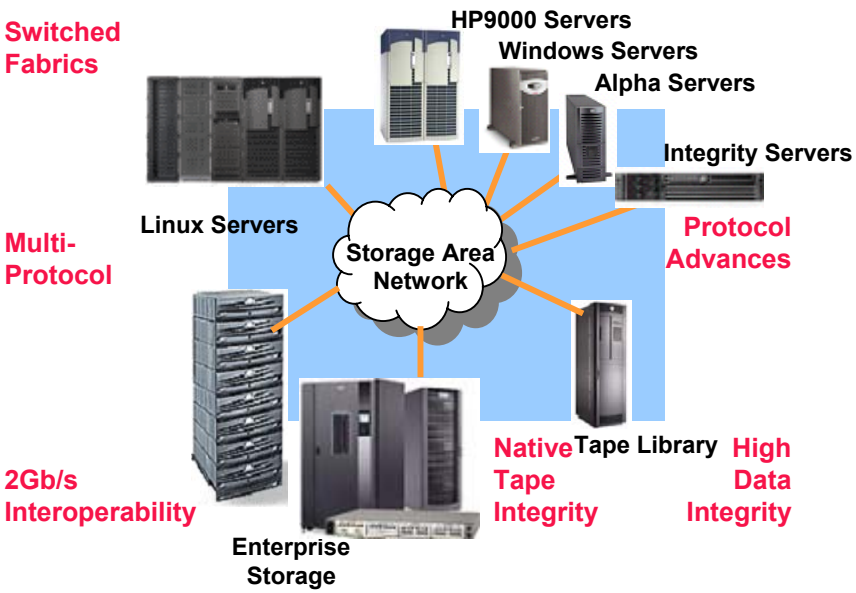

Lower Downtime
Simplified Management
Investment Protection

EMULEX
We network storage

Eliminate Driver Complexity

Unique Emulex Capabilities:

- **Hardware-Independent Driver Architecture**
  - ⇨ **One O/S driver for all products**

1G & 2G:

- PCI
- compact PCI
- low profile PCI
- PCI-X
- PCI-Express
- Sbus
- dual channel

- **Simplifies configuration and maintenance**
- **Enables seamless migration to new technologies**
  - ⇨ **Product Families**
    - ⇨ **LP7000, LP8000, LP9002L, LP9800, LP10000**
  - ⇨ **Technology Transition**
    - ⇨ **2Gb/s, 4Gb/s, PCI-X, PCI Express, 10Gb/s, … etc.**
- **Simplifies and accelerates OEM and partner design cycles and time to market**

Simplifies Management
Reduces Downtime
Lowers Qualification Costs

EMULEX
We network storage

#1 in Fibre Channel HBA Performance

PCI-X HBA *Transaction* Performance

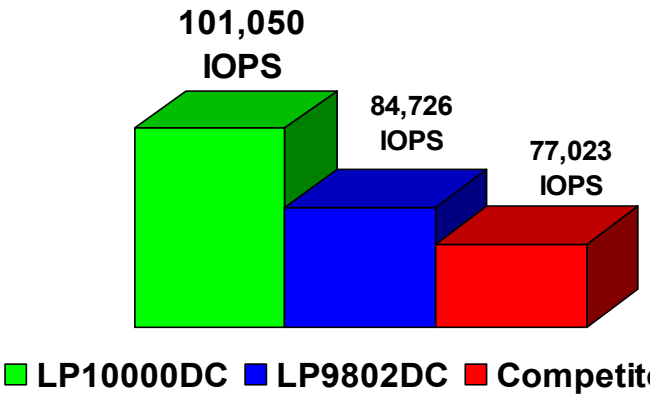

- 31% faster than competition
- 19% increase over previous generation HBA

PCI-X HBA *Throughput* Performance

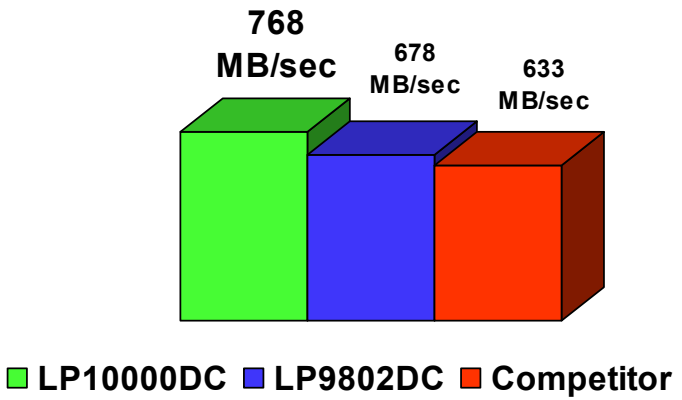

- 21% faster than competition
- 13% increase over previous generation HBA

Note: Transaction performance based on 2K block sequential READ IOPS in a MS Windows 2000 environment; throughput performance is based on 64K block read/write MB/sec in a MS Windows 2000 environment; data represents the total for both channels on dual channel HBA

EMULEX
We network storage

#1 HBA in Mixed Workload Performance

PCI-X HBA *Transaction* Performance

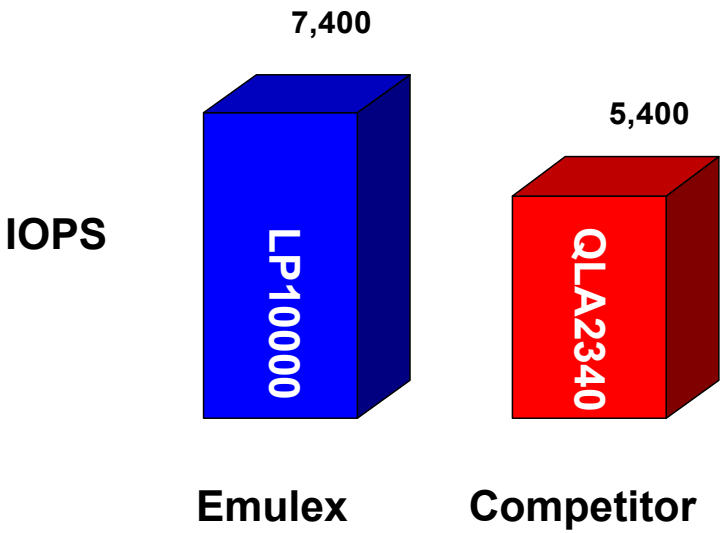

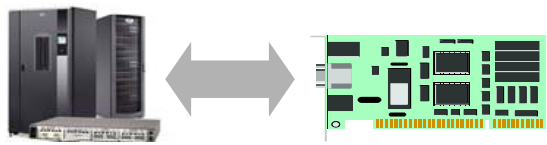

Tape and Disk Storage attached to one HBA

- **Concurrent transaction and backup workloads**
- **Transactions processed at frame rather than exchange level**
  - ⇨ **37% faster transaction performance**

Note: Mixed workload performance based on a workload that is split between 2K block READ transaction and 64K block WRITE backup in a MS Windows 2000 environment

EMULEX
We network storage

Enabling Technology

EMULEX
We network storage

Largest Buffer-To-Buffer Credits

Emulex provides 64 Buffer-to-Buffer credits (2K ea) on LP8/9000, 16 credits on LP850/950. Most others provide 4

Benefit:

- **Higher Performance in Distance Applications**
  - **Allow Devices to Negotiate Maximum Amount of Data In-Flight**
  - **Credits Represent Size of Buffer on HBA**

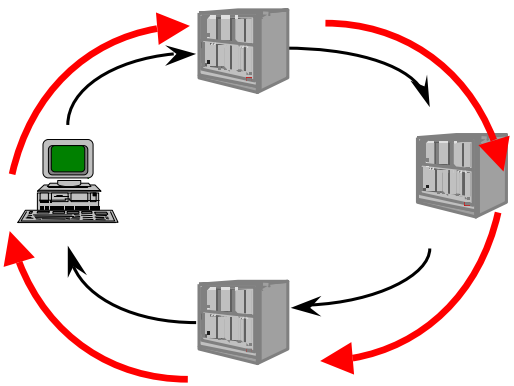

Loop Buffer to Buffer Distance = Entire Distance Between two Devices

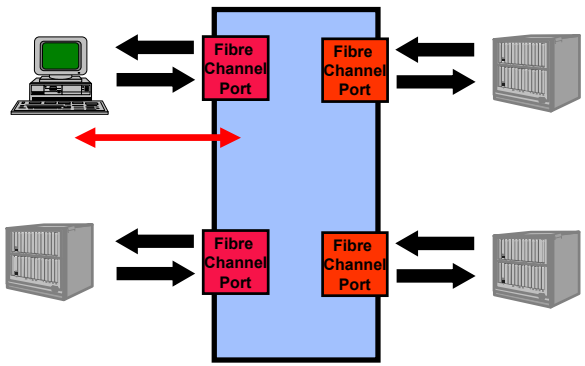

Fabric Buffer to Buffer Distance = Distance Between Device & Switch Port

EMULEX
We network storage

Hardware Context Cache

- Improved Receive Performance in Switched Fabric Applications
  - ⇨ Efficient hardware processing of interleaved exchanges
  - ⇨ Offloads Host PCI Bus CPU (contexts are not spilled/filled across the PCI Bus)
  - ⇨ Critical in Multi-Target Configurations
- Improved Scalability
  - ⇨ Able to process up to 2048 concurrent active exchanges

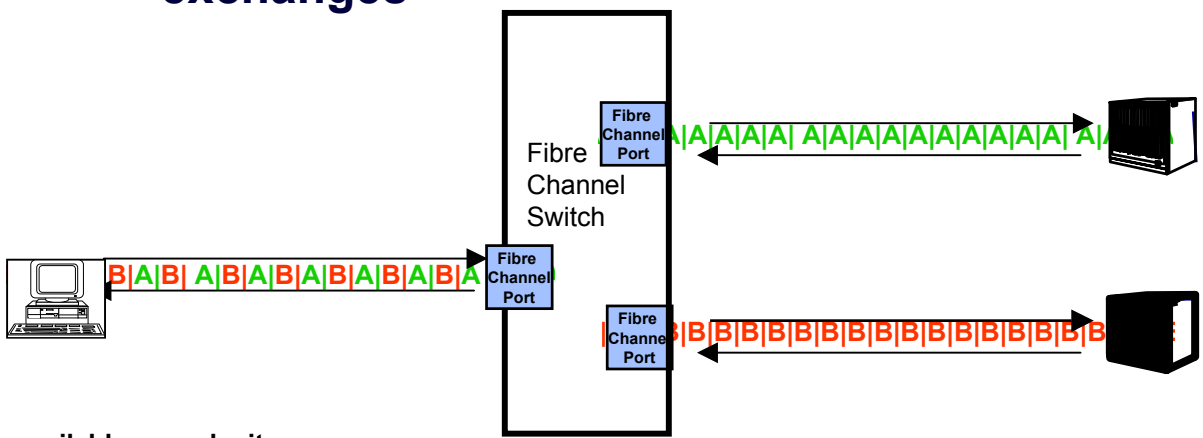

White paper available on web site

EMULEX
We network storage

Superior Product Design

Common drivers simplify HBA integration and management

Independent drivers and upgradeable HBA firmware protects investment

HBA platform continues to scale offering increased performance, new features, flexibility

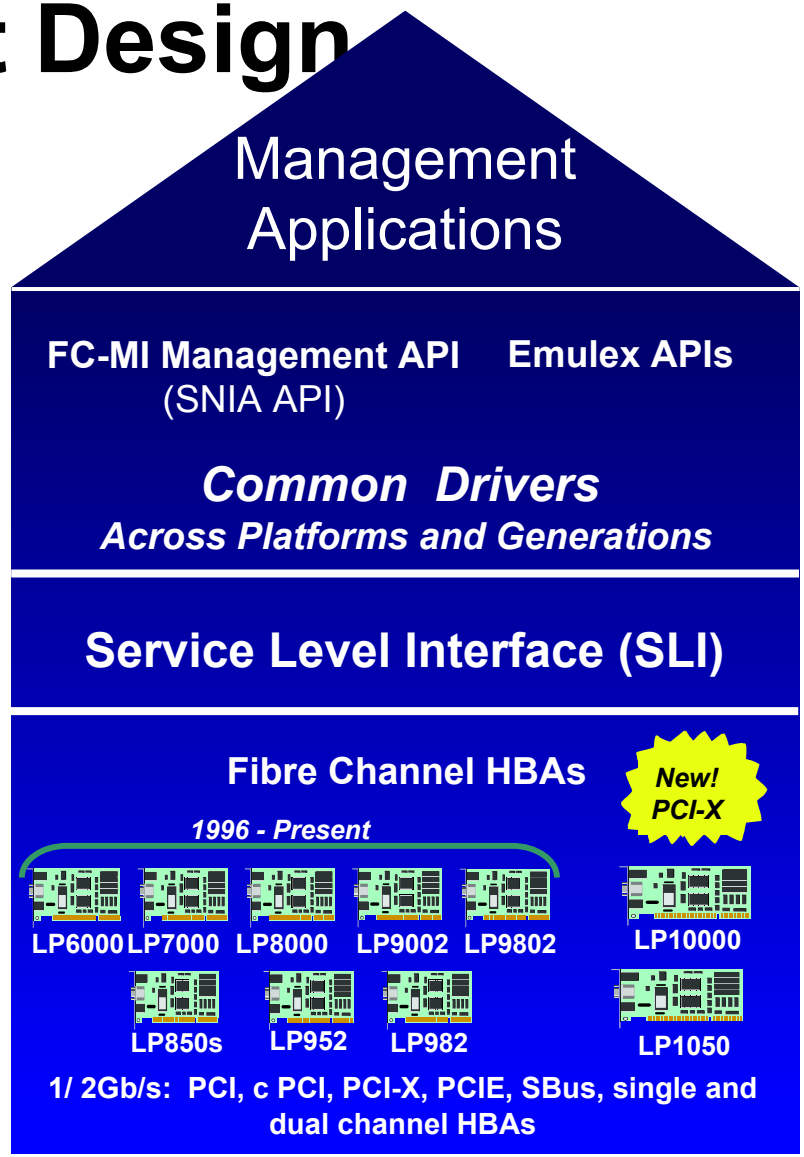

EMULEX
We network storage

Product Overview

- *Host Bus Adapters (HBAs)*
- *Management Applications*
- *Embedded Switching*

EMULEX
We network storage

HBAs Cross Referenced Matrix

Emulex Model	HP Model	Description	Targeted System & OS
LP8000	DS-KGPSA-CA/CX	1-Channel PCI	Tru64/OVMS
LP8000	DS-KGPSA-CB/CY	1-Channel PCI	Windows
LP952L	FCA2101	1-Channel PCI	NAS, Windows & Linux
LP9002L	FCA2354	1-Channel PCI	Tru64/OVMS
LP9002DC	FCA2355	2-channel PCI	Windows & Linux
LP982	A7298A	1-Channel PCI/PCI-X	Windows IPF/IA64 & Linux IA64 (Bulk pack)
LP982	FCA2408	1-Channel PCI/PCI-X	Windows
LP9802	AB232A	1-Channel PCI/PCI-X	Windows IA64
LP9802	336070-001	1-Channel PCI/PCI-X	NAS & IPF (Bulk Pack)
LP9802	FCA2384	1-Channel PCI/PCI-X	Tru64/OVMS, PCI-X for Alpha
LP9802	FCA2404	1-Channel PCI/PCI-X	Windows & Linux, PCI-X for IA32 servers
LP9802DC	FCA2404DC	2-Channel PCI/PCI-X	Windows
LP9002L		1-Channel PCI	HP-UX 11i “Certified for HP-UX”
LP9002DC		2-Channel PCI	HP-UX 11i “Certified for HP-UX”

EMULEX
We network storage

Fibre Channel HBA Product Line

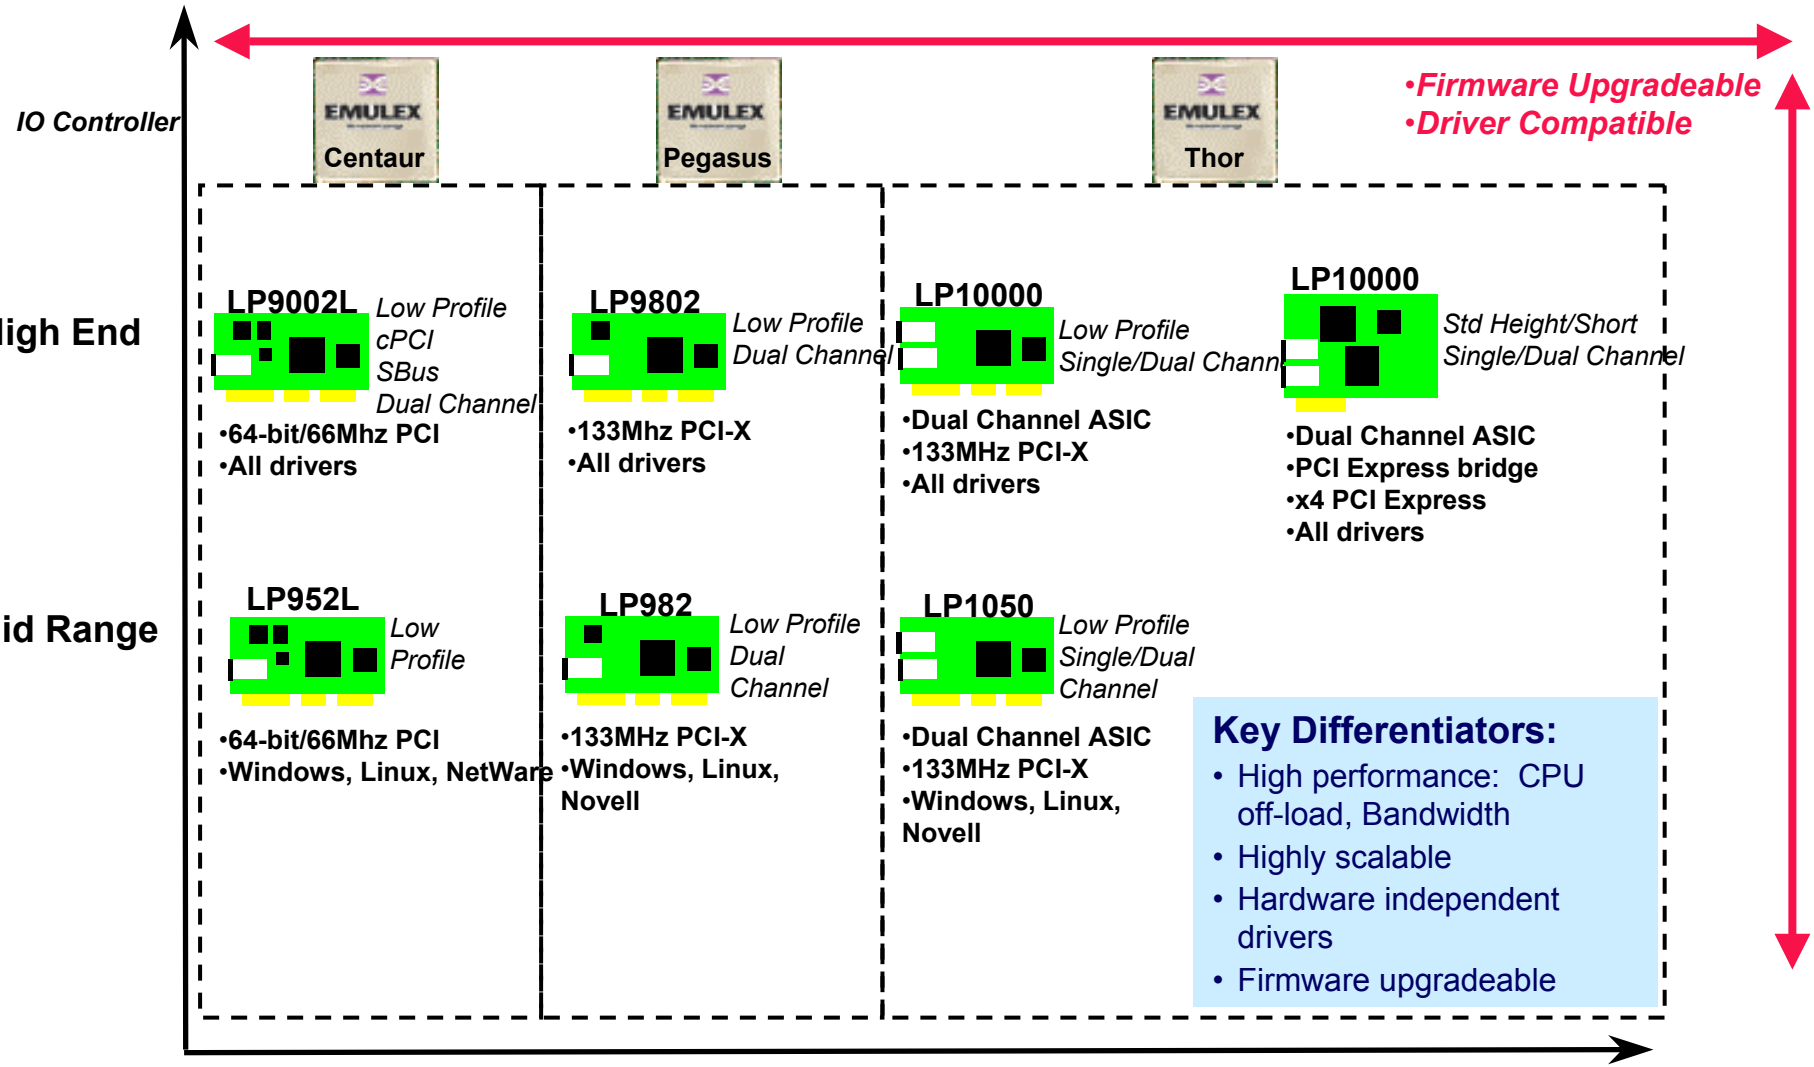

EMULEX
We network storage

LP9802DC Host Bus Adapter

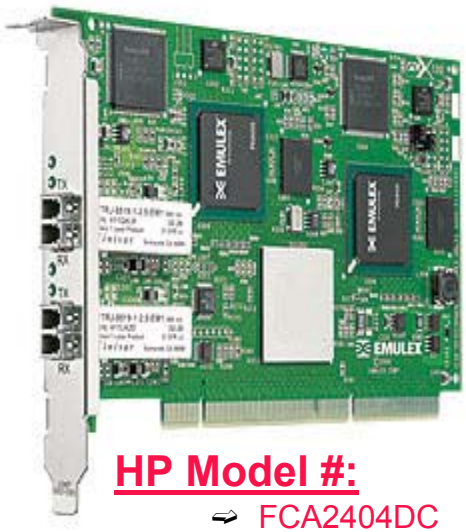

HP Model #:

- FCA2404DC

- 1 standard model
  - LP9802DC-F2

- 64-bit 66/100/133 MHz PCI-X 1.0a compatible
  - 33/66 Mhz PCI 2.2 compatible
  - 3.3 and 5 volts required
  - Dual channel short form factor PCI card
- Standards
  - INCITS FC-AL, FC-FLA, FC-MI, FC-PH, FC-PI, FC-PLDA
  - 1-2Gb/s, FCP, FCP-2, PHP Hot Plug, Class 2 & 3
- Software
  - Supports Windows NT/2000, Linux, NetWare, Solaris and HP-UX
  - Custom drivers through Emulex SLI
  - Supports cHBA API, HBAnyware, MultiPulse
- Architecture
  - Emulex Pegasus ASICs
  - 1024 concurrent on-board exchanges/channel
  - 16 buffer to buffer credits/channel
  - Full Speed PCI-X to PCI-X Bridge

EMULEX
We network storage

LP9802 Host Bus Adapter

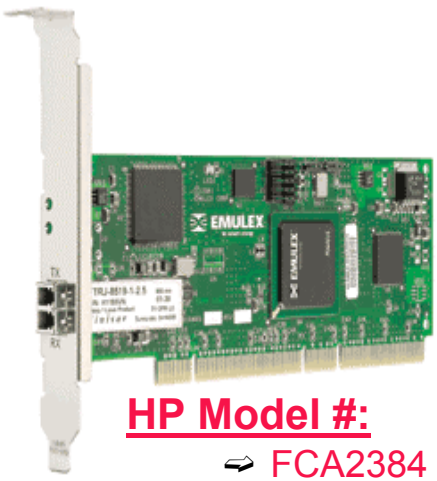

HP Model #:

- FCA2384
- FCA2404
- AB232A
- 336070-001

- 2 standard models
  - LP9802-F2
  - LP9802-X2

- 64-bit 66/100/133 MHz PCI-X 1.0a compatible
  - 33/66 MHz PCI 2.2 compatible
  - 3.3 VDC
  - Short form factor and low profile PCI card
- Standards
  - INCITS FC-AL, FC-FLA, FC-MI, FC-PH, FC-PI, FC-PLDA
  - 1-2Gb/s, FCP, FCP-2, PHP Hot Plug, Class 2 & 3
- Software
  - Supports Windows NT/2000, Linux, NetWare, Solaris, and HP-UX
  - Custom drivers through Emulex SLI
  - Supports cHBA API, HBAnyware, MultiPulse
- Architecture
  - Emulex Pegasus ASIC
  - 2048 concurrent on-board exchanges
  - 64 buffer to buffer credits

EMULEX
We network storage

LP10000 Family

2Gb/s Fibre Channel PCI-X HBAs

Qualification in Process by HP

- Features 20% higher transaction performance and up to 25% lower power than previous generation adapters
- Advanced Emulex HBA architecture
  - ⇨ Driver compatibility across product line
  - ⇨ High availability
  - ⇨ Scalability features
  - ⇨ Simplified SAN Management
- Single- and dual-channel models
  - ⇨ Low profile PCI (MD2) card format
  - ⇨ Enterprise and mid-range versions
  - ⇨ 3.3 volt PCI signaling, 5 volt tolerant

EMULEX
We network storage

LP10000
Enterprise-Class HBAs

- **Enterprise HBA features**
  - Features larger data buffering and on-board context caching for performance and scalability
  - Driver Support for Windows NT/2000/Server 2003, NetWare and Linux
  - Additional driver support for Solaris, HP-UX and custom drivers
  - Robust fabric boot support in x86, IPF and SPARC systems
  - M2 optics supports digital diagnostics interface enabling SAN management packages to monitor and manage critical link data
  - Improved initialization sequence, supporting all Sun dynamic reconfiguration modes

Emulex Exclusive

- **Single- and dual-channel models**
  - LP10000-M2: Single-channel, short-wave optical link
  - LP10000DC-M2: Dual-channel, short-wave optical link
- **Low profile PCI (MD2) card format**
  - Low profile mounting bracket available at no additional cost

EMULEX
We network storage

LP1050
Mid-Range HBAs

Qualification in Process by HP

- General Purpose HBA
  - Low cost adapters supporting Windows NT/2000/Server 2003, NetWare and Linux
  - Linux support now includes full-featured open source driver
  - On-card fabric boot support for x86 and IPF systems
- Single- and dual-channel models
  - LP1050-F2: Single-channel, short-wave optical link
  - LP1050DC-F2: Dual-channel, short-wave optical link
- Low profile PCI (MD2) card format
  - Low profile mounting bracket available at no additional cost for servers utilizing this type of slot

EMULEX
We network storage

HBAnyware™

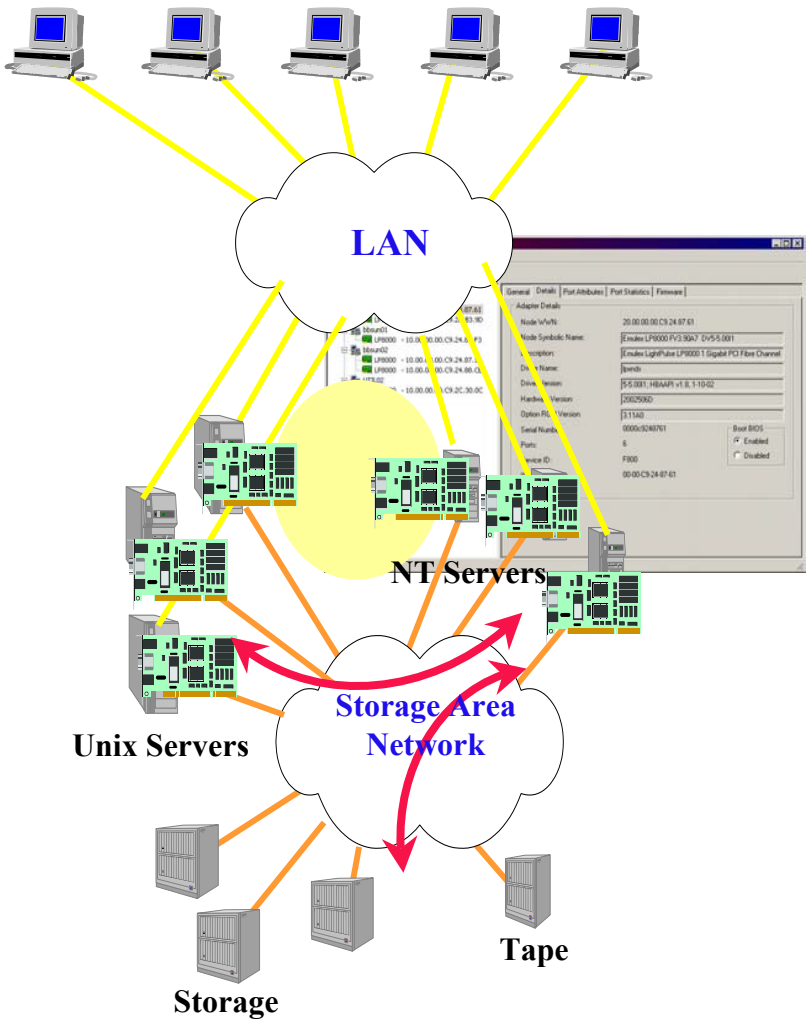

- Centralized HBA management suite designed to simplify and improve SAN management, and lower total cost of ownership for SAN based storage
  - Enables complete HBA management across heterogeneous SANs from a single console
  - Incorporates driver-based technology
- Product features
  - Remote Management API support for SAN management applications integration
  - Windows and Java GUIs and CLI provided
  - Included with standard driver kit
  - Real-time discovery of in-depth HBA data
  - Enables remote firmware upgrades
  - Industry-standard CT based in-band management
- Fully compatible with installed base of over 1,000,000 Emulex HBAs for easy deployment

EMULEX
We network storage

HBAnyware Feature Sets

HBAnyware 1.0

- Infrastructure
  - In-band communication based on Authenticated CT protocol
  - Fabric-wide discovery of Emulex HBAs
  - API strategy for integration into OEM and third-party solutions
    - RMAPI
    - CT extensions
- Management features
  - Real-time capture of in-depth HBA attributes and statistics
  - Remote firmware upgrade
- Standalone GUI utility and CLI
  - Cross-platform, heterogeneous
- Delivery
  - Packaged with standard driver kits
    - 32-bit Windows
    - Solaris 32/64
  - Production deployments since Q1-CY03

HBAnyware 2.0

- Infrastructure
  - Security provisions for secure in-band management
  - Advanced discovery service
    - Improved performance and scalability
    - Real-time changes
  - Extended APIs
- Management features
  - Driver configuration
  - Persistent bindings management
- Enhanced GUI utility and CLI
  - Showcase of framework capabilities
  - Basis for a new utility strategy
- Delivery
  - Expanded platform support to include
    - 64-bit Windows
    - Linux
  - Production deployments begin Q4-CY03

EMULEX
We network storage

MultiPulse™
Intelligent High-Availability Technology

- Unique driver-based technology
- Improves SAN reliability and performance
  - ⇨ Dynamic load-balancing and fail-over
  - ⇨ Improved data availability due to fast response time and intelligent re-routing
  - ⇨ Optimizes I/O traffic across multiple paths across the SAN
- Fully compatible with installed base
- Stand-alone solution or complimentary to OEM products

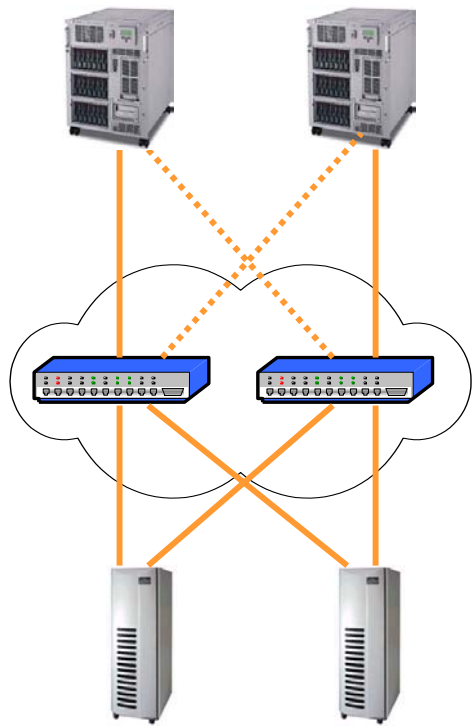

EMULEX
We network storage

InSpeed Technology Defined

Fibre Channel Switch-on-a Chip (SOC)

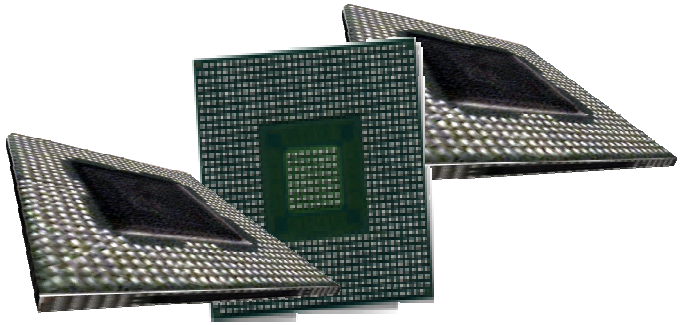

- Unique switching architecture designed from ground-up as an effortlessly managed interconnect solution
- Optimized for the needs and price requirements of storage and server supplies - breaking new cost/performance barriers
  - Increases systems uptime by 20%
  - Increases overall performance by up to 5 times
  - Reduces cost-per-megabyte of storage by 20%
  - Reduces cost of service by up to 15%
- No other technology offers integration at this level

EMULEX
We network storage

Modular RAID Today

Modular Storage Architecture is Shared

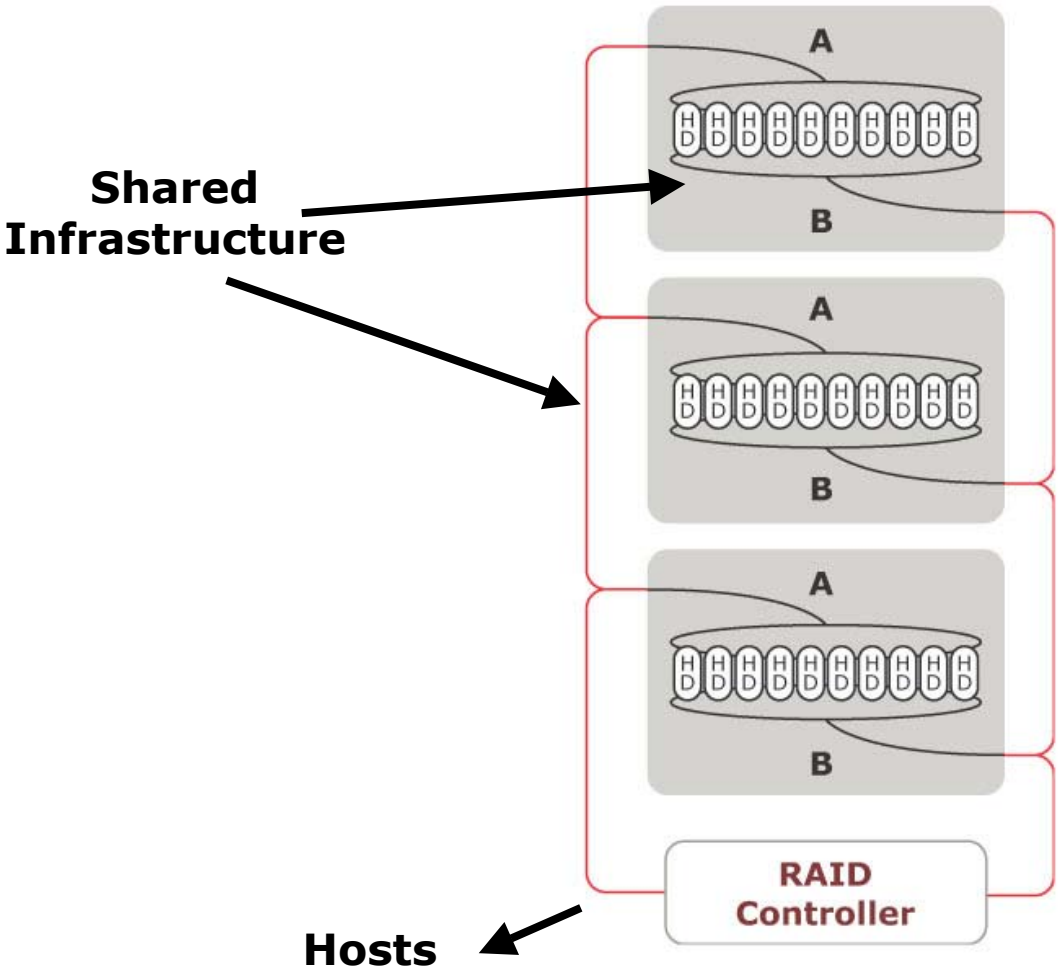

Challenges:

- Shared architecture
- Probability of double loop failures
- Serviceability
- Drive Isolation & Diagnostics
- Latency increases with # HDDs

Switched architectures are improving today's modular designs

EMULEX
We network storage

InSpeed Embedded Switch Solutions

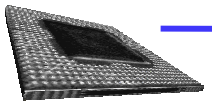

SBOD SOC:

"Switch the Disks"

- **Convert JBODs t SBODs**
- **Drive monitorin**
- **Drive isolation**
- **Performance**
- **Customers include Network Appliance, Fujitsu, Xyratex**

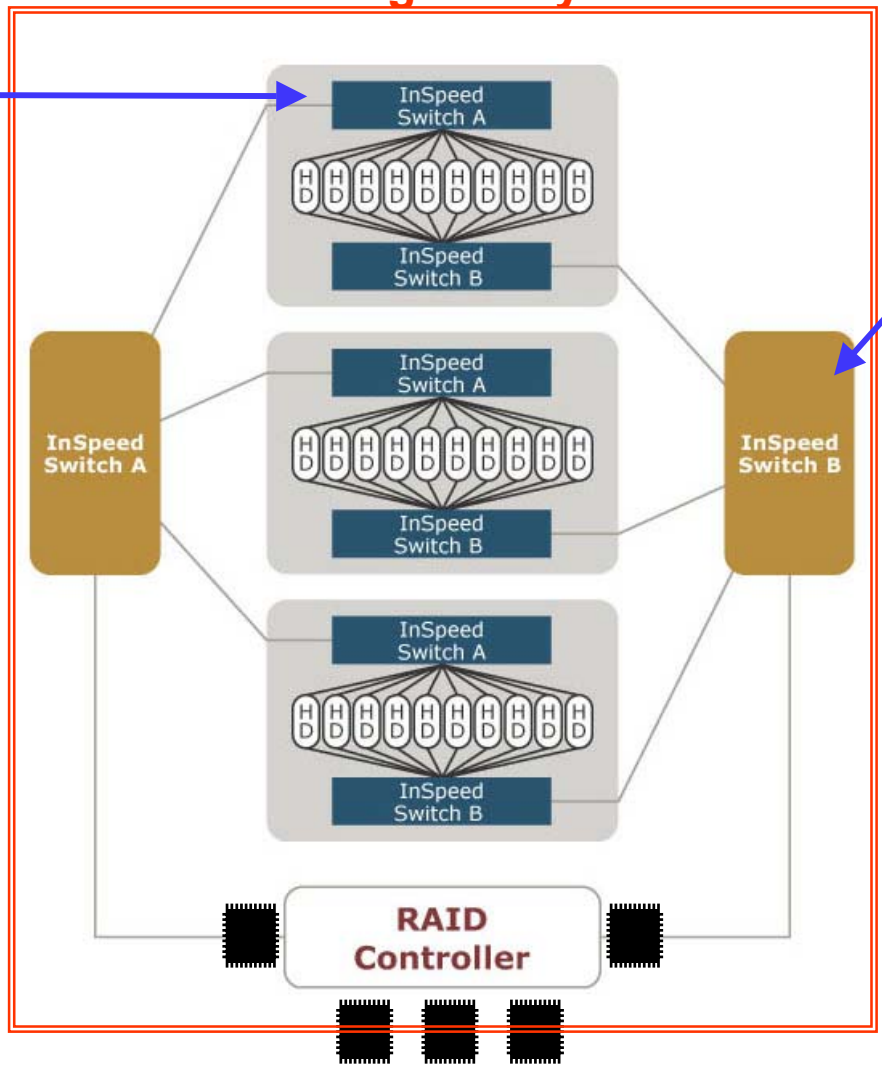

Root Switch:

"Switch the Drawers"

- **Direct data path**
- **Scalability**
- **Serviceability**
- **Performance**
- **Customers include Hewlett-Packard, BlueArc, Quantum**

EMULEX
We network storage

HP's EVA Deployment

InSpeed Back-end Root Switch Implementation

InSpeed Embedded Storage Switches Model 310

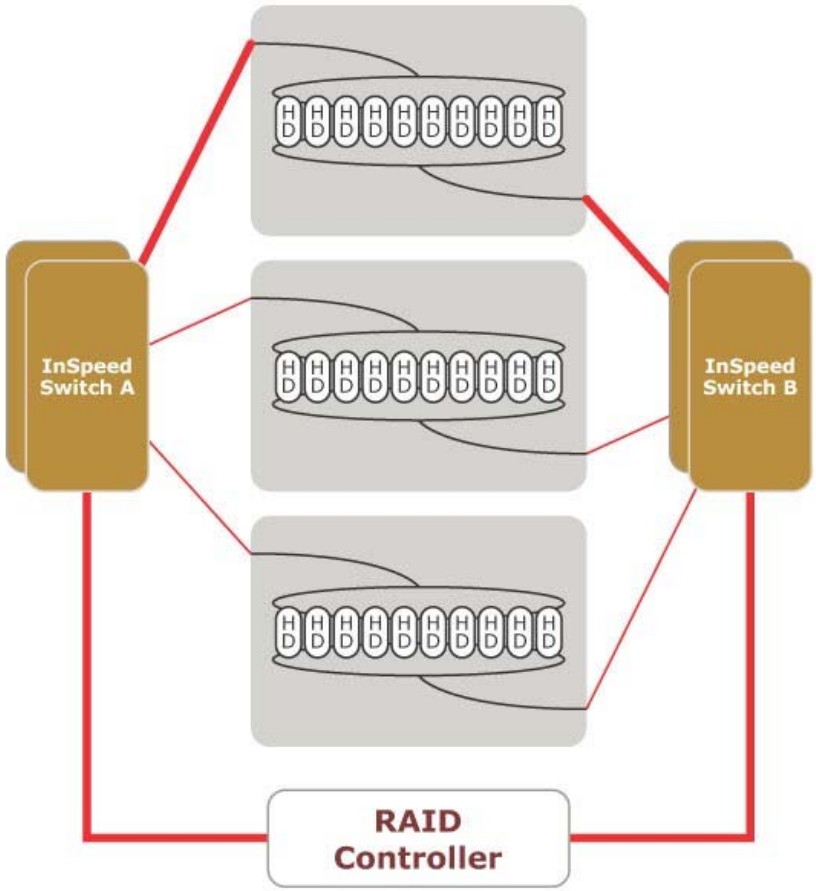

Advantages:

- Performance
  - Outperform competition by 50% in SPC benchmark
  - Turbo charged solution with up to 5x performance improvement
- Reliability
  - Increased RAS by isolating JBOD shelves
  - Decrease MTTR by field service
- Drop in solution – no code rewrite

EMULEX
We network storage

Product Qualification Matrix

EMULEX
We network storage

Servers & Storage Matrix

	HBA Model #	HBA Model#	Environment
 ProLiant Series (ML/DL and BL)	DS-KGPSA-CB/CY FCA2404DC FCA2408 FCA2355 FCA2404	LP8000 LP9802DC LP982 LP9002DC LP9802	Windows NT, 2000 and Server 2003 (IA32-bit)
HP 9000 Superdome Server   »rx5670  »rx4640 HP Integrity Series  »rx2600	AB232A A7298A	LP9802 LP982	Windows Server 2003 (IA 64-bit)
 StorageWorks Modular Storage Array Series   StorageWorks RAID Array Series StorageWorks XP Storage Array Series  StorageWorks Enterprise Backup Solutions (EBS)    StorageWorks Enterprise Virtual Array Series StorageWorks VA Storage Array Series	AB232A FCA2404 FCA2101 FCA2354 FCA2355 FCA2404DC FCA2408	LP9802 LP9802 LP952L LP9002L LP9002DC LP9802DC LP982	Windows NT, 2000 and Server 2003 (IA32-bit) Windows Server 2003 (IA 64-bit)

Related Product Resources

- **Emulex Webpage:**
  - Main Page
    - http://www.emulex.com
  - Products Page
    - http://www.emulex.com/products/fcswitch/index.html
  - Technical Whitepapers
    - http://www.emulex.com/products/white/index.html
  - Partners Page
    - http://www.emulex.com/partners/index.html
  - Online Training Page
    - http://www.emulex.com/training/index.html
- **HP Webpage**
  - Storage Products Page
    - http://welcome.hp.com/country/us/en/prodserv/storage.html
- **Microsoft Webpage**
  - Storport in Windows 2003
    - http://www.emulex.com/products/white/fc/Storport1.pdf

EMULEX
We network storage