Storage Options for Blade Server Environments

Rob Davis

Vice President of Advanced Technology
QLogic Corporation

HP WORLD 2003
Solutions and Technology Conference & Expo

Blade Server Storage Challenges

- What storage technologies to use?
  - Fibre Channel (FC), serial ATA (sATA), Ethernet (iSCSI), InfiniBand (IB), serial-attached SCSI (SAS)
  - NAS
- What storage connection architecture to use?
  - Single technology, dual technology
- What about next generation?
  - OS boot
  - 10 Gb / aggregation of bandwidth
  - Higher performance gateways
  - sATA and SAS fabric migration

Server Storage Technology Options

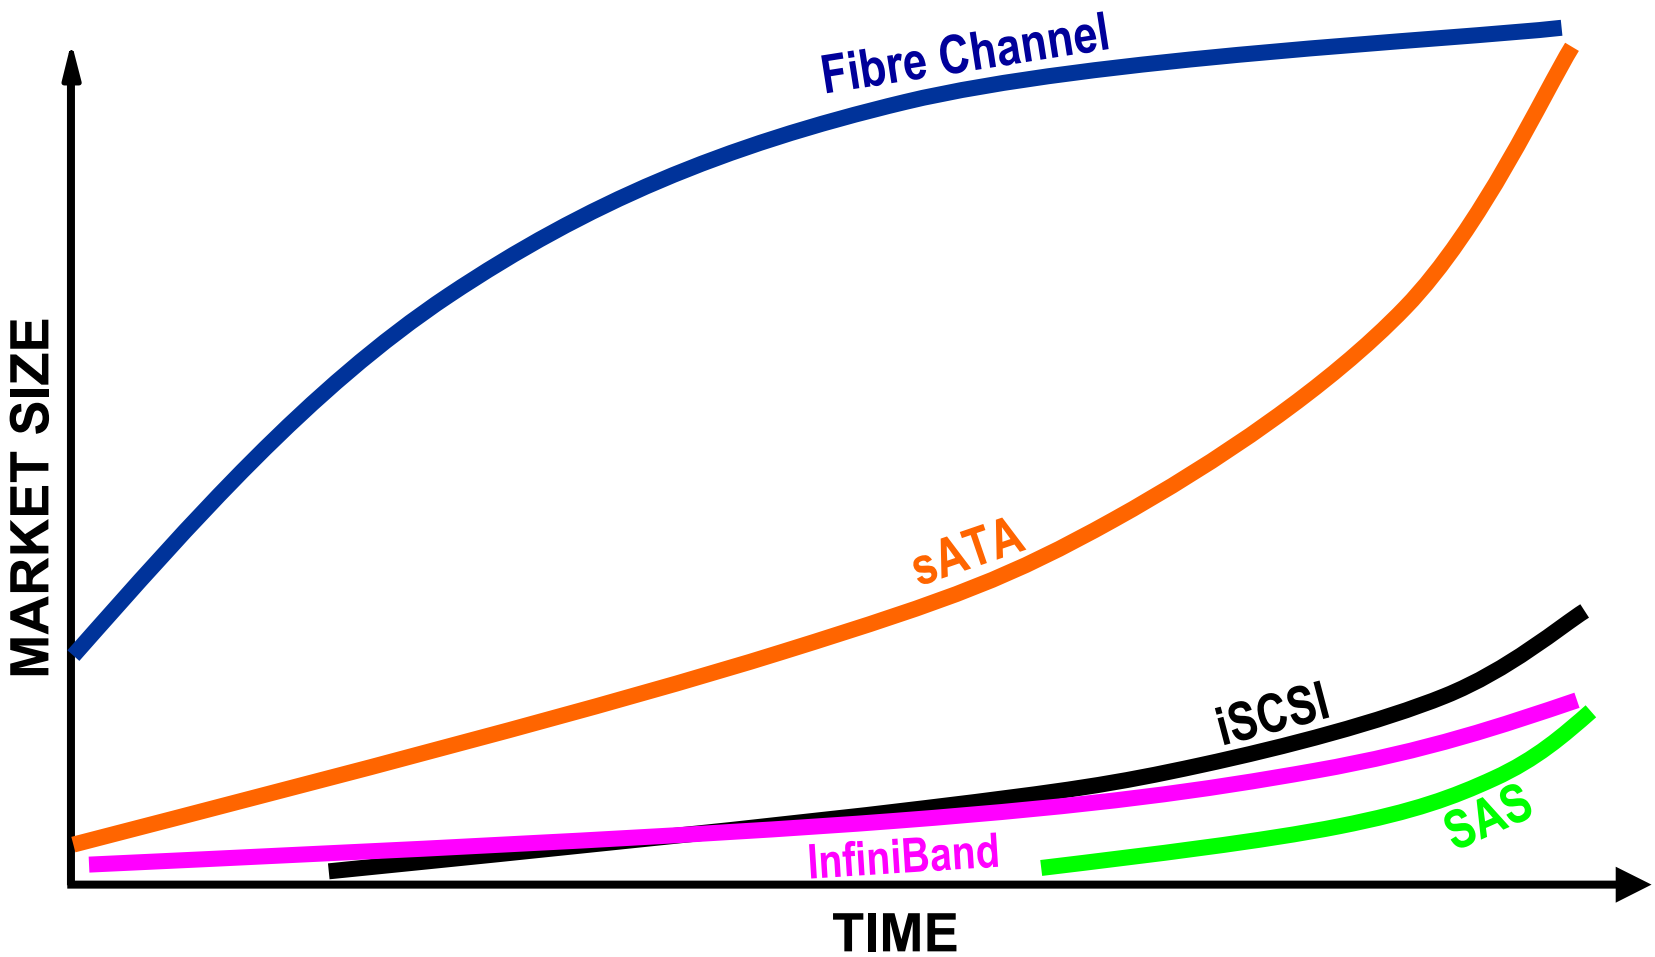

Fibre Channel

- Originally developed in the late ‘80s as a higher-performance LAN alternative to FDDI
  - ANSI T.11 spec.
- Adopted in the early ‘90s as a higher-performance disk drive alternative to pSCSI
  - FCP - SCSI protocol over FC
- Began to gain market acceptance in the mid-‘90s on disk drives and high-performance DAS
- Became a technology used for early storage networking in the late ‘90s
- Dominates SAN technology in early ‘00s

Storage Area Network (SAN)

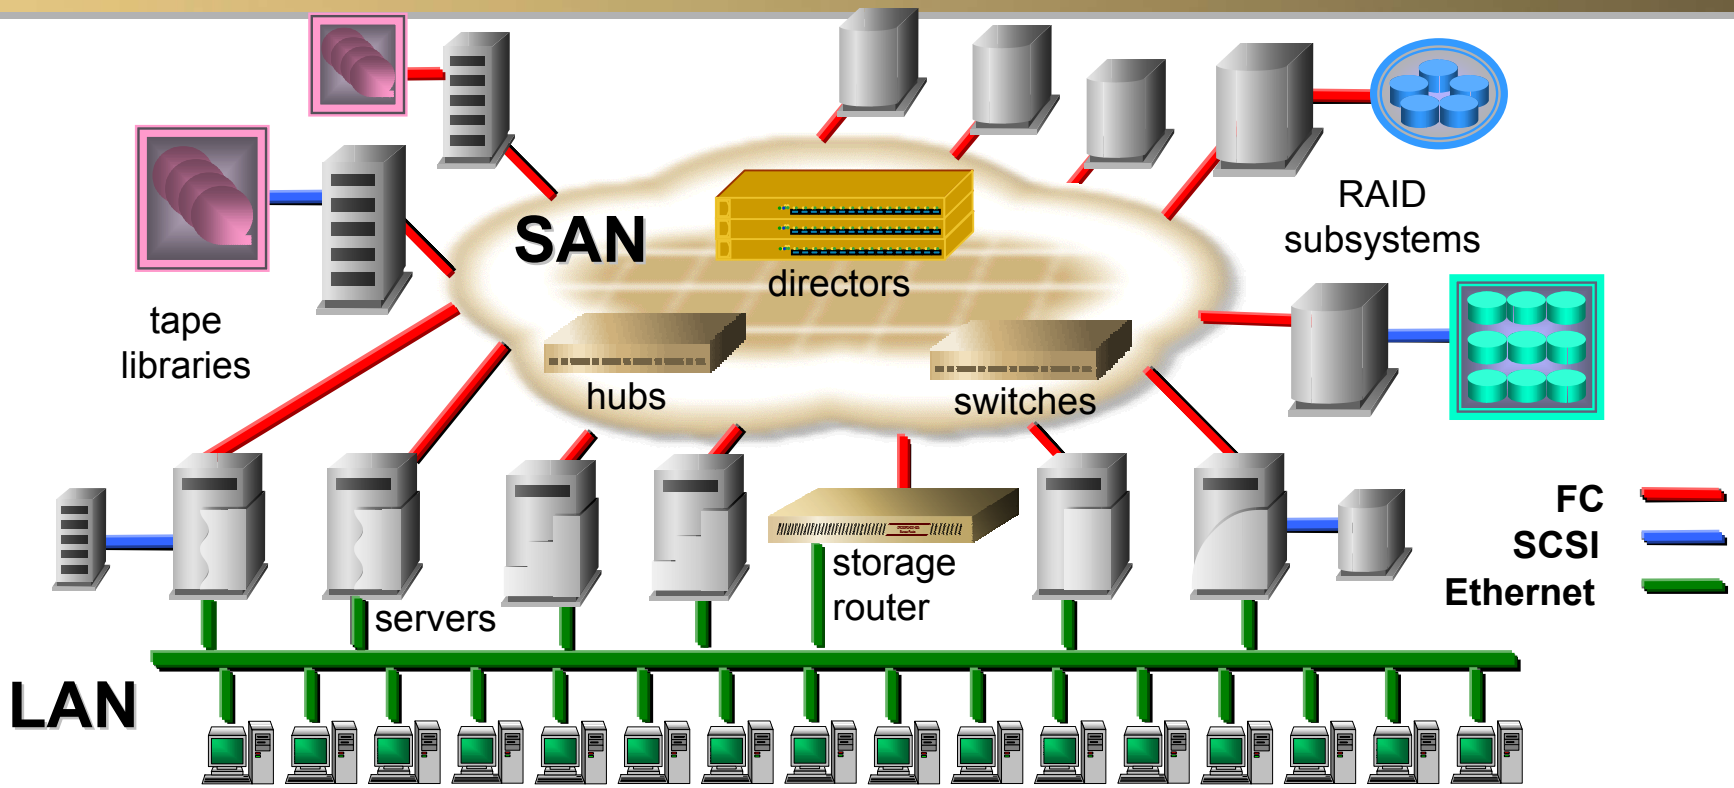

- Fibre Channel is the base SAN technology
- Almost all network storage infrastructure technology today is Fibre Channel
- SAN connectivity required for enterprise servers

InfiniBand™

- Developed as a single unifying industry-wide server I/O interconnect
  - IBTA spec.
  - SRP - SCSI protocol over IB
- Designed as simultaneous network, storage, and IPC interconnect
- InfiniBand has the potential to replace FC as the box-to-box server storage interconnect
- InfiniBand faces some challenges
  - Trouble keeping industry traction
  - Tough economic times for new technologies
  - Requires both new hardware and software

InfiniBand™ Storage Fabric

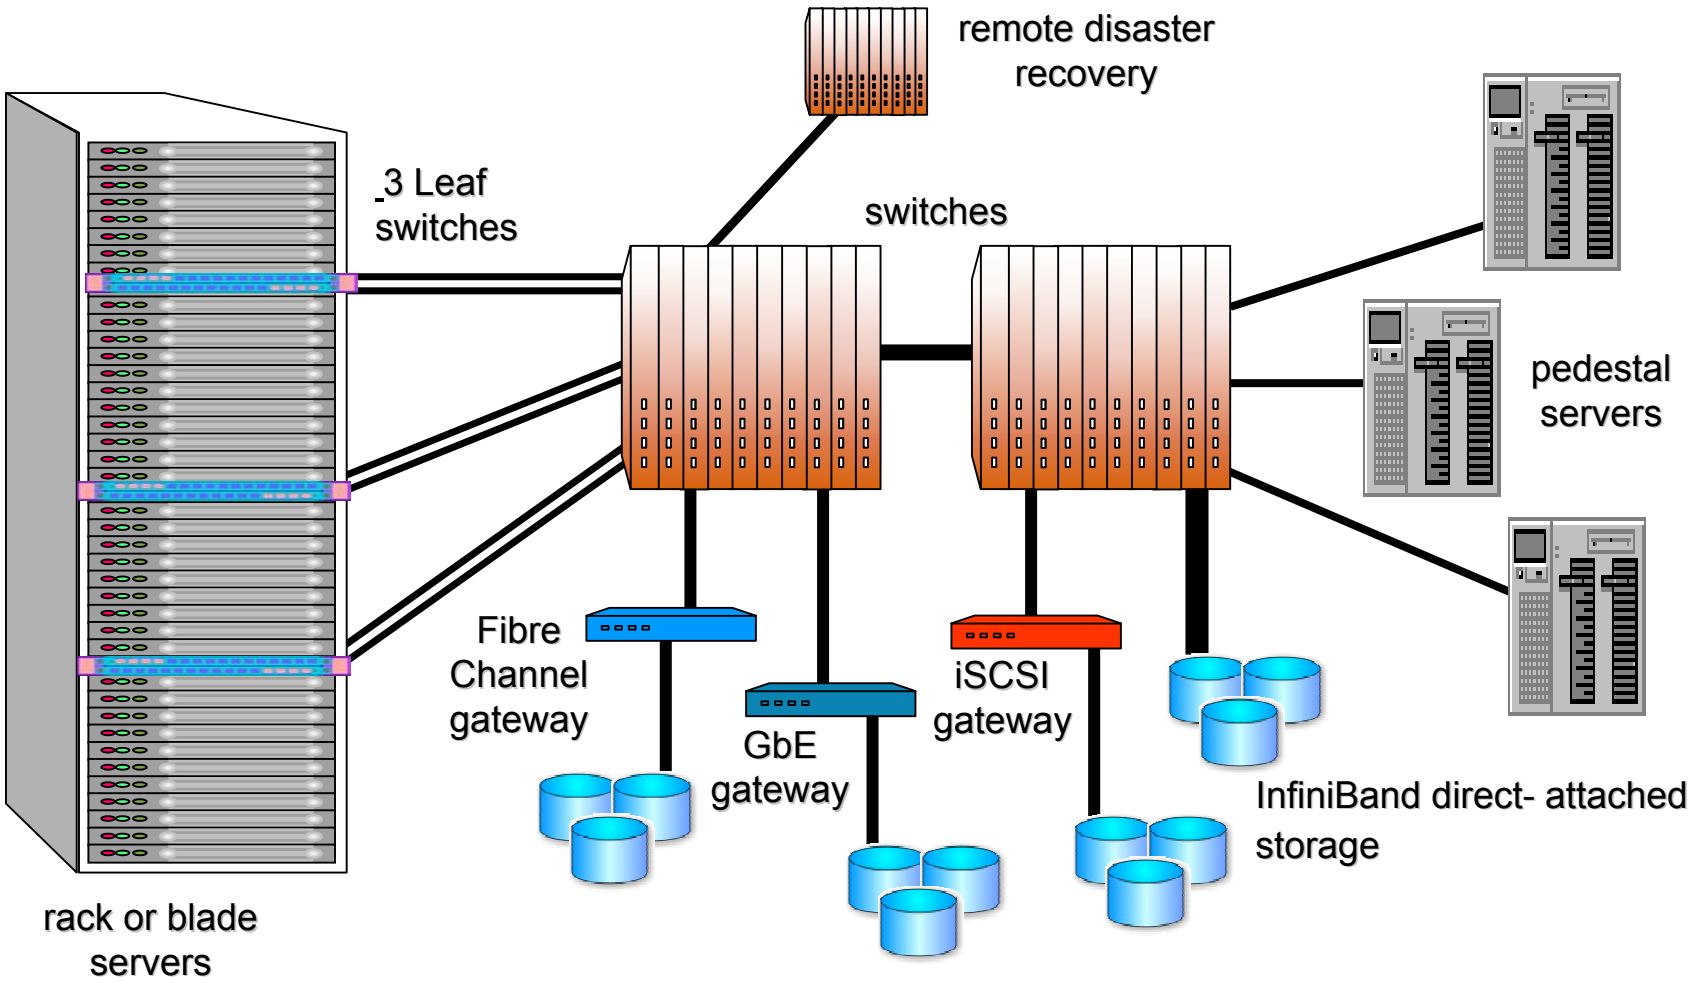

Serial ATA

- Serial version of IDE or parallel ATA disk drive interface used in PCs
  - ANSI T.10 spec.
  - ATA protocol
- Widely accepted that sATA will replace pATA/IDE in desktops and low-end servers in the next few years
  - Pretty good bet this will also happen in bladed servers
- sATA is also showing up in low-cost storage subsystems as drive interface
- Point-to-point technology

iSCSI

- Storage over Ethernet
  - IETF spec.
  - SCSI protocol over Ethernet
- iSCSI has potential to replace FC as box-to-box server storage interconnect
- Under R & D at most large server and storage subsystem vendors - low-end storage interconnect
- Unlike IB, iSCSI is able to leverage an existing infrastructure of hardware and software
- Still has many challenges ahead
  - Customer awareness, cost/performance, acceptance, reliability, and interoperability

Potential Future Storage Area Network

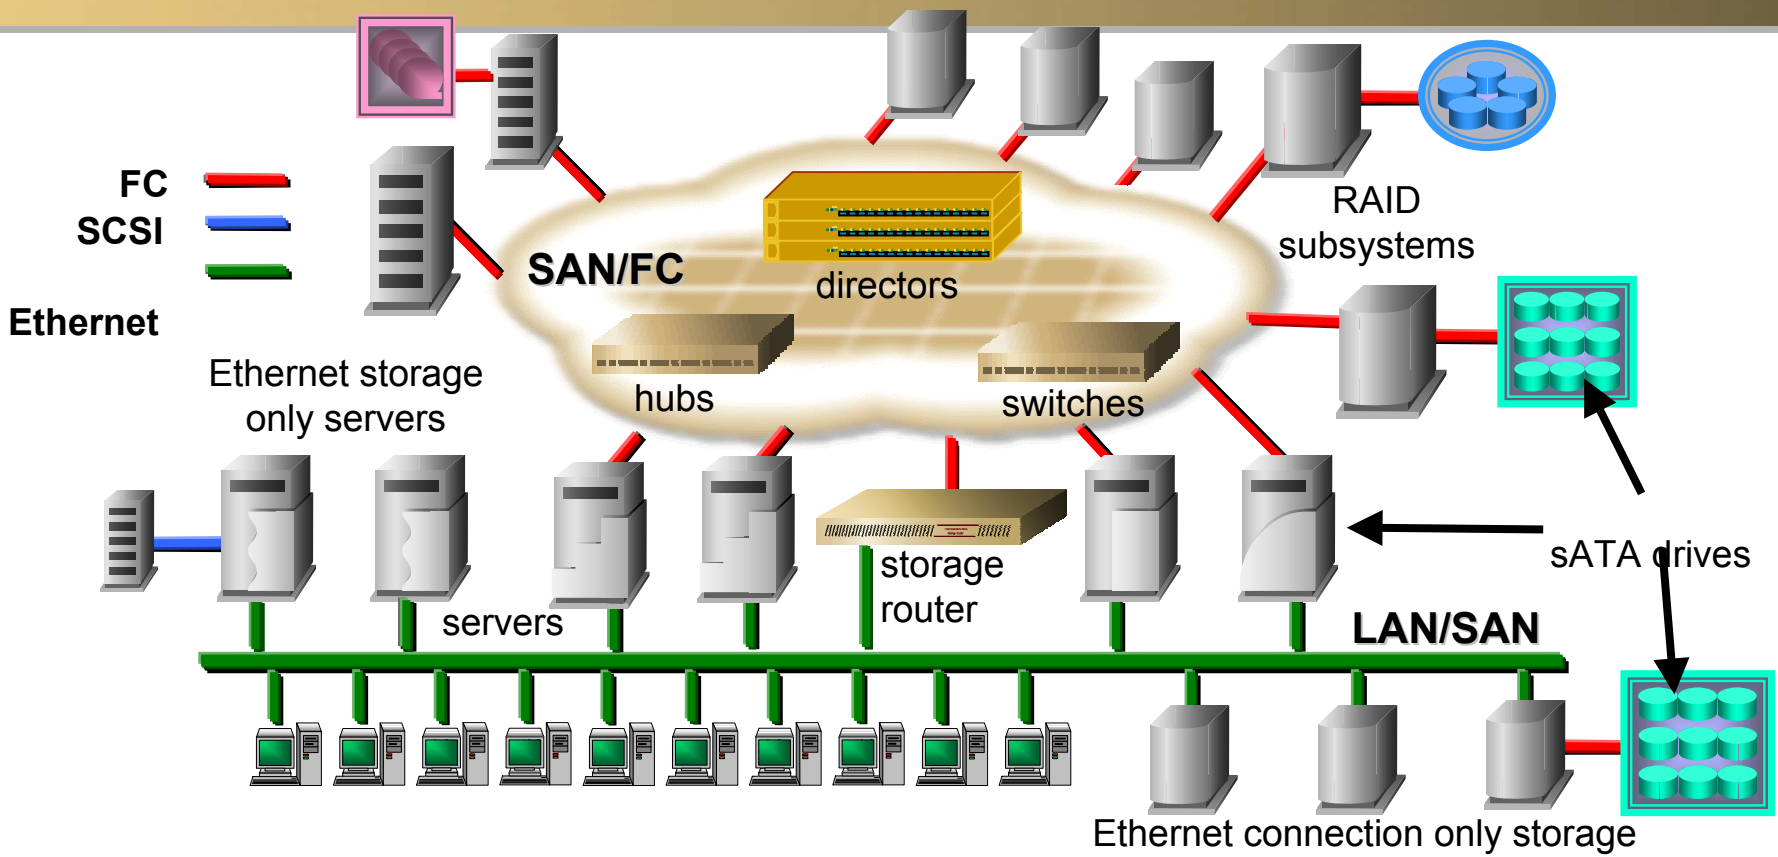

- Some SAN devices are connected via Ethernet only
- sATA drives for lower-performance servers and RAID Subsystems

SAS

- Serial-attached SCSI is the roadmap for pSCSI
  - ANSI T.10 spec.
  - SCSI protocol over ATA Phy.
- Like sATA is to IDE/ATA, this is the serial version of pSCSI
- SAS has the potential to be co-designed onto processor blade servers with sATA
  - Manufacturers could use sATA drives for lower-cost blades and SAS for higher-performance blades
- Point-to-point technology

Architectures

HP WORLD 2003
Solutions and Technology Conference & Expo

Ethernet Bladed Servers – NAS/DAS

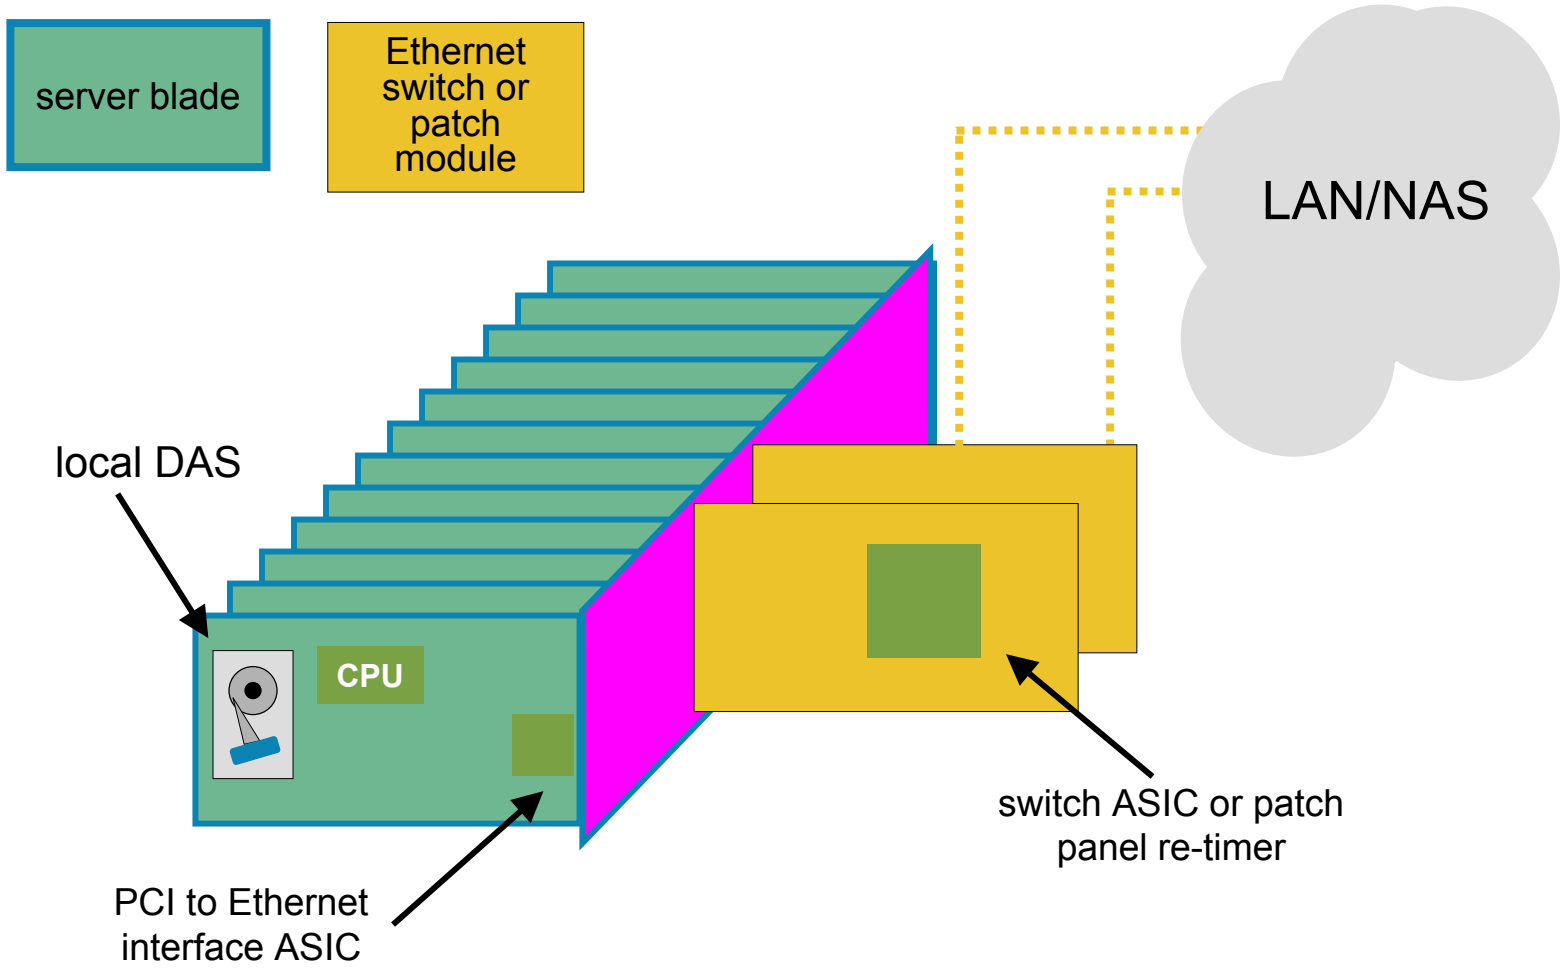

System Architecture

Chassis view

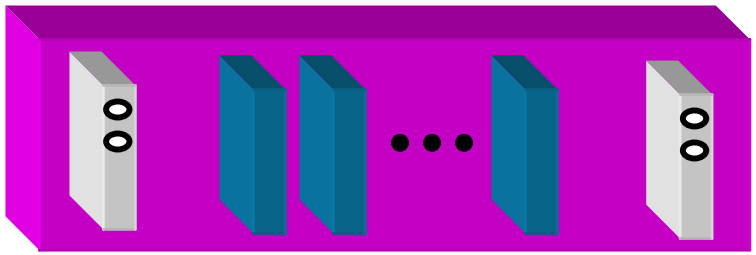

Ethernet topology view

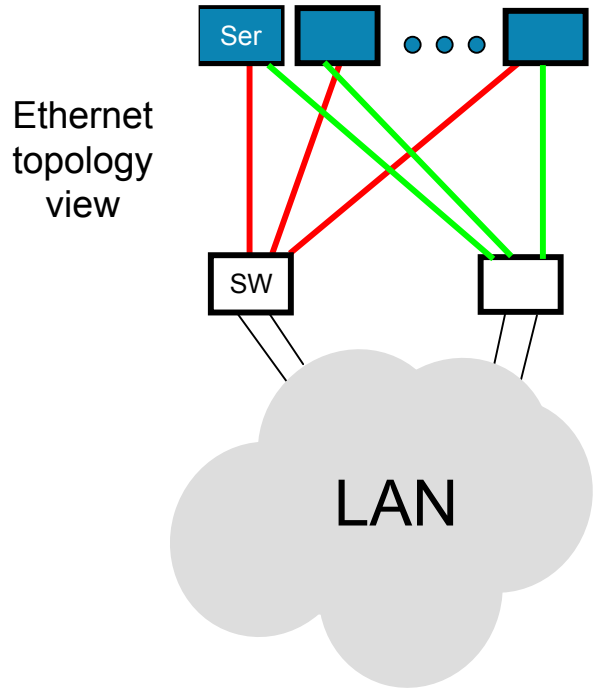

External Data Paths
Primary Data Paths
Secondary Data Paths

SW Switch Blades

Ser Server CPU Blades

Ethernet Technology Bladed Servers

- Advantages
  - Cable consolidation with internal switch only, not patch panel
  - Free storage interconnect for NAS
    - Need LAN anyway
  - Connects directly to existing LAN infrastructure
    - Able to use existing “off the shelf” components
    - No software development required
- Disadvantages
  - No SAN connection
  - Only internal DAS for block storage
    - iSCSI SAN is not widely available or full function
  - Low-performance IPC
  - Bandwidth limited to 1 Gb Ethernet
    - 10 GbE not yet cost effective

Dual Technology Bladed Servers

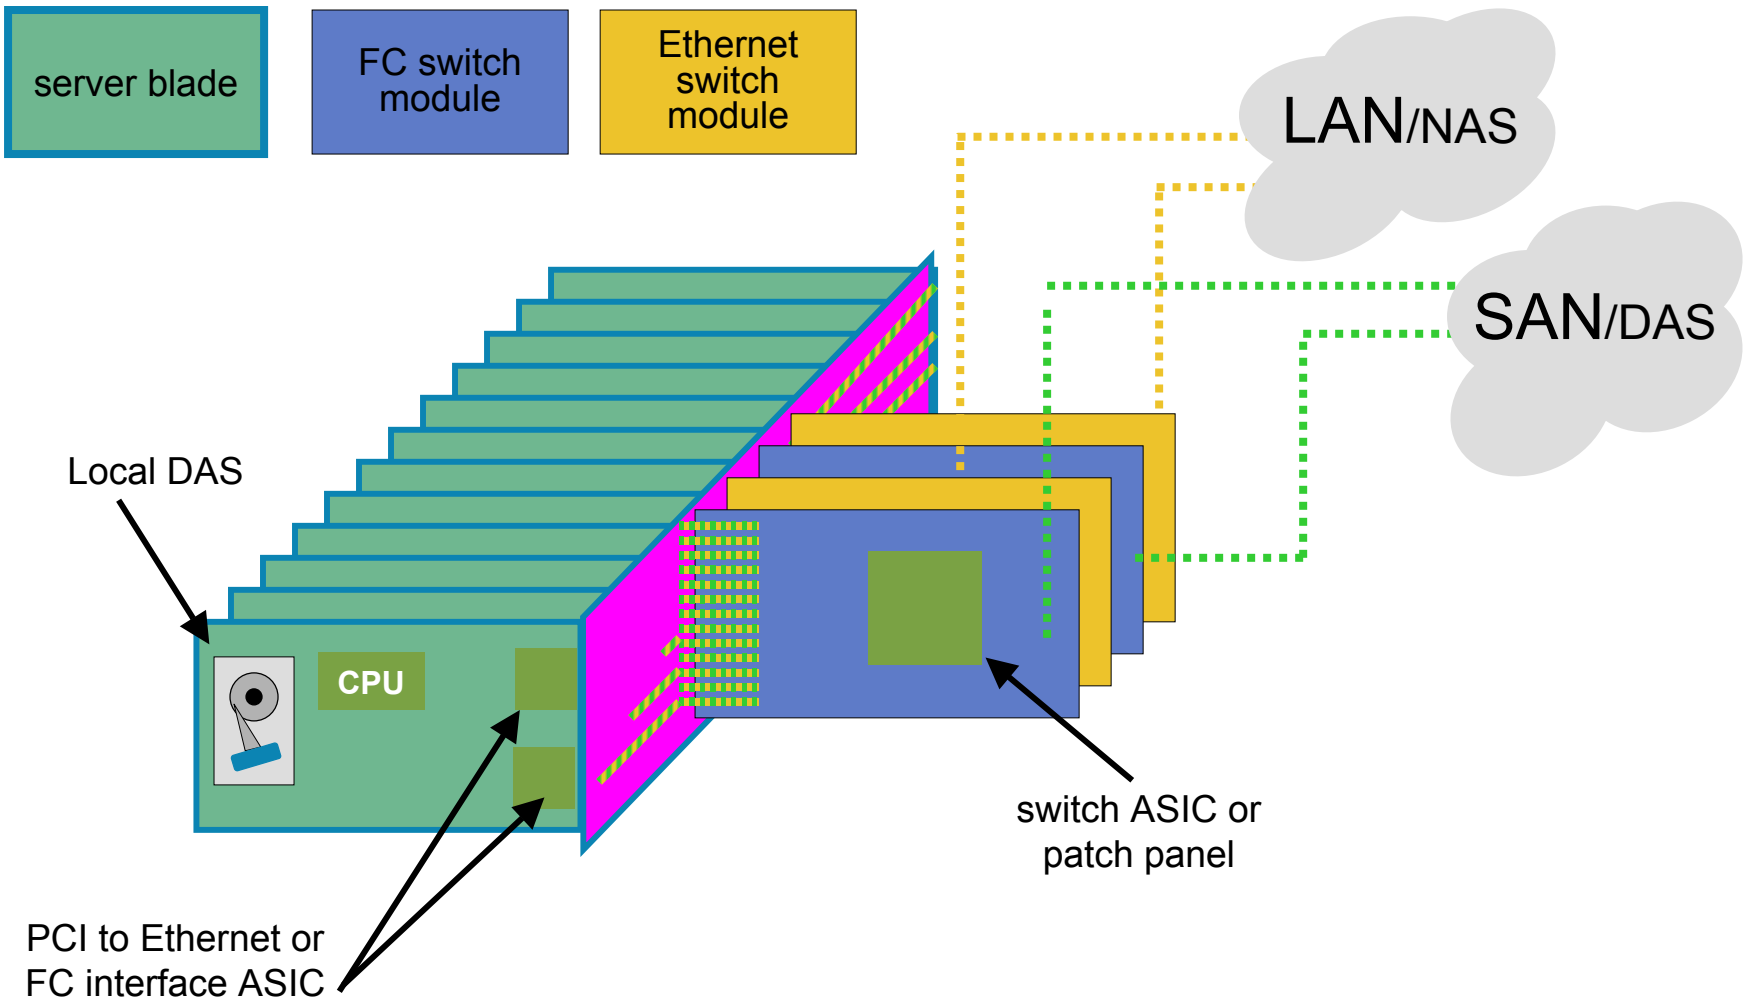

System Architecture

Chassis view

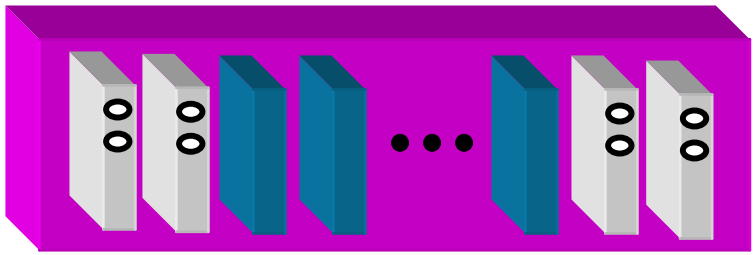

Ethernet topology view

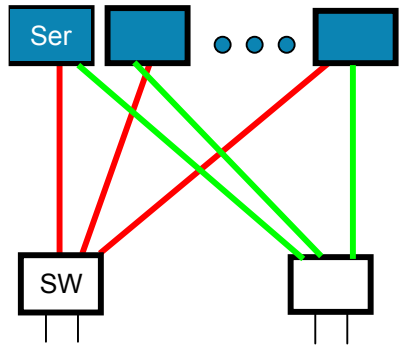

LAN

Fibre Channel topology view

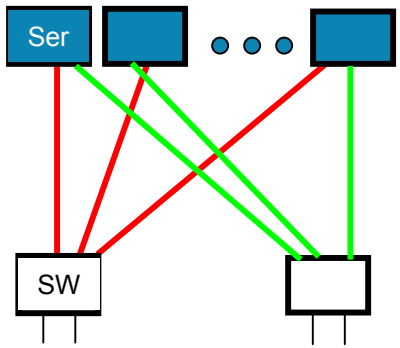

SAN

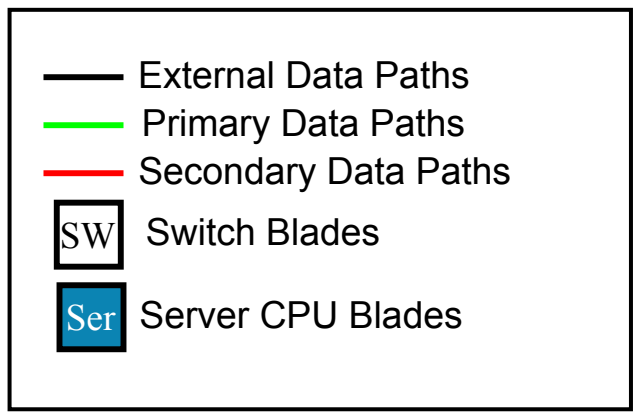

Dual Technology Bladed Servers

- Advantages of this design
  - Provides cable consolidation
    - With internal switch version
  - Connects directly to existing SAN, DAS, LAN, and NAS infrastructure
    - Able to use existing “off the shelf” components
    - No software development required
- Disadvantages
  - Requires dual backplane technologies
  - Requires two switching infrastructures
    - With internal switch version
  - Bandwidth limited to 1 or 2 Gb
    - 10 GbE not cost effective yet
    - 10 Gb FC is not available yet

Single Technology - IB Dense Server

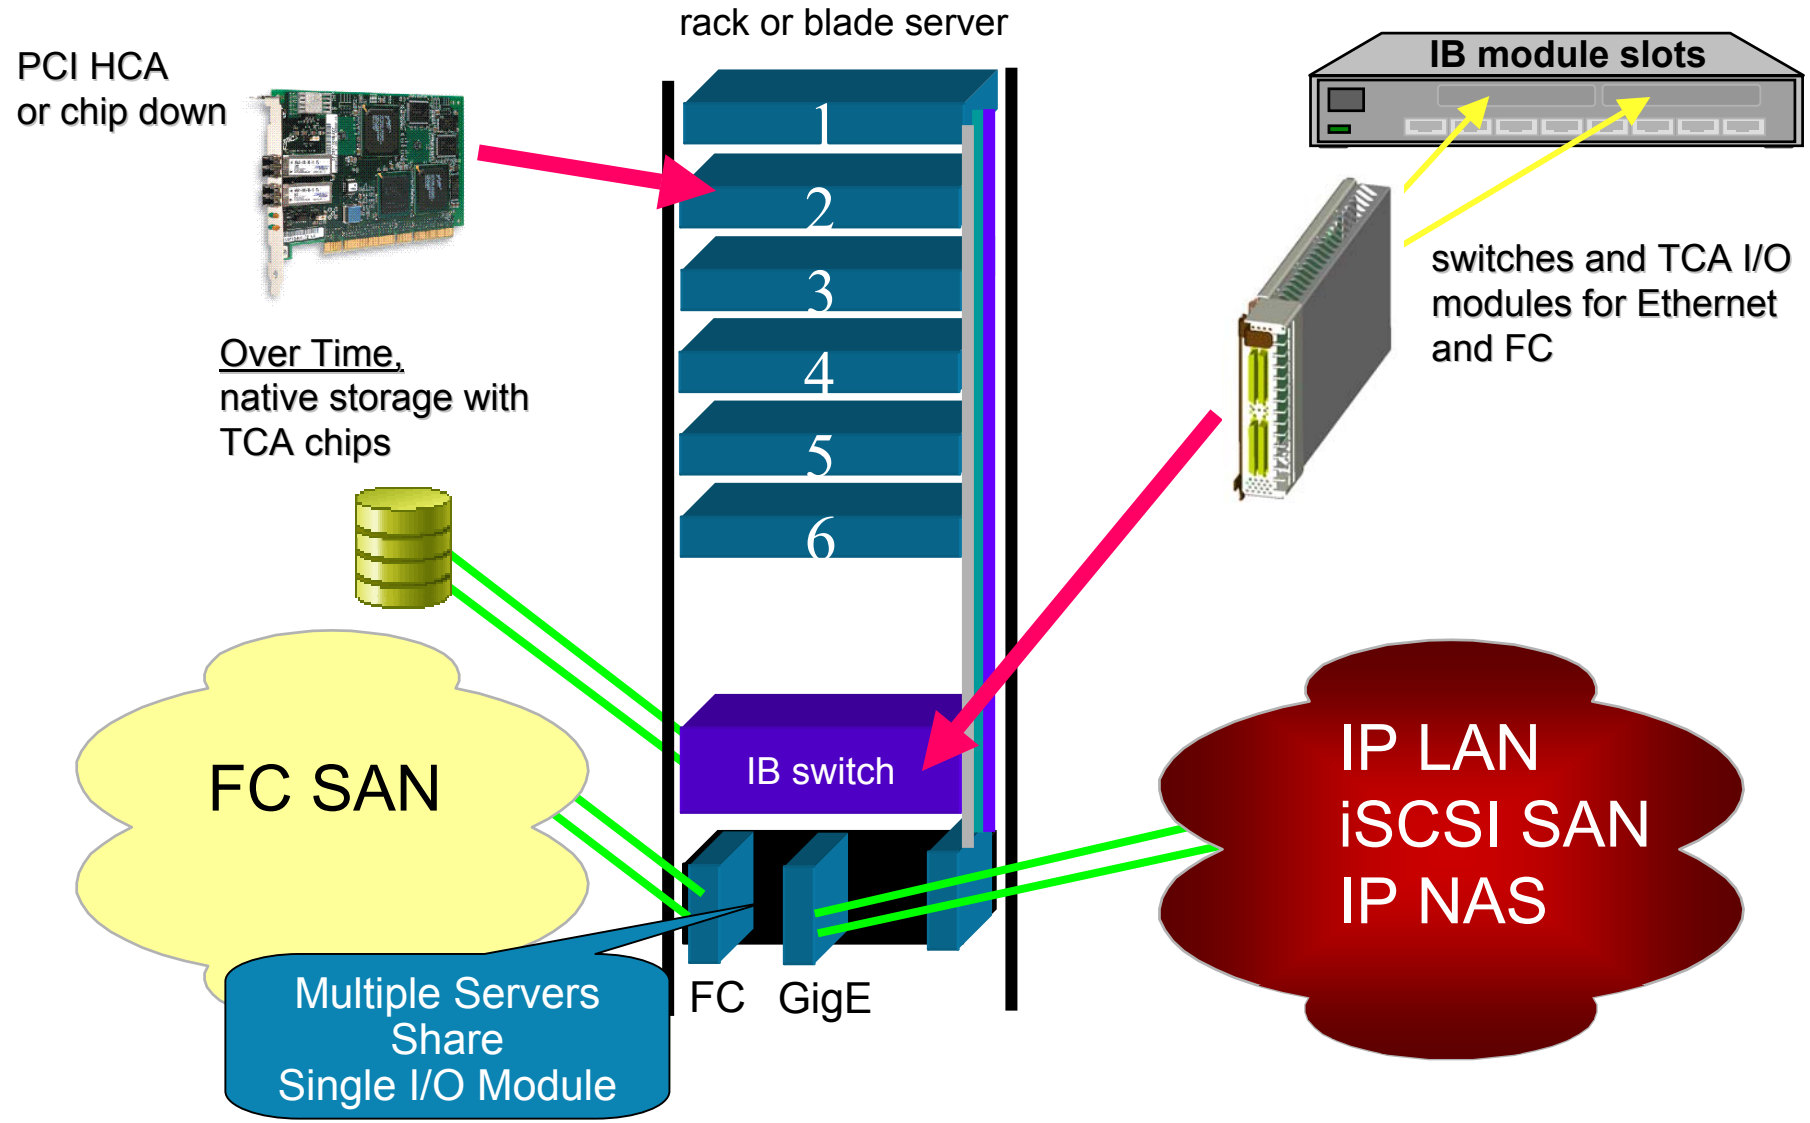

IB to FC TCA Architecture

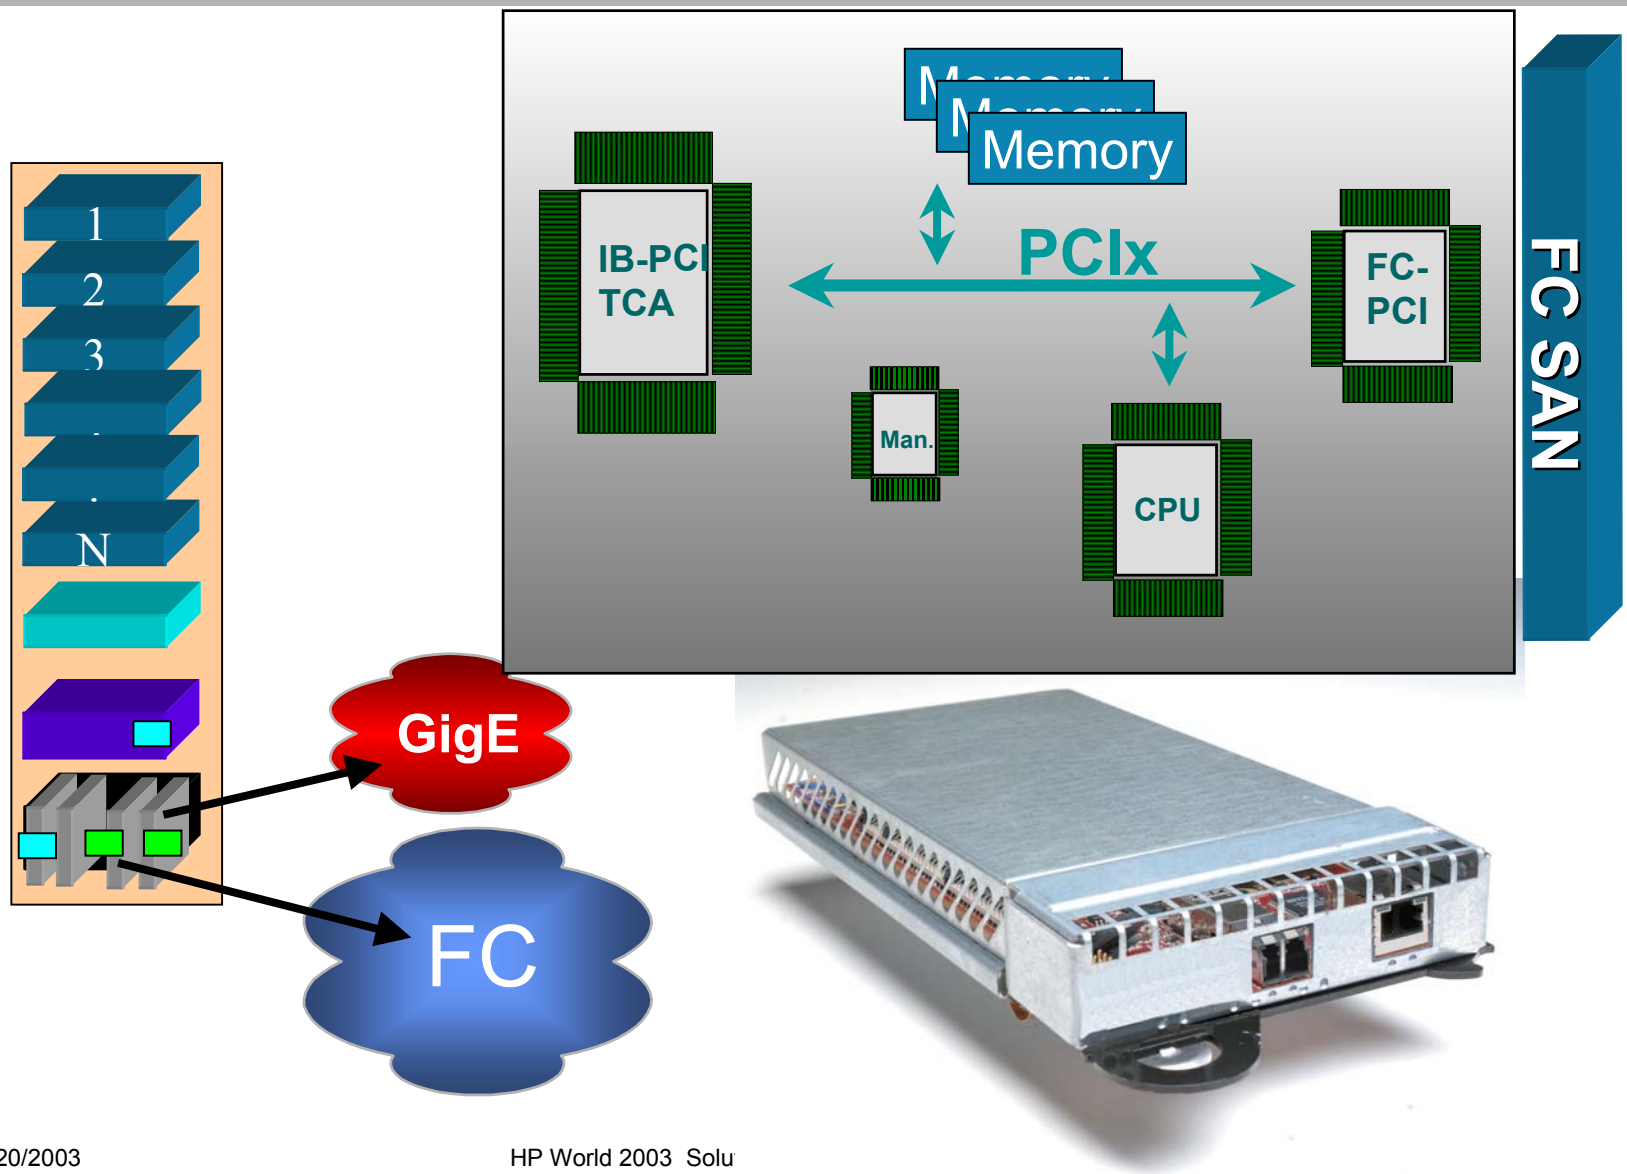

IB to Ethernet TCA Architecture

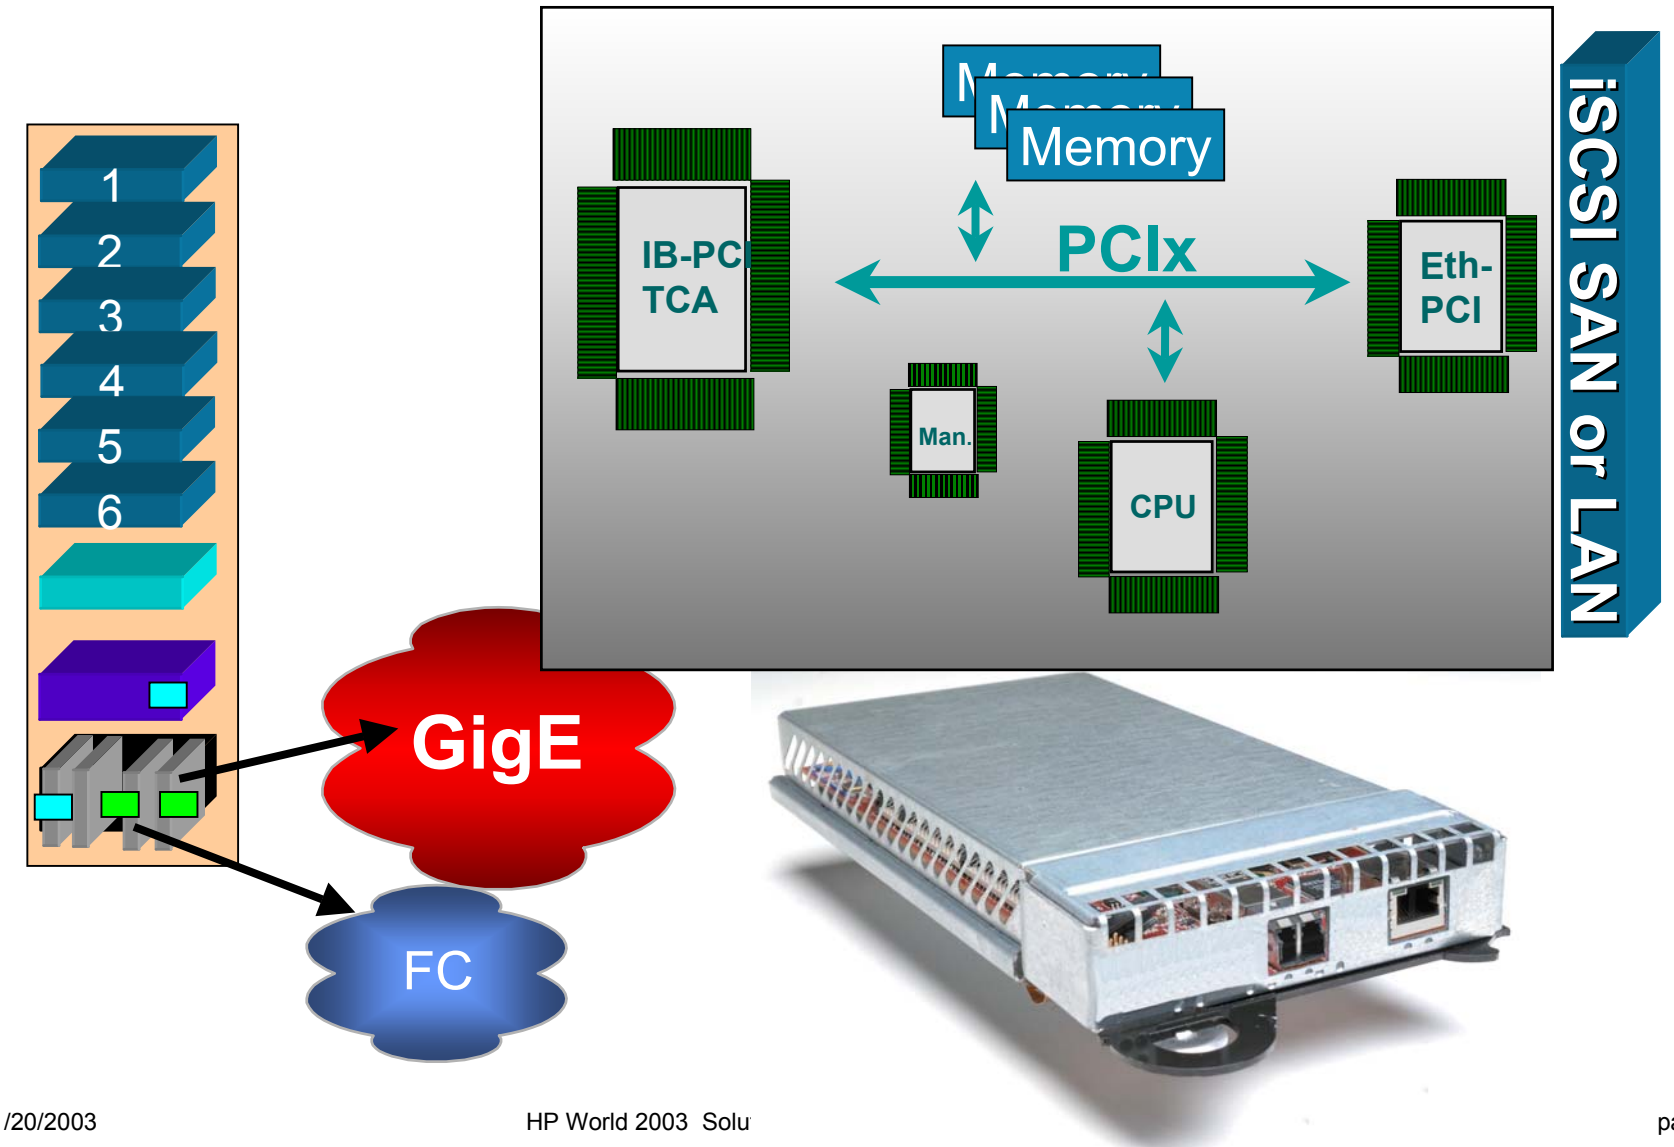

Single Technology Bladed Servers

- Advantages
  - Provides cable consolidation
  - Provides single backplane technology
  - Only need one switching infrastructure
  - Have 10 Gb external connection for aggregation
- Disadvantages
  - Doesn't connect directly to existing SAN, DAS, LAN and NAS infrastructure
    - Direct connect IB storage not available yet
    - This probably causes cost/performance challenges
  - Needs newly-developed hardware components
  - Needs newly-developed software

Next Generation

HP WORLD 2003
Solutions and Technology Conference & Expo

Better Boot from SAN

- Improved ability to boot processor blades from remote storage subsystems across the SAN
  - Current technology has issues due to memory paging
- Advantages of this technology
  - Allows you to potentially remove disk drives from processor blades
  - Allows blades to use many different operating systems
  - Greatly simplifies OS version management
- Solution
  - Technology improvement available soon

10 Gb / Aggregation of Bandwidth

- Ability to aggregate processor blades with lower bandwidth to higher bandwidth external chassis connections (10Gb)
- Advantages of this technology
  - Some applications require top performance
    - Streaming media
- Solution
  - 10 Gb FC and Ethernet are becoming available
  - Allows 1 or 2 Gb connection to processor blades and 10 Gb aggregation connection to SAN/LAN

10 Gb System Architecture

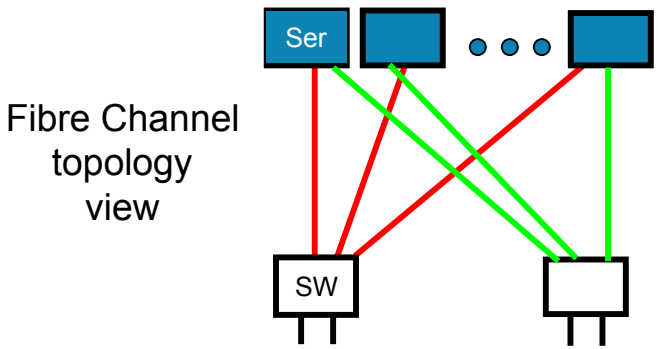

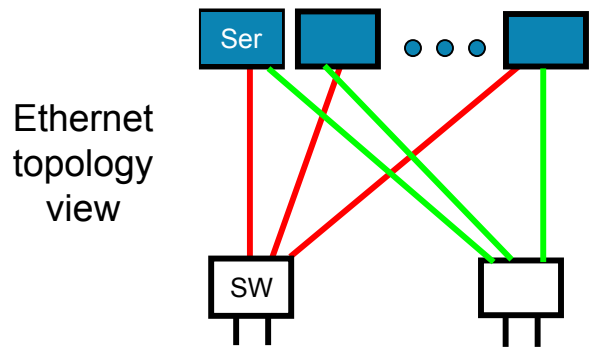

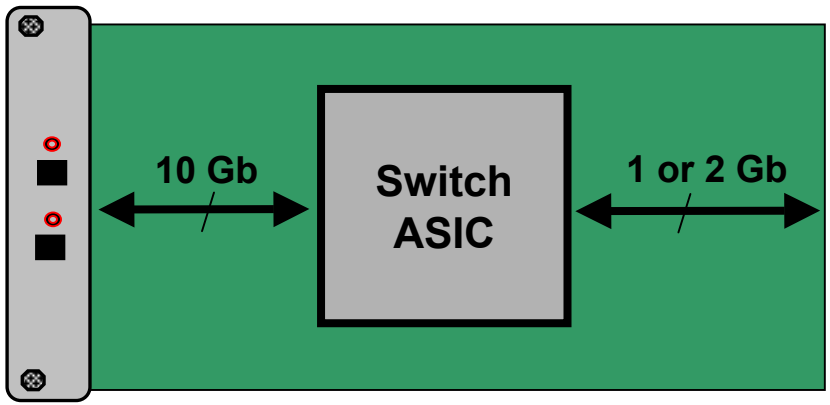

switch blade architectural view

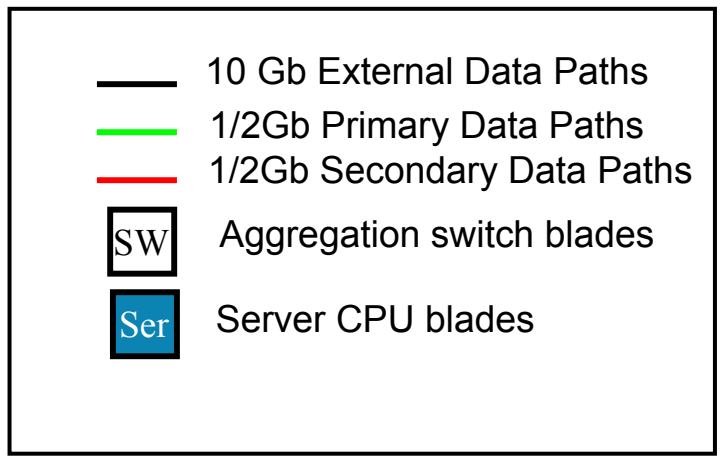

Single Backplane Storage Performance

- Storage processor technology offers good potential to provide single backplane designs a much better cost/performing storage solution

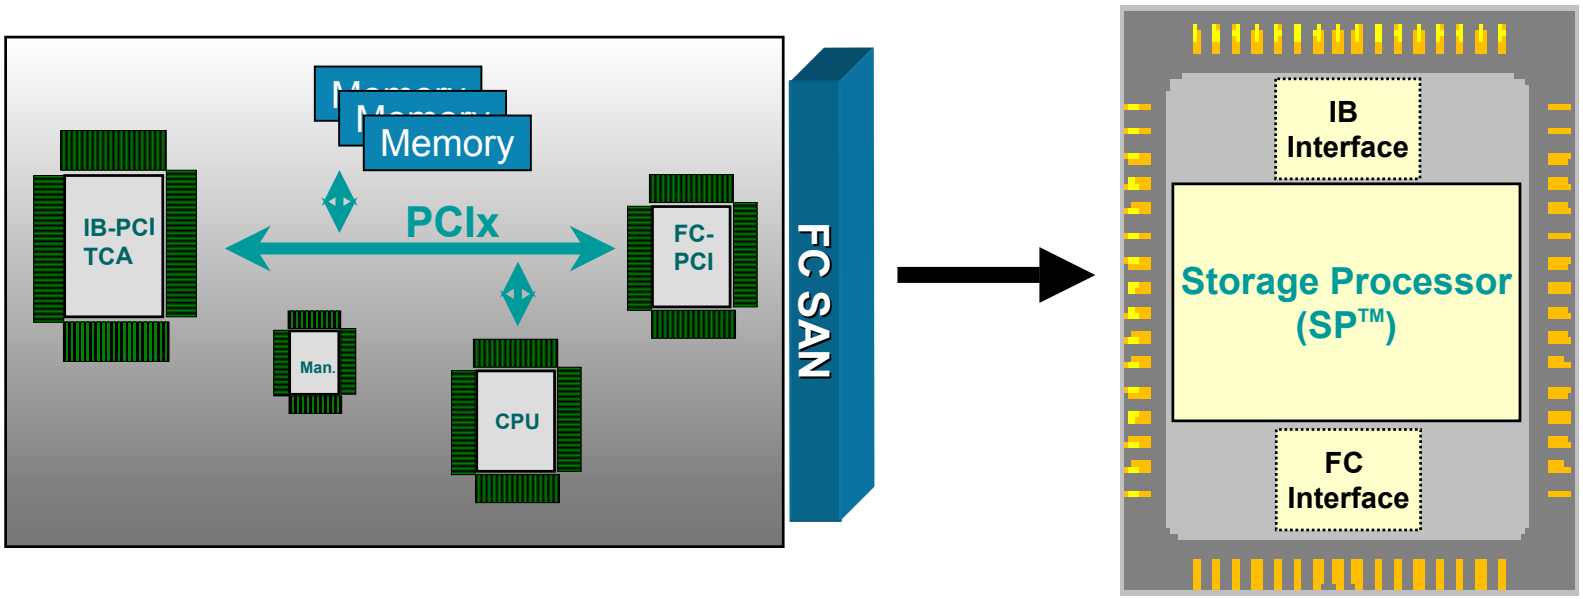

New Storage Fabric Technology

- Both sATA and SAS have future plans to add switching
  - sATA-2
  - Later SAS revs
- These would have potential use as low-end blade server storage fabrics

Summary

- Storage technologies
  - Many choices: FC, sATA, iSCSI, IB, SAS
  - Best: Fibre Channel for SAN/DAS and Ethernet for NAS
- Blade server architecture for storage
  - Single or dual technology backplane
  - Best: dual with Fibre Channel and Ethernet
- What's next to help fill the gaps
  - OS boot
  - 10 Gb

Interex, Encompass and HP bring you a powerful new HP World.