Introduction to Networked Storage

- Howard Goldstein
- President Howard Goldstein Associates, Inc.
- WWW.HGAI.Com
- Howard.Goldstein@HGAI.Com

Agenda

- The rising importance of Data
- History of Storage
- Storage Components
- Understanding the I/O Path
- SAN and NAS in Networked Storage
- Other Alternatives
- Future of Networked Storage

Data as a Business Asset

- Information Age
  - The amount of stored data is increasing at a compounding rate
- Data Warehousing
  - Information analysis is increasingly viewed as a management tool and a competitive weapon
- E-business imperative
  - Businesses realize that the electronic marketplace is unforgiving, data plays a central role in providing customer service
- Re-evaluating the value of data
  - Business people empirically understand the value of data but IT people struggle to assign a dollar value

History of Storage – Part 1

- Host to channel attached storage
  - Host system ‘owns’ the disks
  - Storage access is physically based
  - All storage management is via the host
- Mainframes move to channel attached ‘pools’ of storage
  - Disks are grouped into pools of storage
  - Storage is allocated logically
  - Access moves to logical
  - Storage management moves to separate applications
- Open Systems begin to take hold

History of Storage – Part 2

- Introduction of RAID
  - Data Protection through mirroring and parity checking
    - Mirroring – copies data during writes
    - Parity checking – XOR calculations
  - Increased physical speeds through disk striping
- Common RAID Levels
  - 0 Striping without parity (no redundancy)
  - 1 Mirroring
  - 3 Synchronous striping with dedicated parity disk
  - 5 Independent striping with distributed parity disk
  - Combinations – 0+1 or 10
- RAID Advisory Board (www.raid-advisory.com)

History of Storage – Part 3

- RAID disks are still mostly direct attached
- Multi-port storage systems
  - Allow sharing of resources
  - Support storage based services and management within the storage system
- Networked storage protocols
  - NFS (Network File System) - introduced with Unix operating systems
  - CIFS – (Common Internet File System) -
- Storage networks emerge
  - SUN – includes support for fibre channel in O/S
  - IBM – Serial Storage Architecture

Storage Components

- Storage Services
  - Backup and Restore
  - Disaster Recovery
  - Fault Tolerance
- Storage Management
  - Space Allocation
  - Quota Management
  - Security
- Storage Hardware
  - Disk, tape, etc.
  - Host Bus Adapters (HBA)
  - Interconnect media

Tracing the I/O Path (simplified)

- Server-side processing
  - Application code generates request
  - CPU to I/O controller
  - I/O controller to internal bus
  - Internal bus to channel adapter
- Storage-side processing
  - Channel adapter to disk controller(s)
  - Storage system controller to internal bus
  - Internal bus to disk drive(s)

I/O Path Variables

- Block versus file level requests
- I/O controller functionality
- Channel alternatives
- Disk controller functionality
- Storage system internal bus alternatives
- Drives – speeds and feeds, etc.

The Availability Challenge

- Backups or other data operations frequently require data be taken ‘off-line’ to be fully protected, causing interruptions to the applications
- Increasingly data must be available 24 X 7
- Two alternatives exist to overcome this challenge
  - Remote Mirroring – copies data to other storage systems, includes facilities to ‘break’ the mirror and resynchronize later
  - Snapshots – creates point-in-time copy of data either physically or logically (directory)
- Snapshots or remote mirrors can serve as additional copies of the data for other processing.

Networked Storage – A Story of Two Technologies

- NAS – network attached storage
  - Commonly regarded as a plug-and-play approach to providing additional storage capacity
  - Utilizes industry standard protocols to serve files over existing networks
  - Can provide multi-O/S access to the same files
- SAN – storage area network
  - Regarded by many as the future of storage technology
  - Provides a many-to-many connection between servers and storage components
  - Removes storage traffic from application networks
- NAS and SANs are complementary technologies

NAS – Topological View

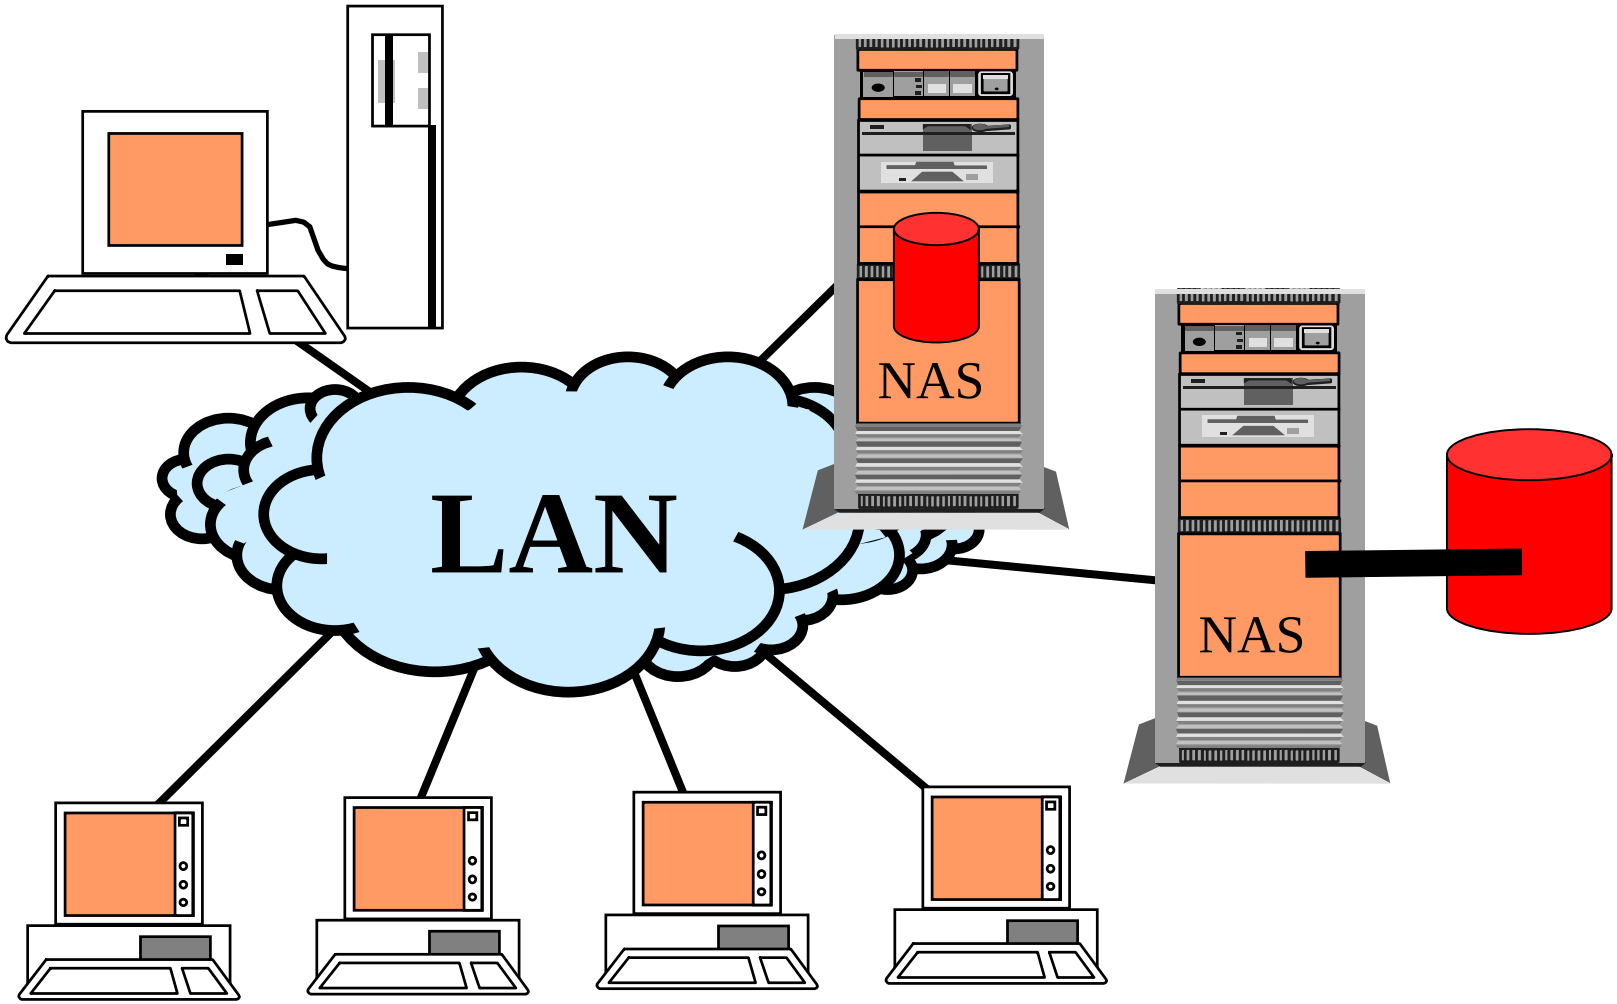

NAS Alternatives

- Appliance devices
  - Plug into the network and auto-configure
  - Generally smaller in size, under 100GB
- Server-based solutions
  - Utilize a Unix or Windows server as host
  - May offer additional O/S based functions and manageability
- Full-featured solutions
  - Include built-in management and storage services
  - Provide high capacities – up to 1Tb or more
  - Features may include advanced caching, fail-over, etc.

NAS Applications

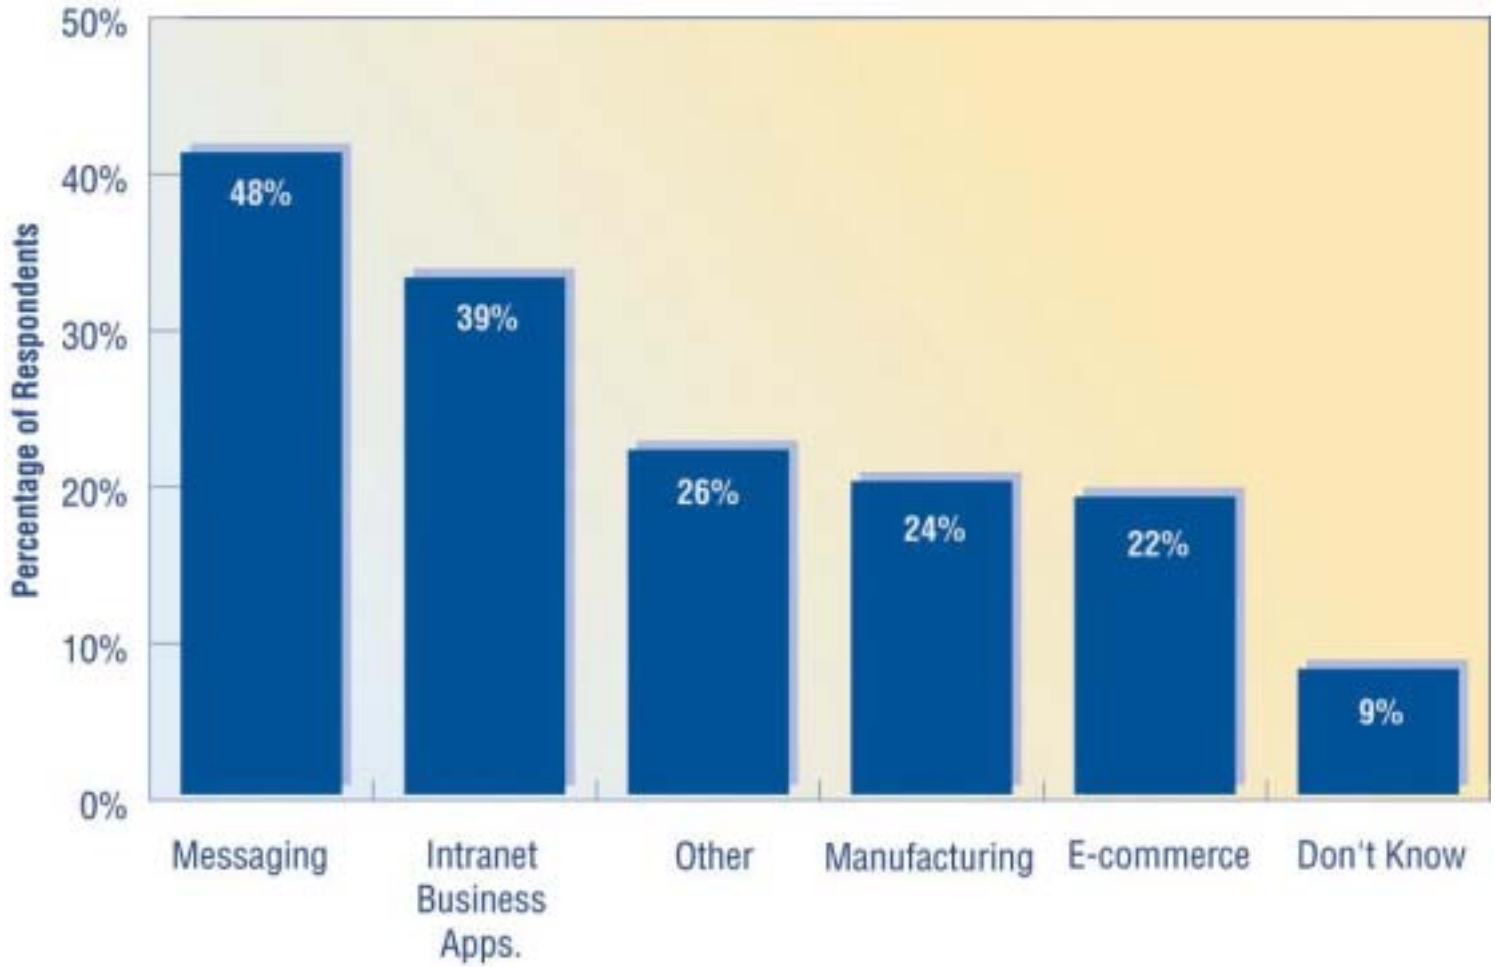

SAN – Topological View

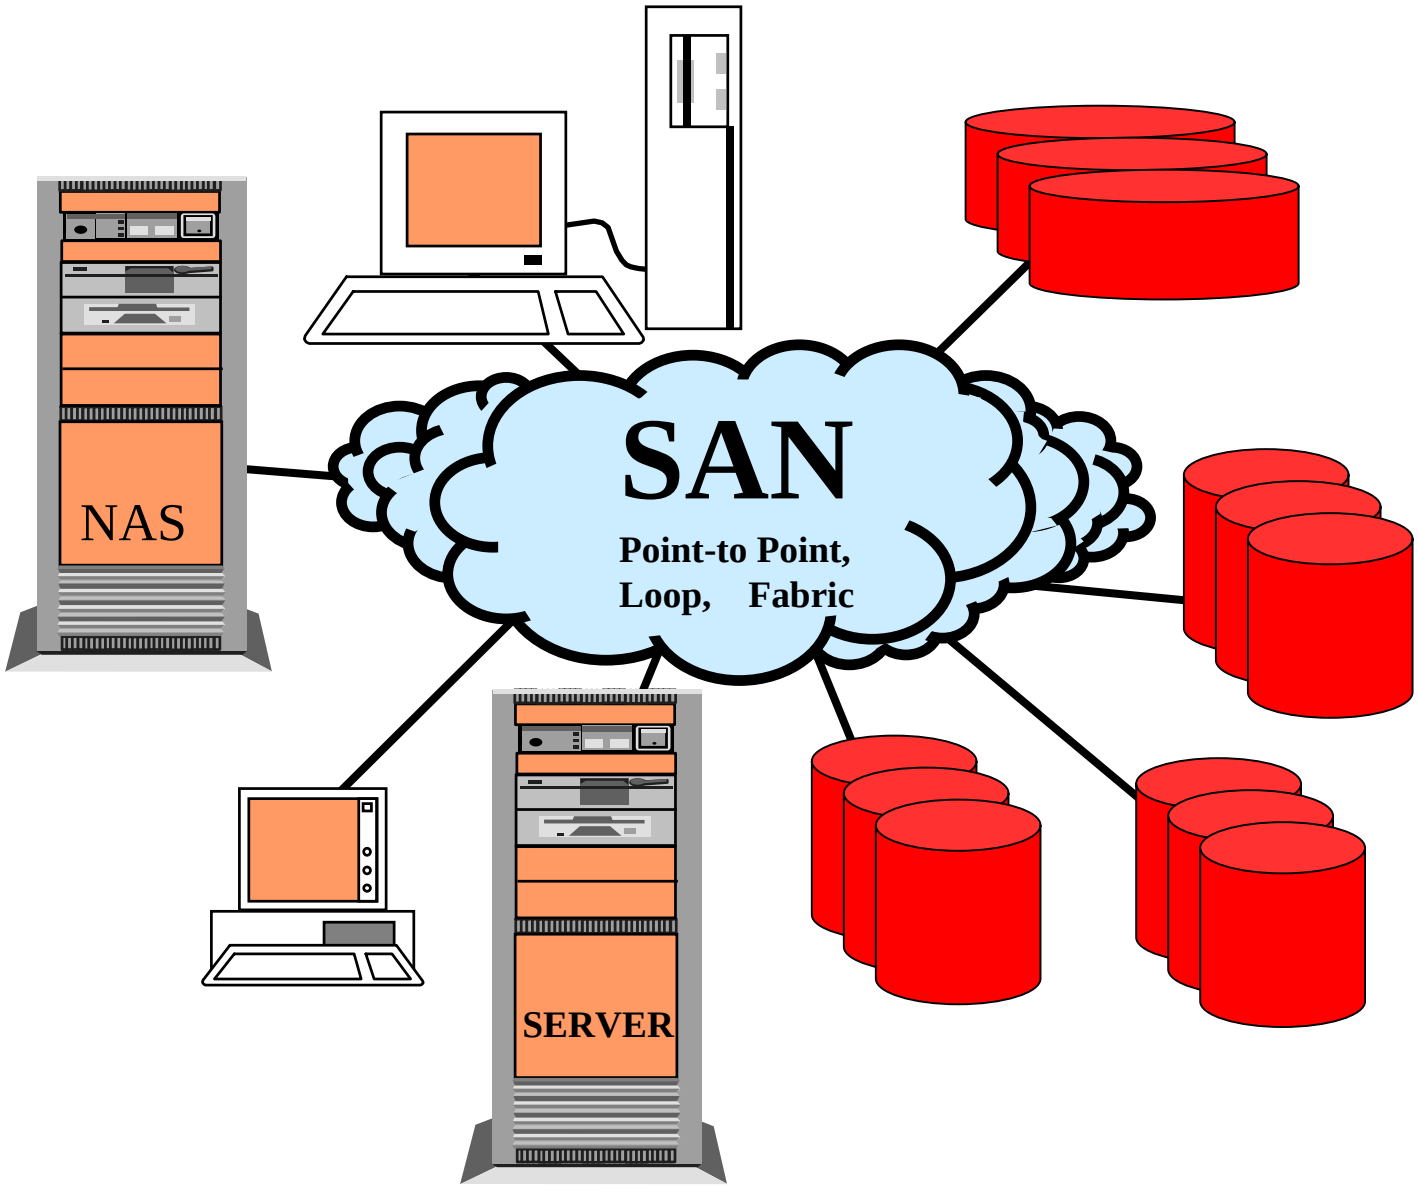

SAN Applications

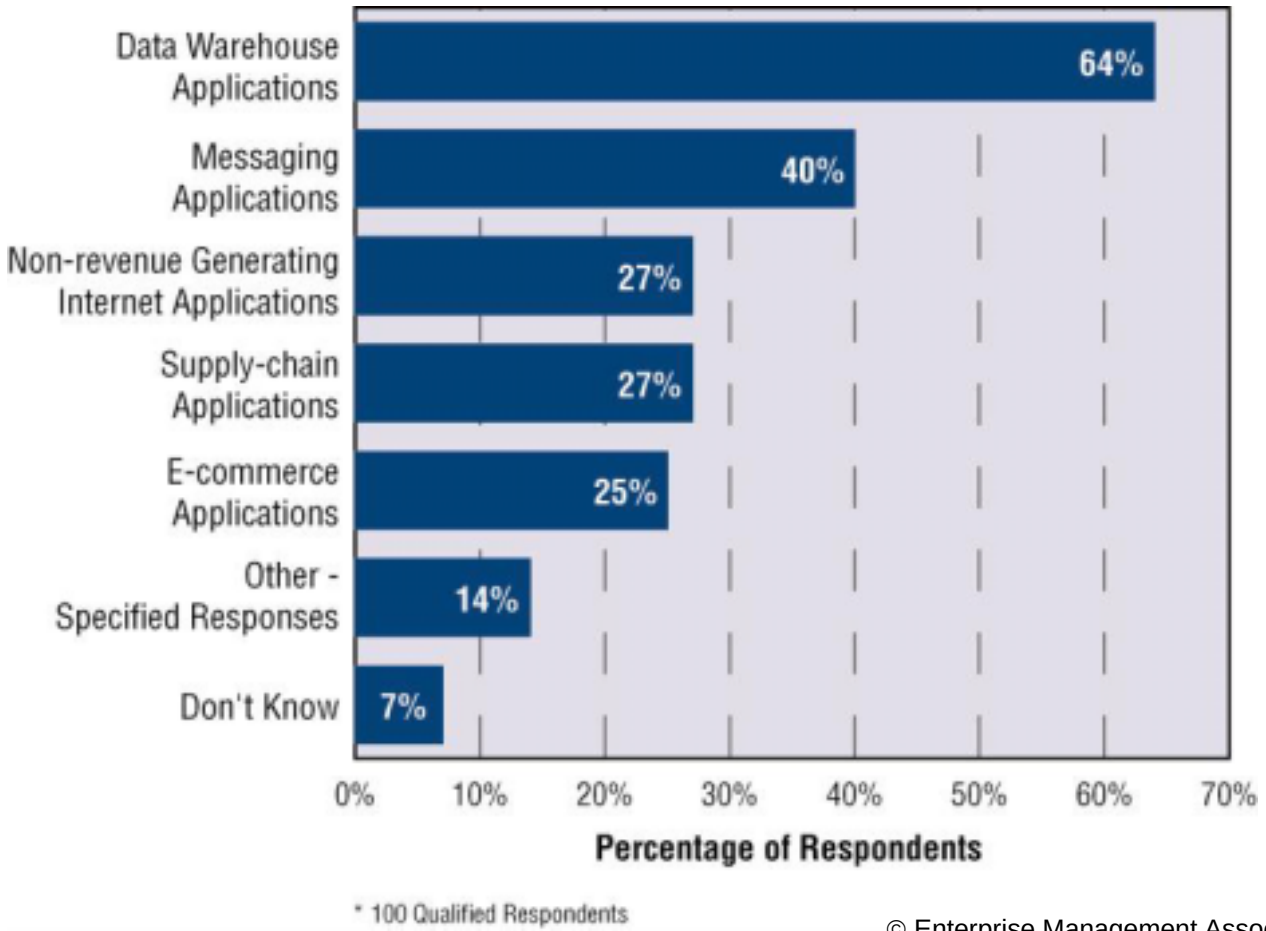

SAN Alternatives

- Fibre Channel –
  - Industry standard protocol
  - Switched – utilizes name services for device login and provides one-to-one connections using non-blocking switches or directors
  - Arbitrated Loop – shared loop configuration using arbitration to service requests
- Gigabit Ethernet
- SCSI -
  - native or SCSI over IP
- Others
  - VIP, Infiniband

SAN Alternatives – Making the Choice

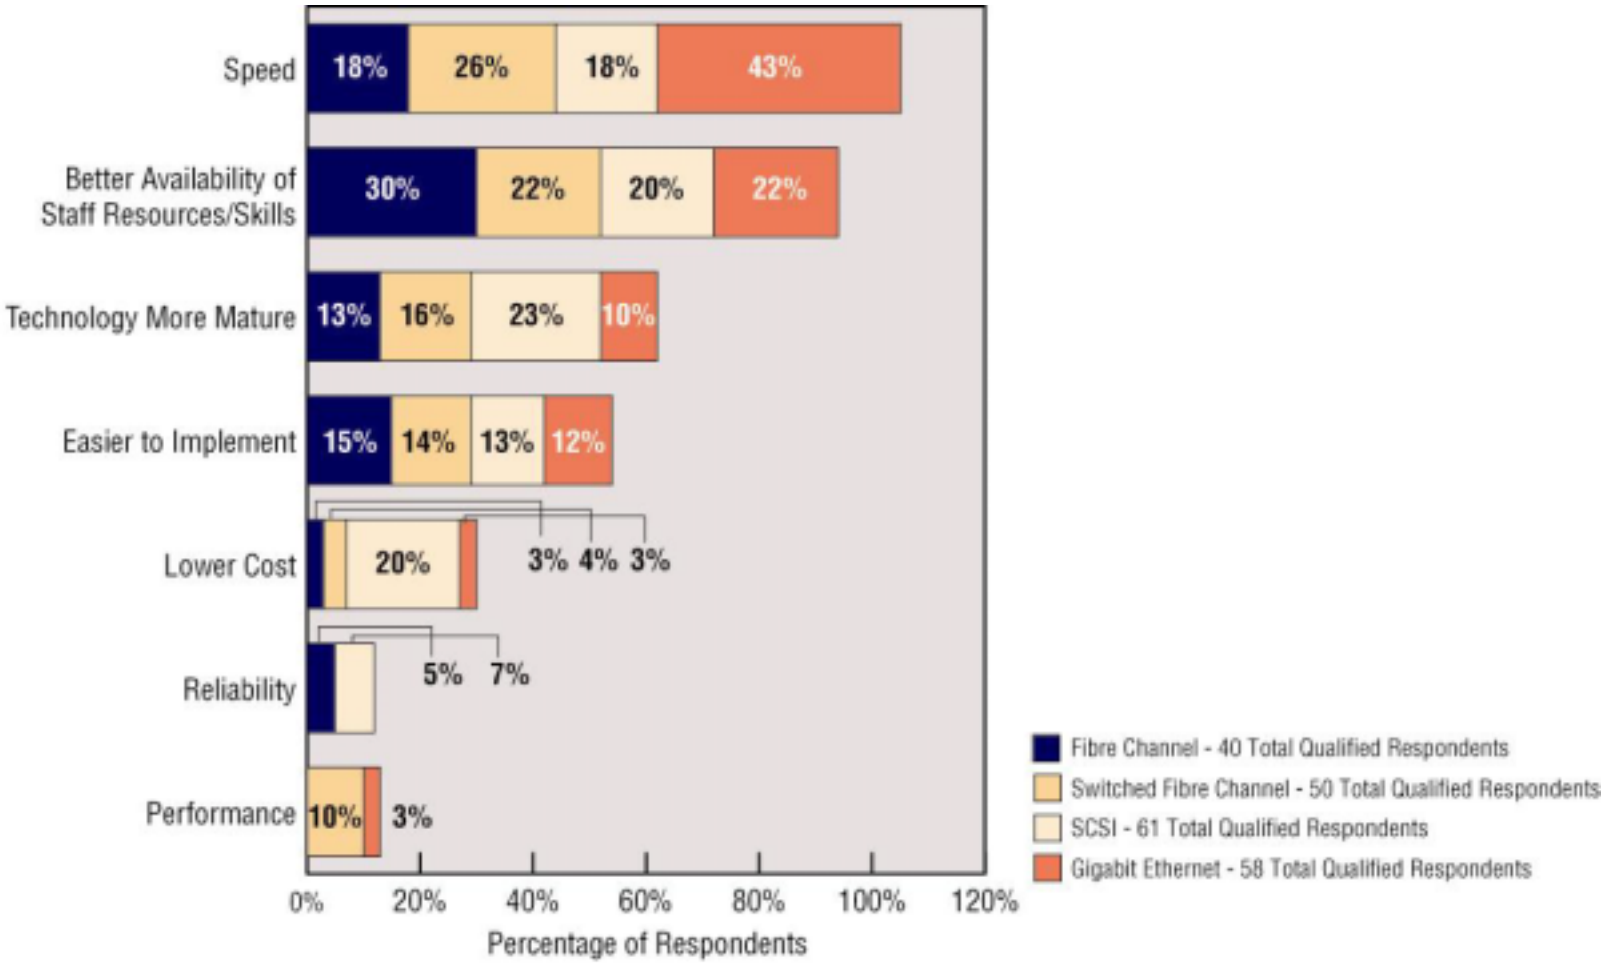

Software Based Storage Networks

- Use software to intercept storage I/O requests from the server
- Reroute requests through storage manager server
- May include server based or storage network based storage systems
- Include caching and other advanced storage processing
- Provides a virtual storage network

Storage Service Providers

- Cost per Megabyte
- Service to Customers
  - Via high speed storage interconnect
- Service via third parties
  - NSP, ISP, ASP, etc.
  - Transparent to customer

Networked Storage Today

- Storage Networks are used in many open systems environments
- Users anticipate their SANs will solve storage problems such as:
  - Higher data availability
  - Easier to share and manage storage resources
  - Faster data access
  - Relieve constraints during backups
  - Scalability

Future of Networked Storage

- Independence from Hosts
- Disk Growth or Changes without application interruption
- Networking Innovations

Future of Networked Storage

- Storage services software
  - LAN-free backup – moving the backup traffic from the application network while maintaining shared resources and central management
  - Server-less backup – removing the backup service impact from the application server
- Storage management software
  - Policy based
  - Capacity on demand
  - Automated storage tasks

Future of Networked Storage

- New Filing Systems
- Cross system record locking
- Intelligent data objects
  - XML, HTML, others

Networked Storage – Physical Transport Choices

- Fibre Channel
- Storage over IP
- Gigabit Ethernet
- SCSI Bus
- ESCON
- Infiniband
- iSCSI
- Topology Options

OSI & IPS (TCP/IP)

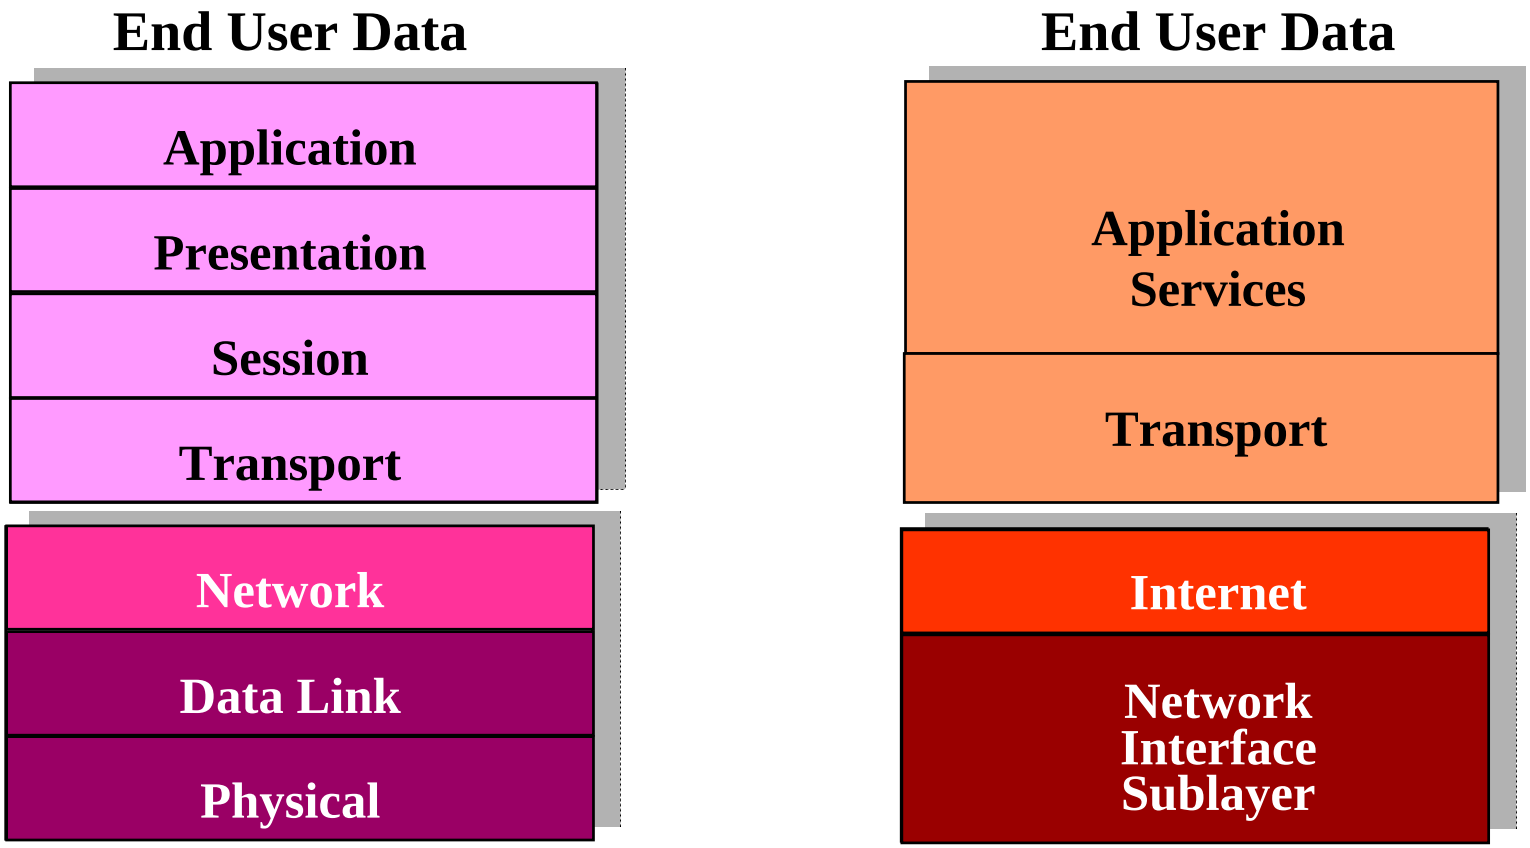

OSI & Fibre Channel Architecture

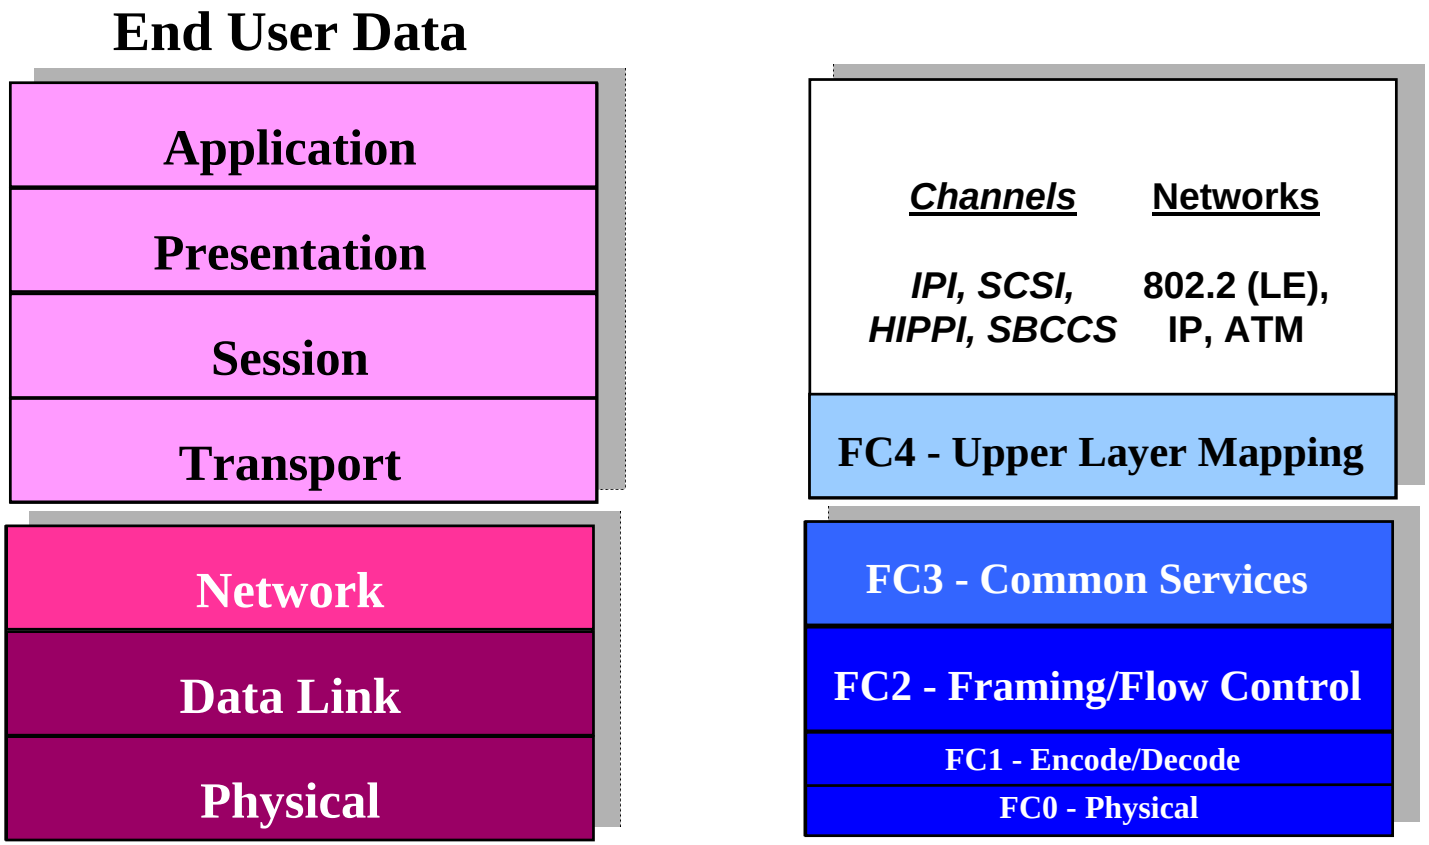

SCSI Bus

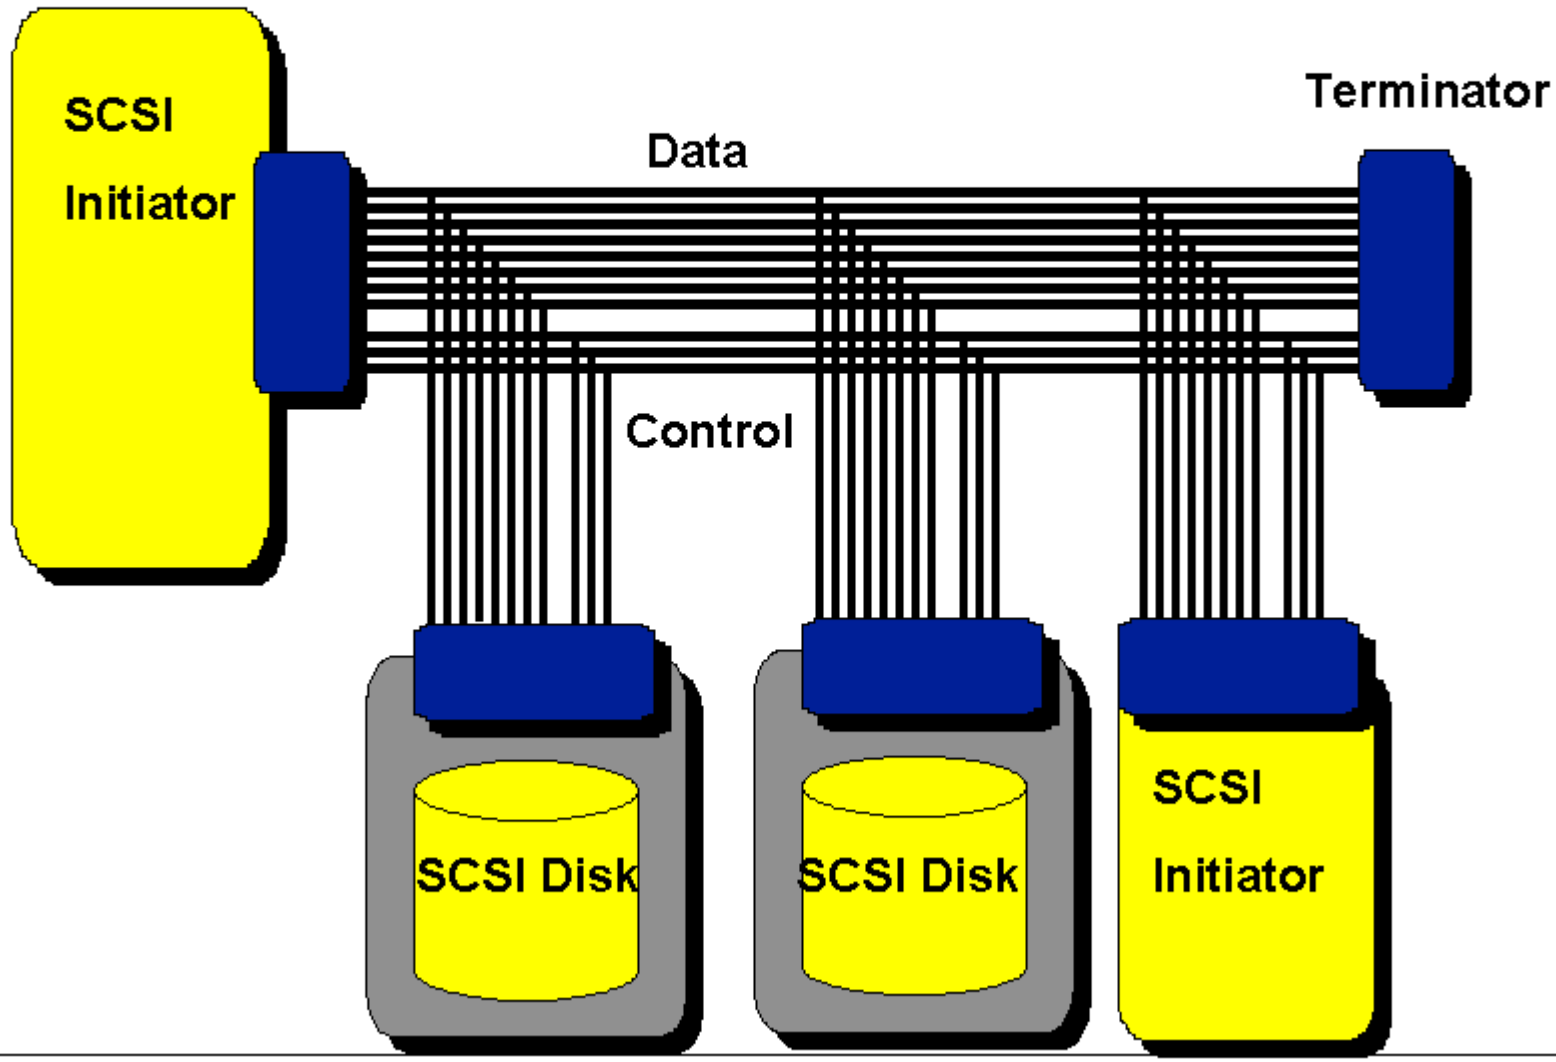
SCSI
Initiator
Terminator
Data
Control
SCSI Disk
SCSI Disk
SCSI
Initiator

SCSI

- Small Computer System Interface
- Small systems connectivity?
- Open systems workstation and servers
- Limited distance 12 – 25 meters
- Connectivity 16 devices
- Speed
  - 80 MBPS Ultra-2
  - 160 MBPS Ultra-3

ESCON

- Mainframe /enterprise computing
- Proprietary
- 20 MBPS
- Extended distance 43 KM
- Extended Connectivity

ESCON

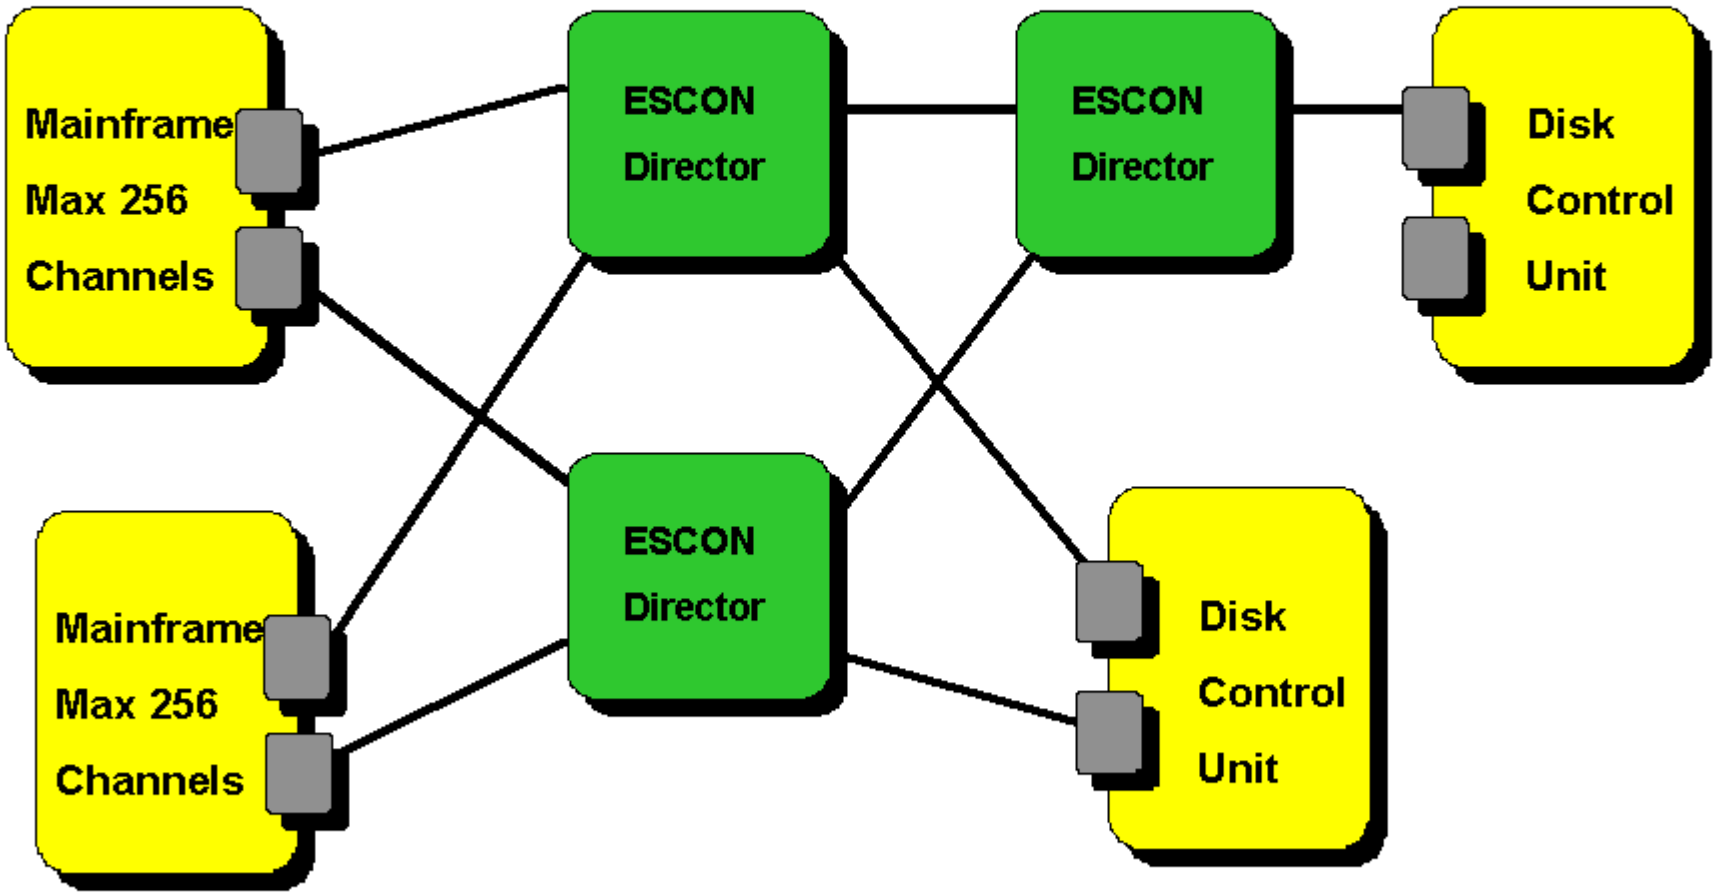
Mainframe
Max 256
Channels
Mainframe
Max 256
Channels
ESCON
Director
ESCON
Director
ESCON
Director
Disk
Control
Unit
Disk
Control
Unit

SAN Fibre Channel Topology

- A Topology is a set of hardware components such as media, connectors and transceivers connecting two or more node ports together
  - Point-to-Point, Arbitrated Loop, Fabric

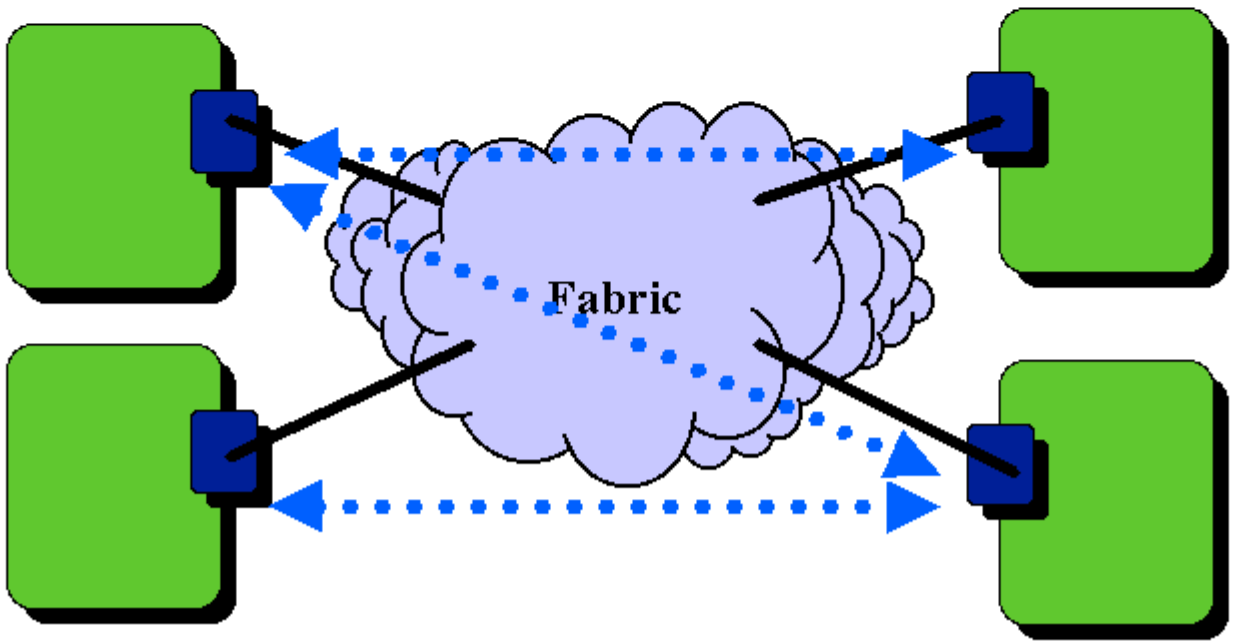

Topologies: Point-to-Point

- Point-to-Point
  - Bi-directional point-to-point serial full duplex channel
  - Fundamental Fibre Channel communications model
  - Peripheral storage applications
  - Extended distance

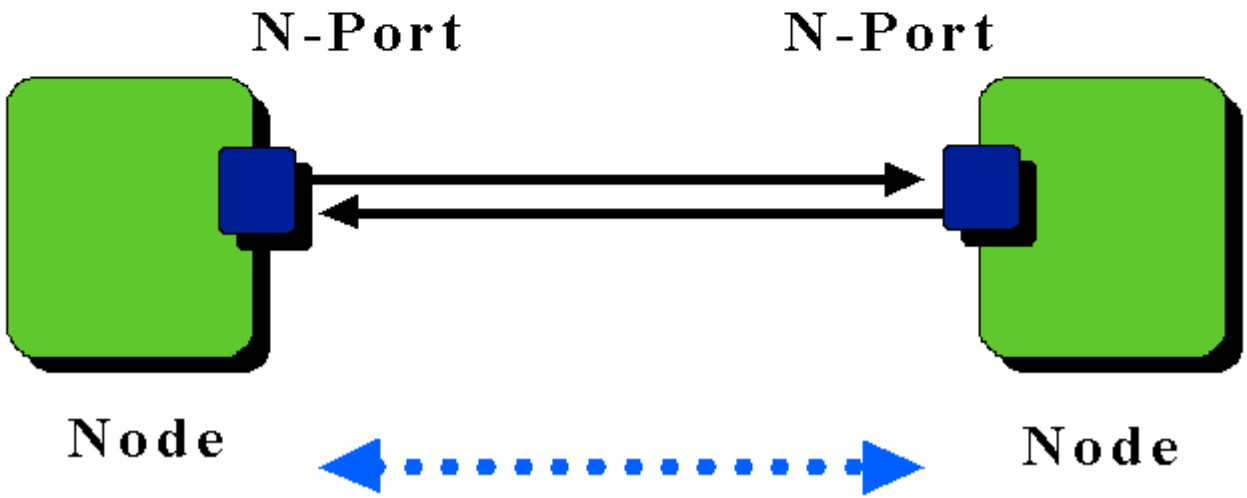

Point-to-Point

- Point-to-Point
  - Limited backbone
  - High performance full duplex system
  - Workstation, server to storage
  - Extended distance mapping
  - Video distribution
  - Specific application optimized
  - Simple implementation
  - Staging for later topologies

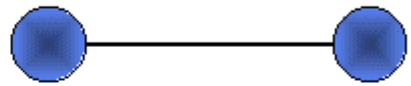

Topologies: Arbitrated Loop

- Arbitrated Loop
  - Max of one virtual point-to-point connection at a time!
  - Activities
    - » Initialize
    - » Arbitrate
    - » Open
    - » Information Flow
    - » Close

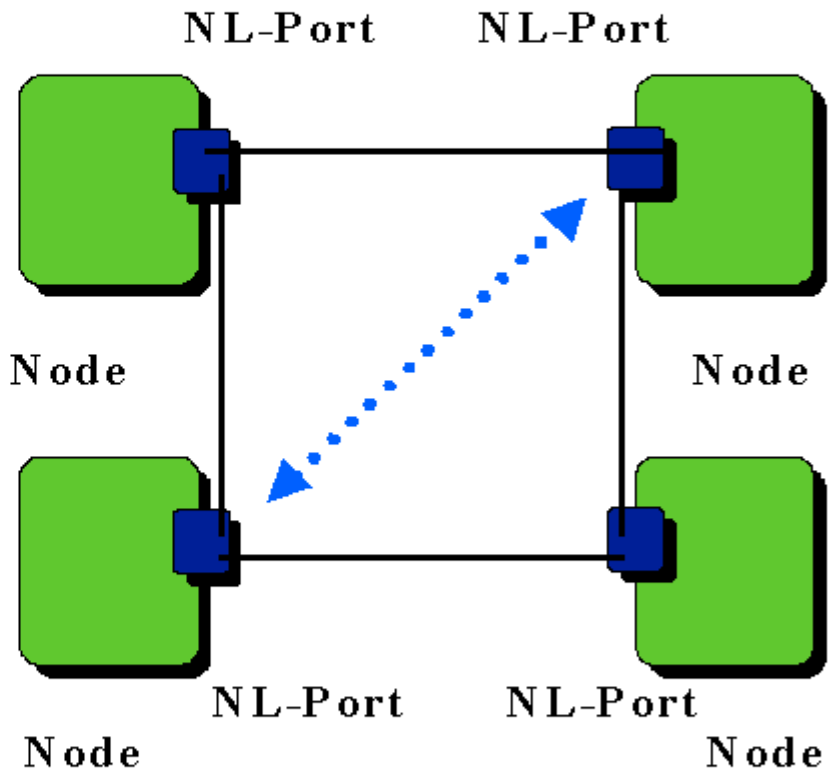

Arbitrated Loop

- Arbitrated Loop
  - Low cost, efficient interconnect
  - Storage subsystems (internal SAN)
  - Connectivity improvements
  - Shared bandwidth performance
  - Intelligent and non-intelligent hubs
  - Performance considerations
    - » Multiple initiators and targets
    - » Multiple protocol affects
    - » Multiple types of information flow
  - Hybrid loop extensions

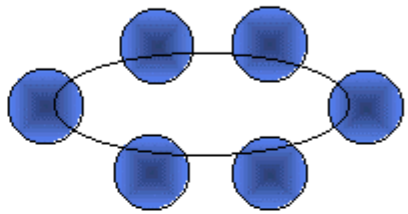

Topologies: Fabric

- Fabric
  - Multiple concurrent virtual point-to-point connections
  - Activities
    - » Initialize
    - » Information flow between partners
      - Logins
      - Commands
      - Data
      - Status

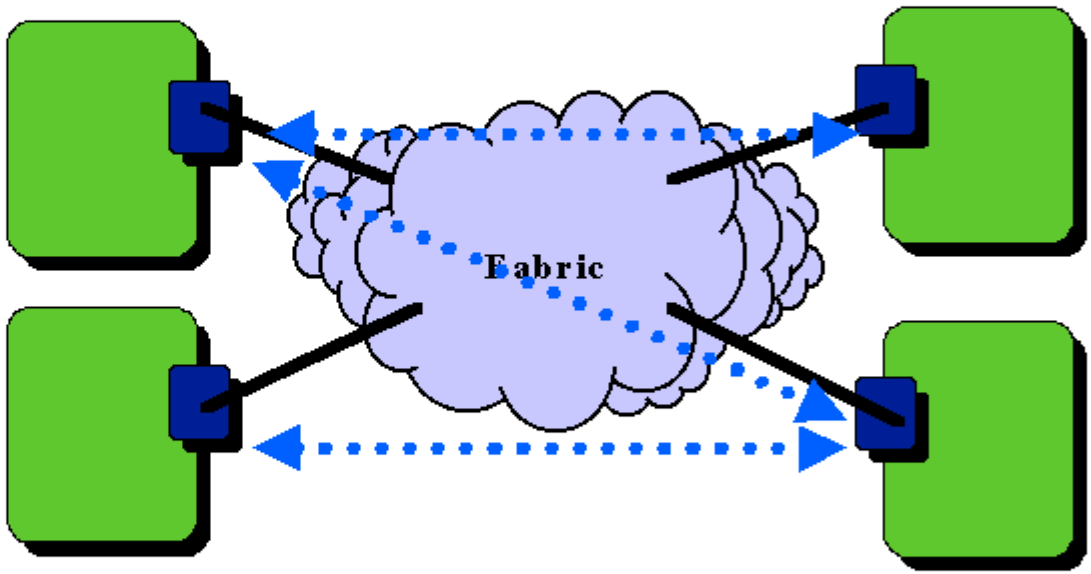

Fabric

- Fabric
  - Distributed systems
  - Multiple backbones
  - High availability
  - Multiple full bandwidth paths
  - Non-blocking performance
  - Higher cost

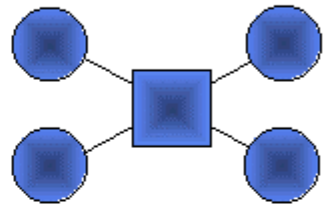

Port Types: FL Ports

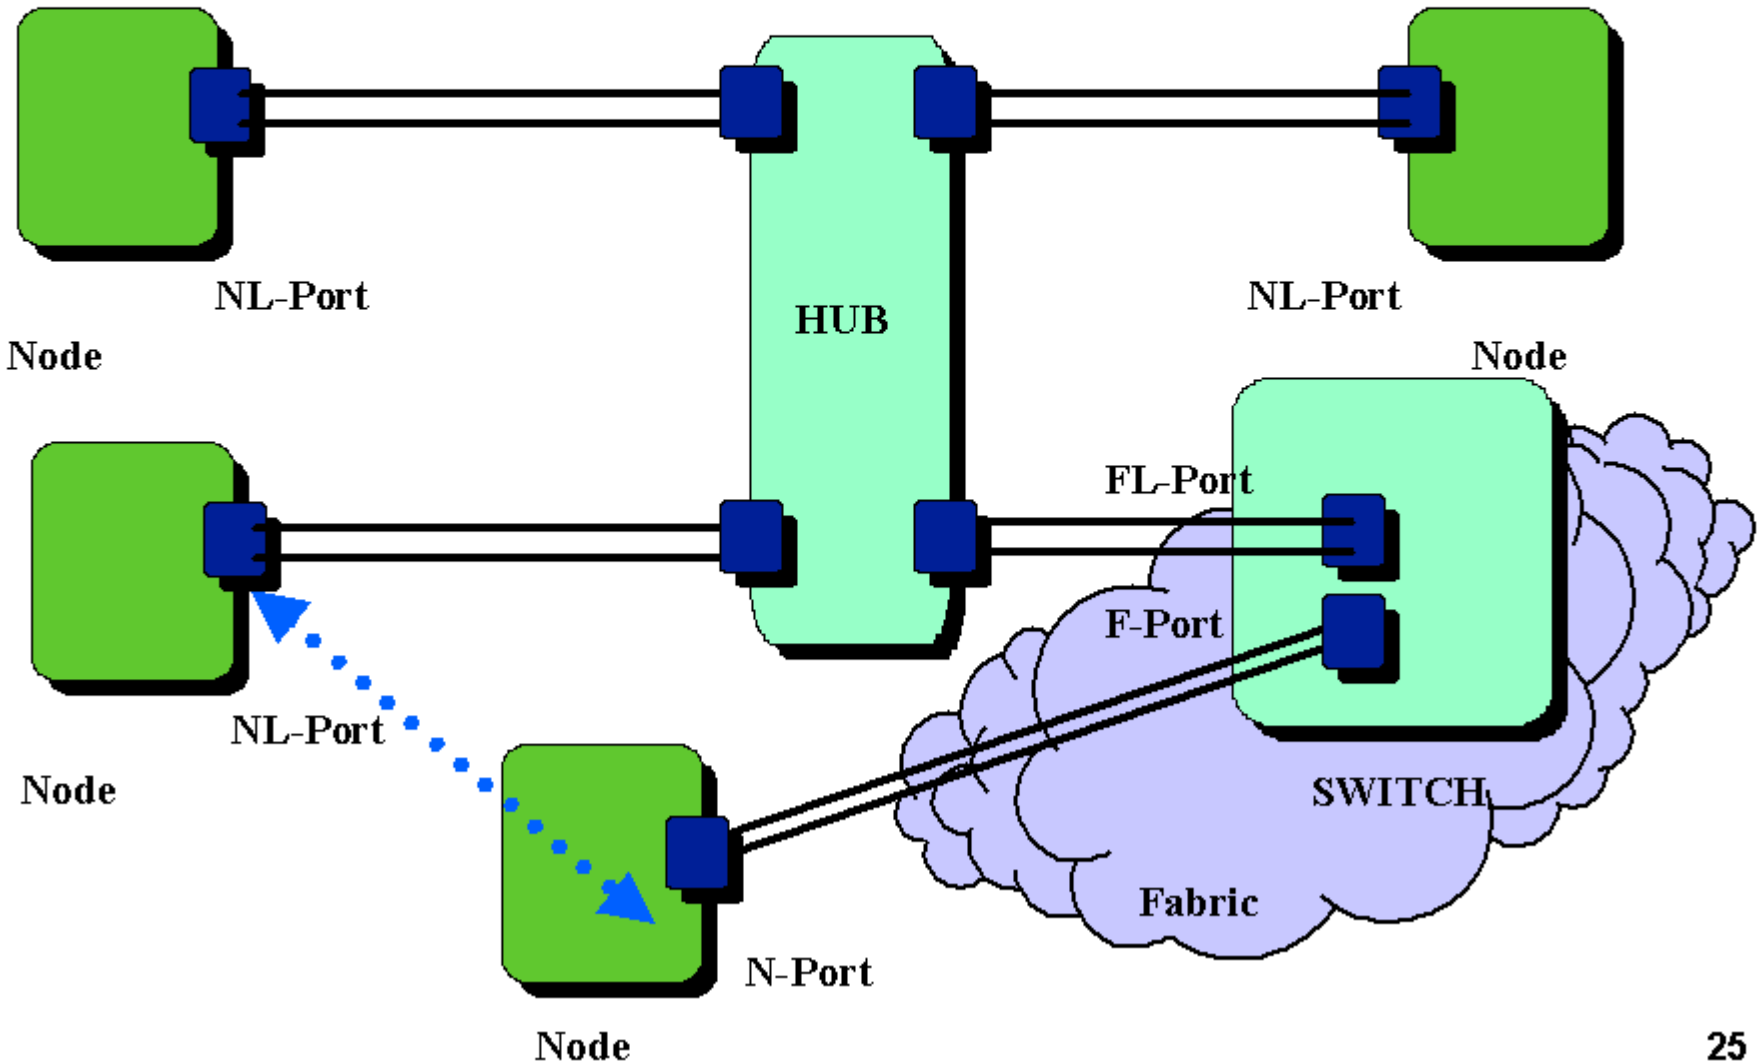

Fibre Channel vs GE

Category	Fibre Channel	Gigabit Ethernet
Bandwidth	100-200 MBps	100-200 MBps
Maximum Addressable Nodes	127 FC-AL 16 Million Fabric	2^{48} MAC adresses
Distance	33m Copper 500m MM Fibre 10Km SM Fibre	25m Copper 260m Fiber
Data Loss Due to Congestion	No Class 1, 2, 4, 6 Credit based flow control	Yes, however new standard 802.1 PQ
Throughput	High	High
Mode	Full Duplex Serial	Full Duplex Serial
Protocol	SCSI, IP, Hippi, ESCON, etc.	Multiple Network Layer Protocols

Fibre Channel over IP
IP over Fibre Channel

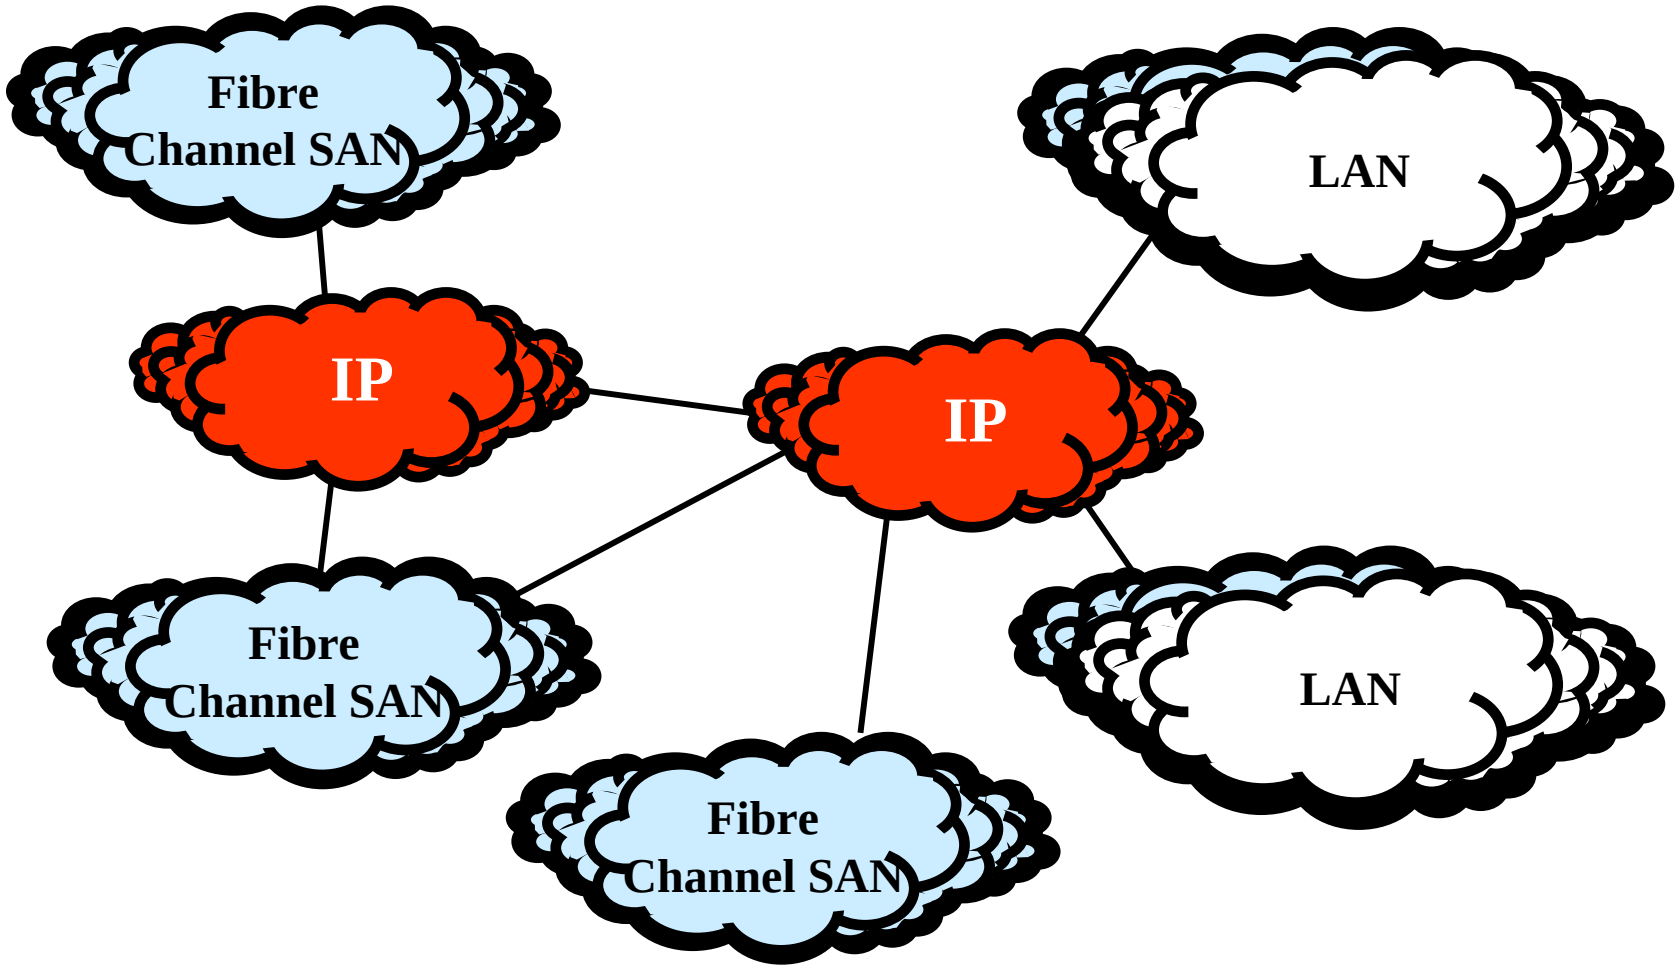

Networked Storage Future

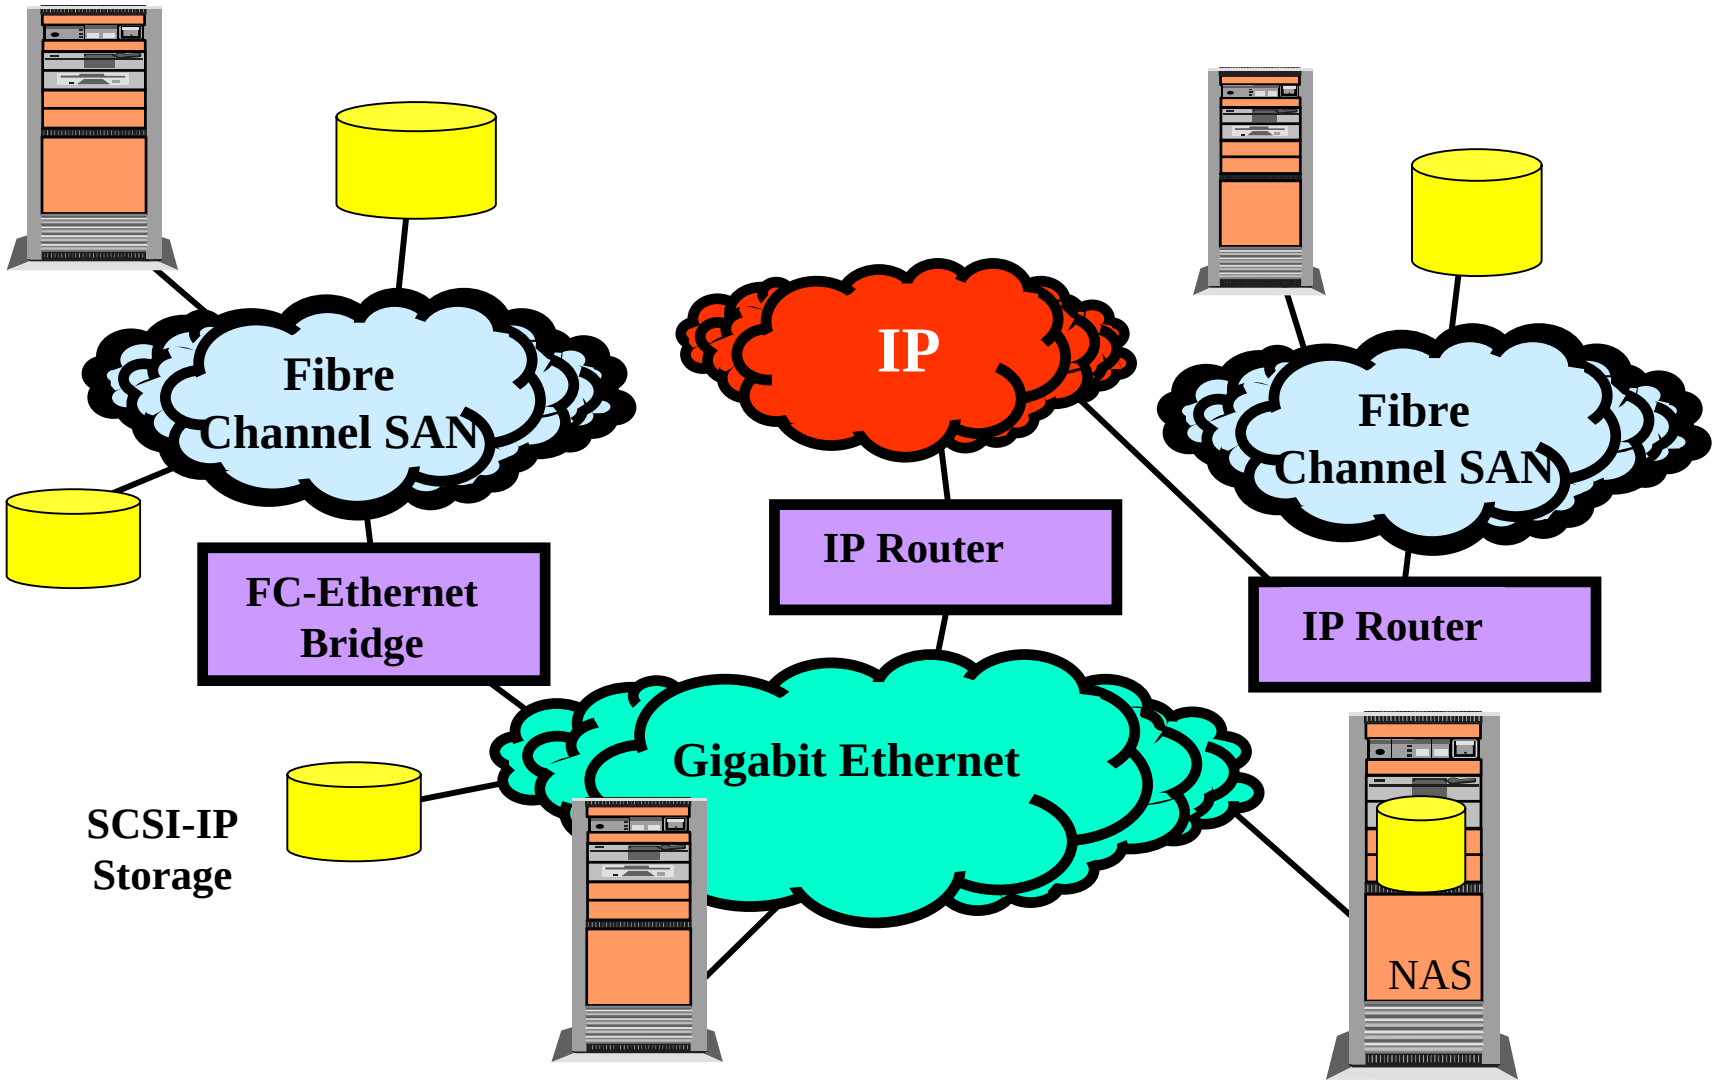

iSCSI and EtherStorage

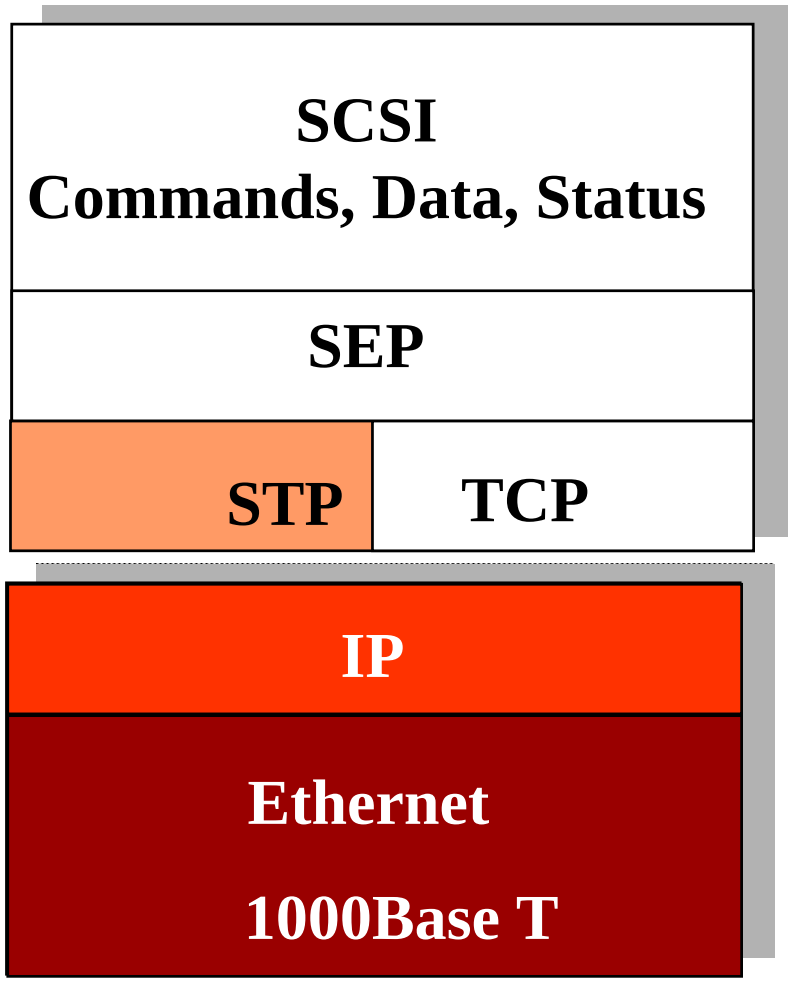

Fibre Channel vs IPS

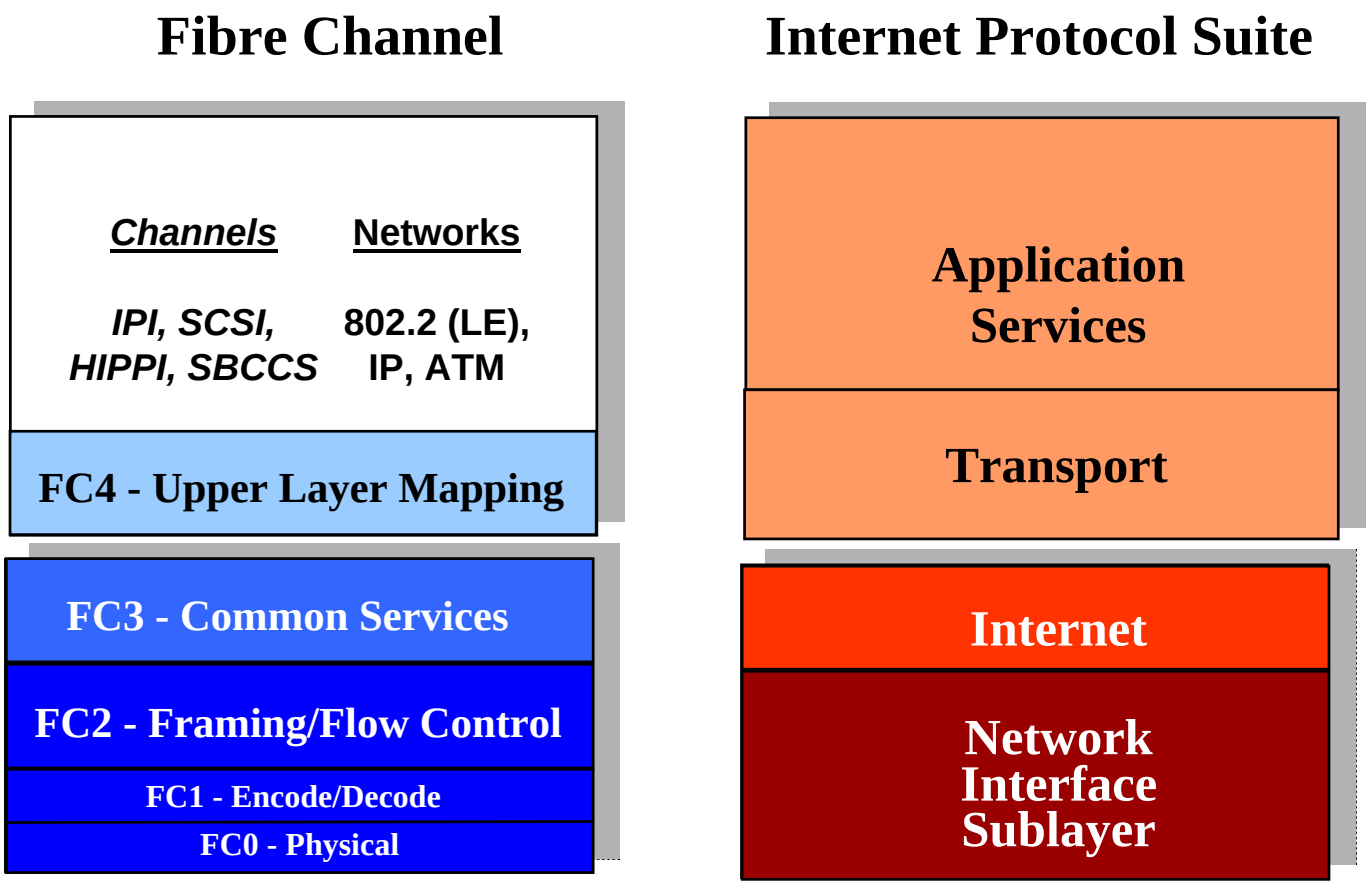

Traditional

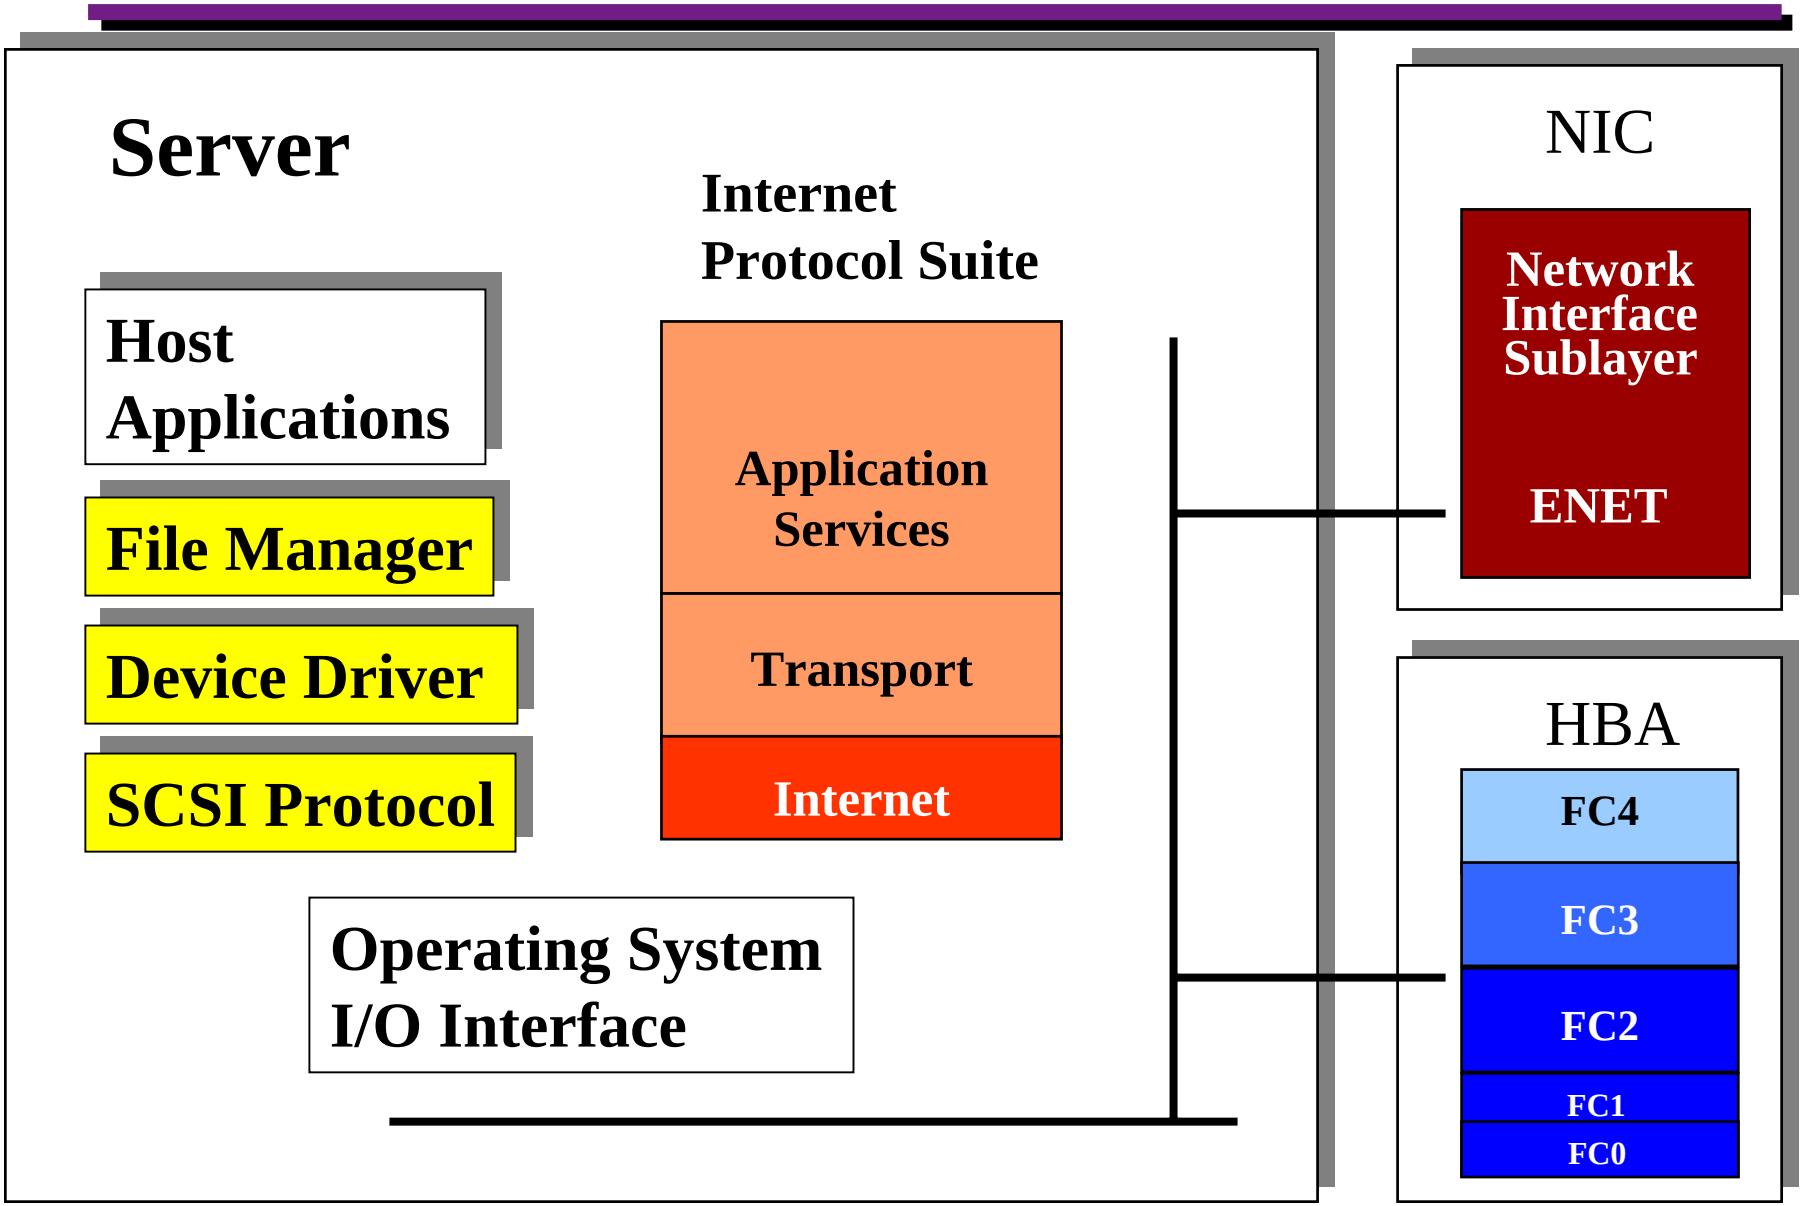

iSCSI

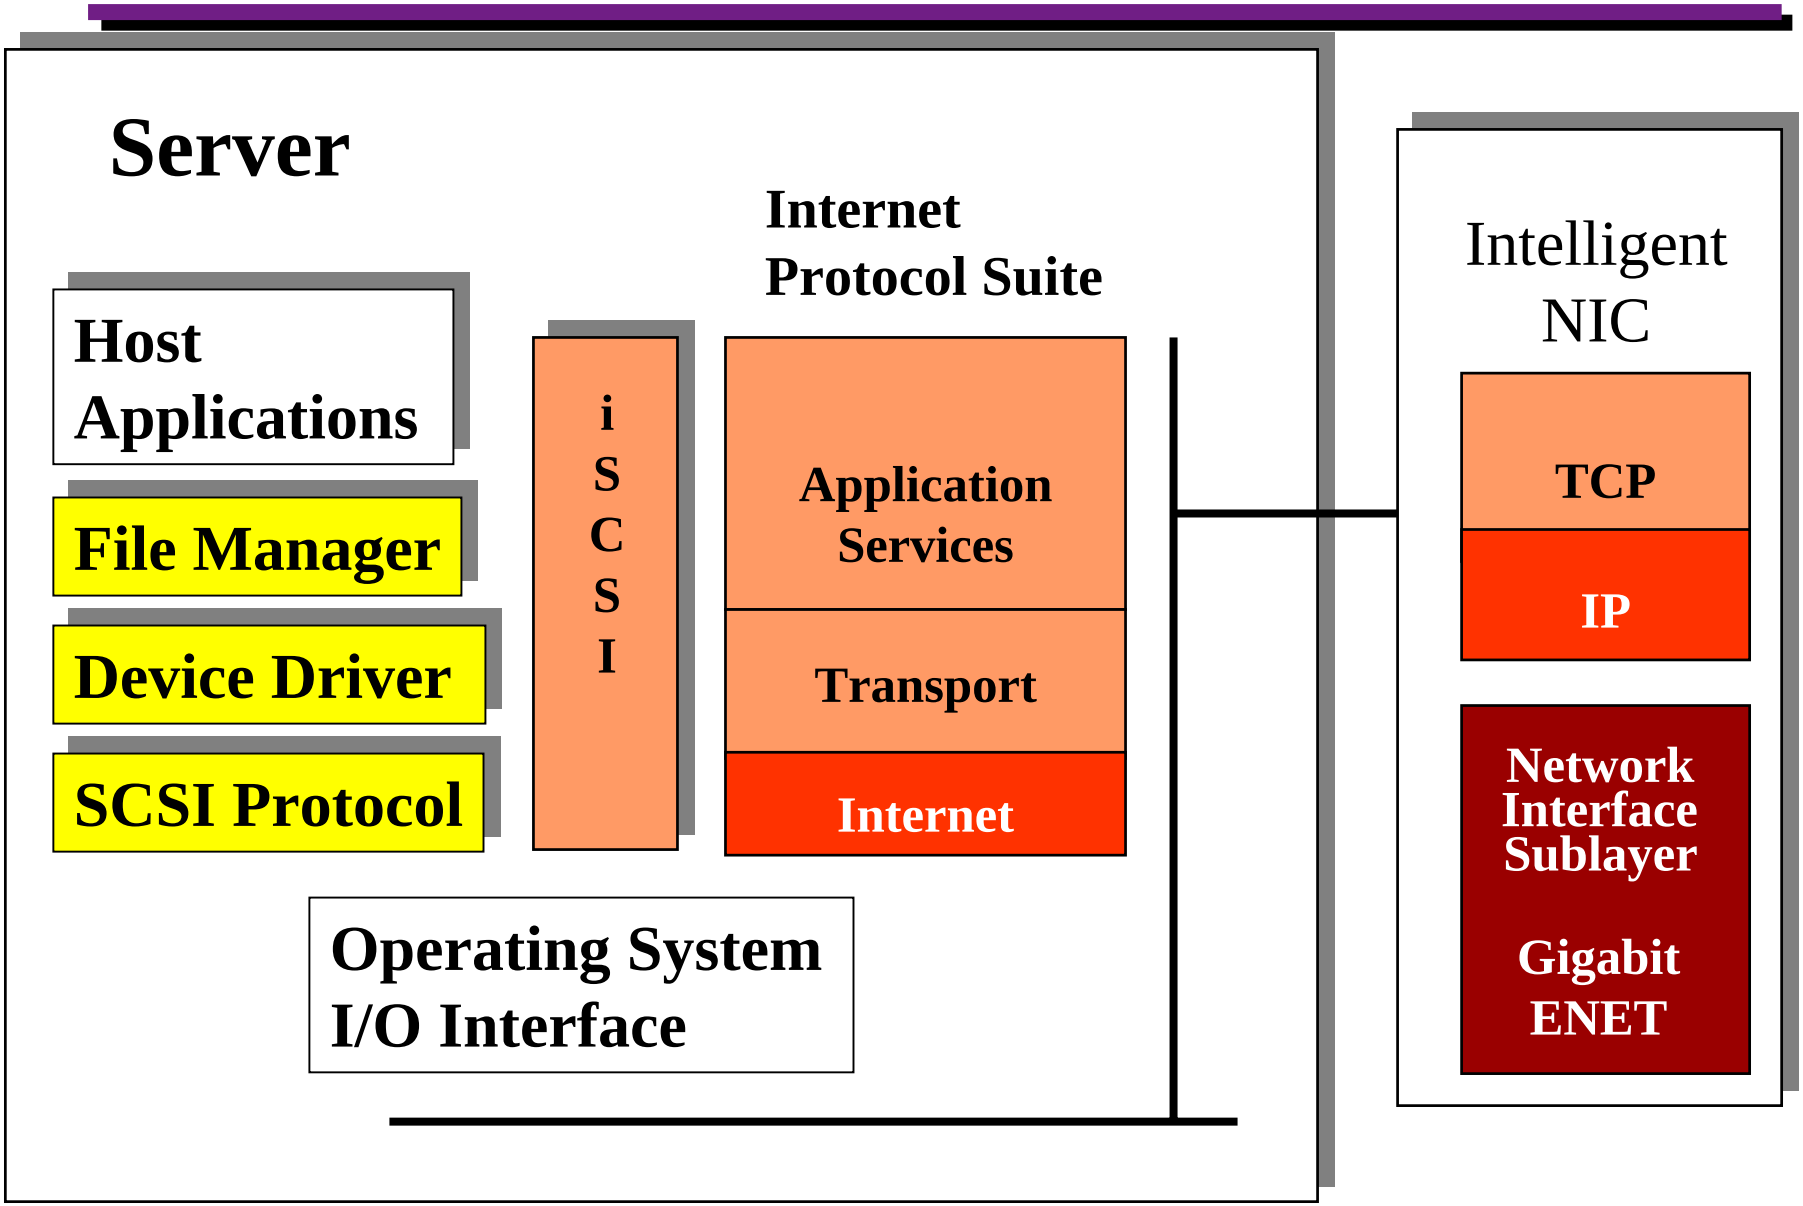
Server
Host Applications
File Manager
Device Driver
SCSI Protocol
i S C S I
Internet Protocol Suite
Application Services
Transport
Internet
Operating System I/O Interface
Intelligent NIC
TCP
IP
Network Interface Sublayer
Gigabit ENET

IP Storage Methods

Name	Integration Approach	Legacy Storage Compatibility	IP Network Compatibility	Companies Involved
IP Tunneling	FC frame encapsulation	FC	(Tunneled)	Gadzoox, CNT, Brocade, Cisco
Ether Storage	SCSI mapped to STP (Proprietary L4)	SCSI	Layer 1-3	Adaptec
I-SCSI	SCSI mapped To TCP		Layer 1-4	IBM, HP, EMC, Cisco, others
SoIP	SCSI/FCP adaptation	SCSI & FC	Layer 1-4	Nishan, others

Server I/O PCI Bus

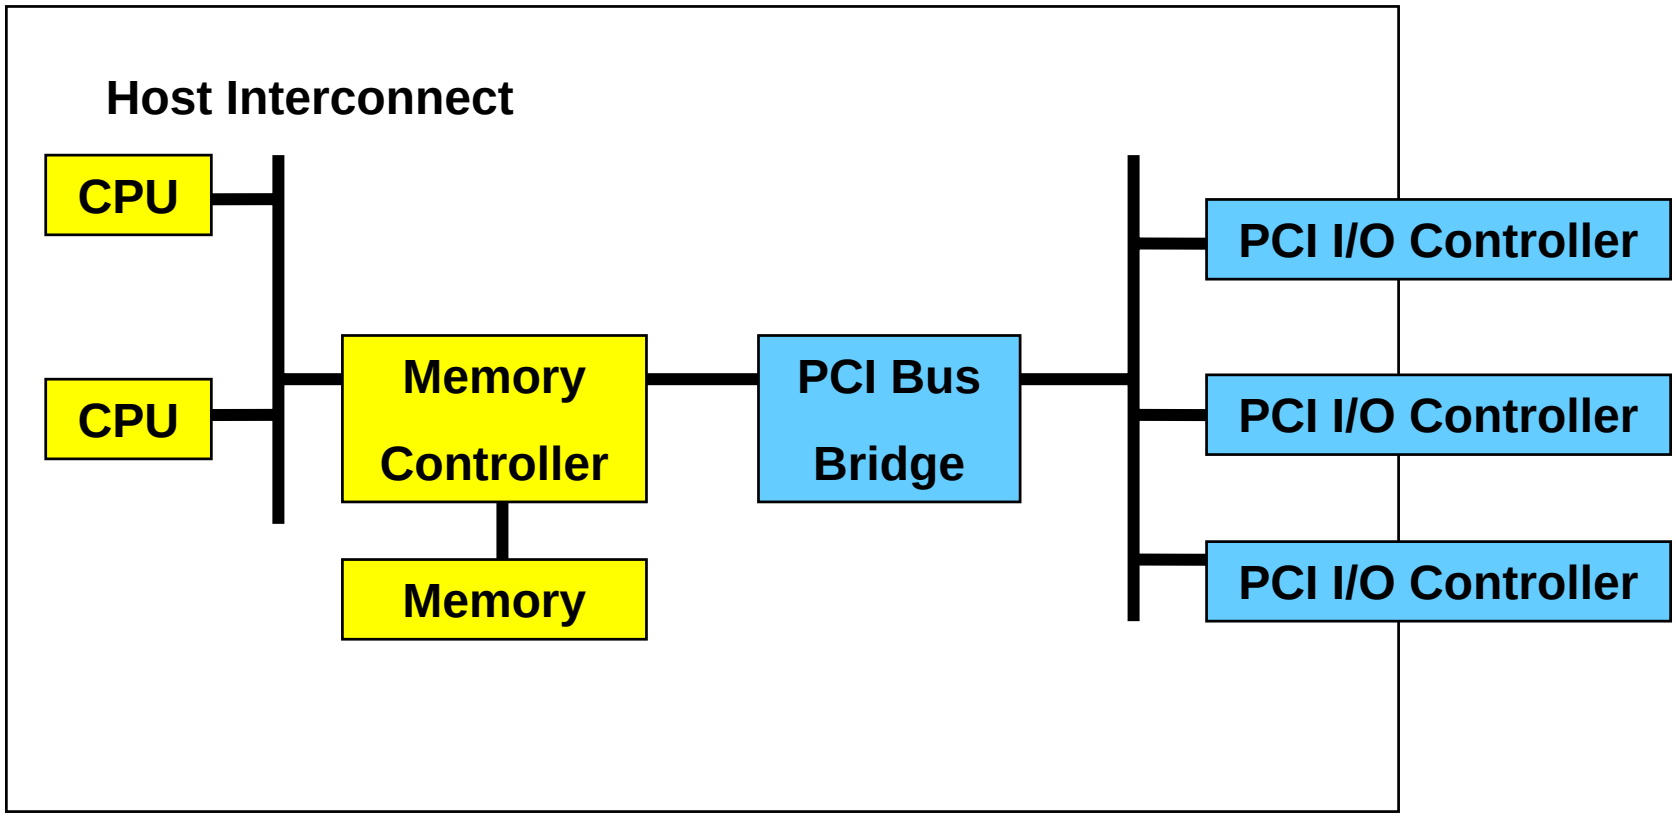

Infiniband Architecture IBA

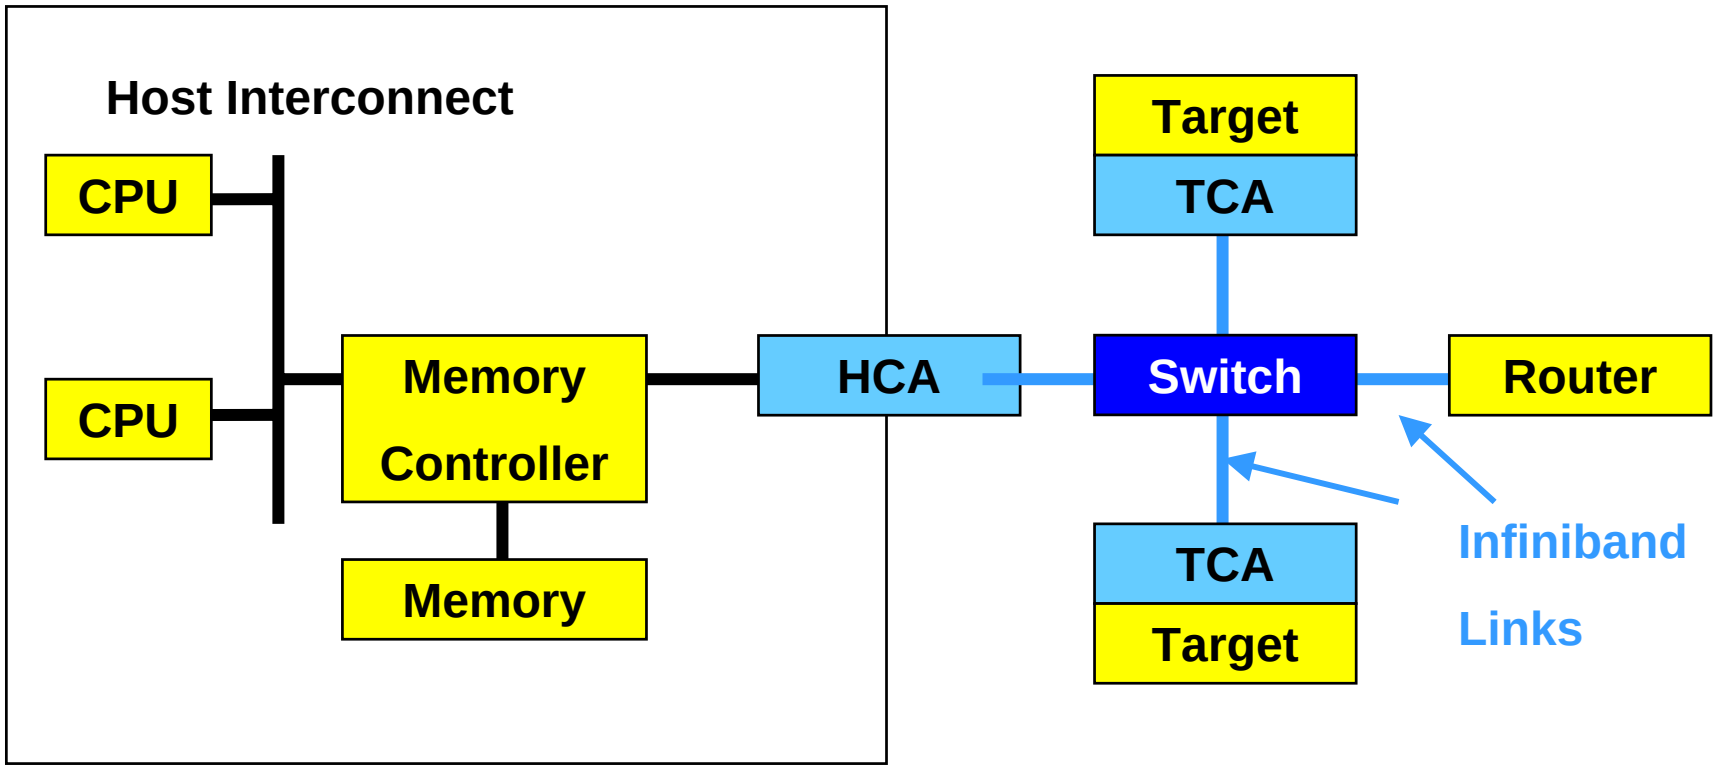

Targets: SCSI Bus, Fibre Channel, Ethernet

Infiniband Positioning

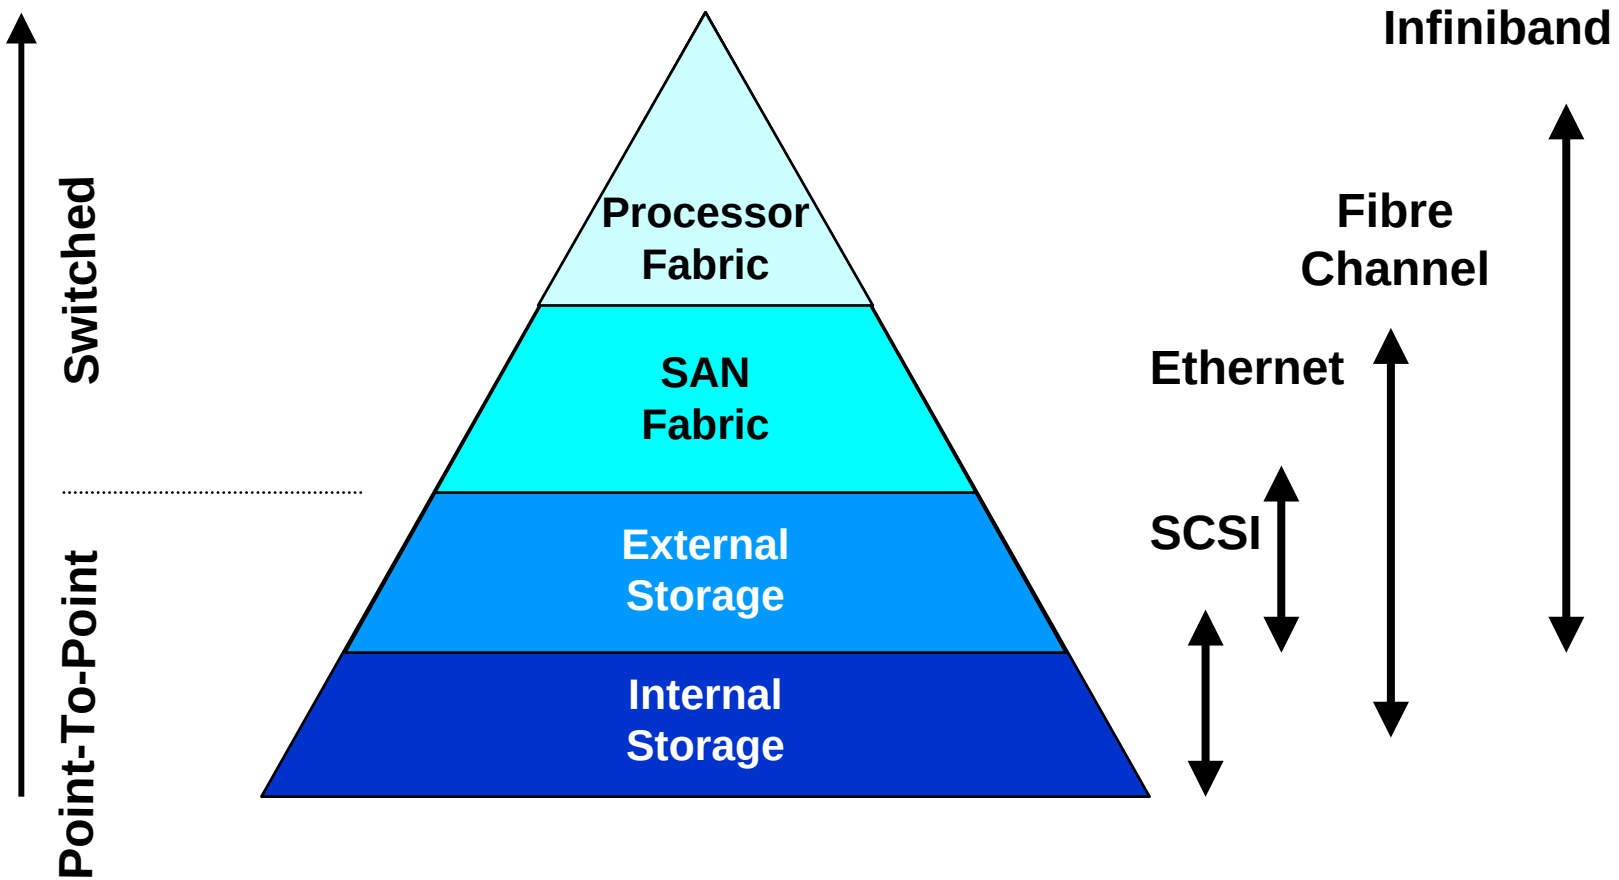

Networked Storage Futures

- Greater process login for multiple upper layer protocols
- Class 4 fractional bandwidth
- Class 1 dedicated simplex
- Stacked connect and buffered class 1
- Data compression and encryption
- Greater usage
- Co-existence with many underlying architectures
- Edge routing to and through technologies

Networked Storage Futures

- Upper layer protocol changes
  - Lazy and Hyper reads and writes
  - System SAN awareness
- Higher bandwidth rates for aggregation and higher speed processors and adapters
- SAN Over
  - Fibre Channel
  - IP
  - GE
  - Infiniband
- NAS & SAN integration

Networked Storage

- Thank you for attending
- Good luck with Networked Storage Technologies

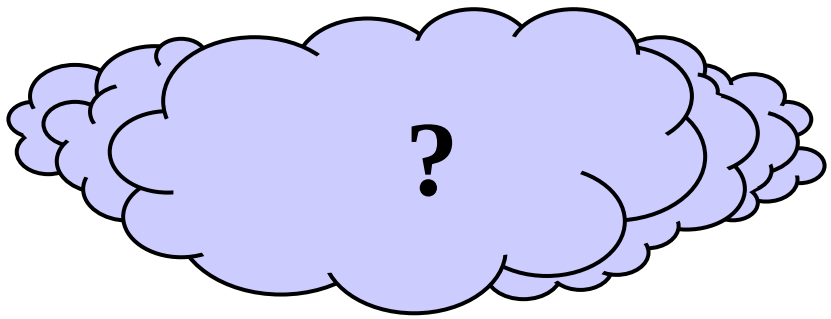

For more information

- SCSI Trade Association
  - www.scsita.org
- Fibre Channel Industry Association
  - www.fibrechannel.com
- Storage Network Industry Association
  - www.snia.org
- Building Storage Networks;
  - by Marc Farley
- SAN Solutions and Technologies Tutorial
  - NetWorld + Interop