Session 1111

Christopher Greene
Manager Storage Services
Hewlett-Packard Services

HP WORLD 2003
Solutions and Technology Conference & Expo

Agenda

- Introduction
- Traditional Layered Models
- Storage Infrastructure Model
- Application of SIM
- Q & A

Introduction

Introduction

- The Networking Industry has long used layered models to better conceptualize, administer, and fault isolate networks of all types and topologies.

- Storage Networking has evolved from a Channel-Attach concept to a full functioning network. As Storage Networking has evolved, models have been developed for specific technologies or protocols of Storage Networking.

Introduction

- There are no all-encompassing models in Storage Networking that illustrate the storage network in its entirety as in Traditional Networking Models.
- Storage Networking has much to gain through the use of layered models and, in specific, the Storage Infrastructure Model.

Traditional Networking Models

Traditional Network Layered Models

- Traditional layered networking models break network communications into several layers allowing for the separation of functions and responsibilities.
- Each layer operates independently
- This approach provides the ability to conceptualize components and functions necessary for network communications
  - Greater understanding of network for administrative purposes
  - Greater understanding of network for design and expansion purposes
  - Faster isolation in fault isolation scenarios

Open Systems Interconnect OSI Model

The most referenced layered model

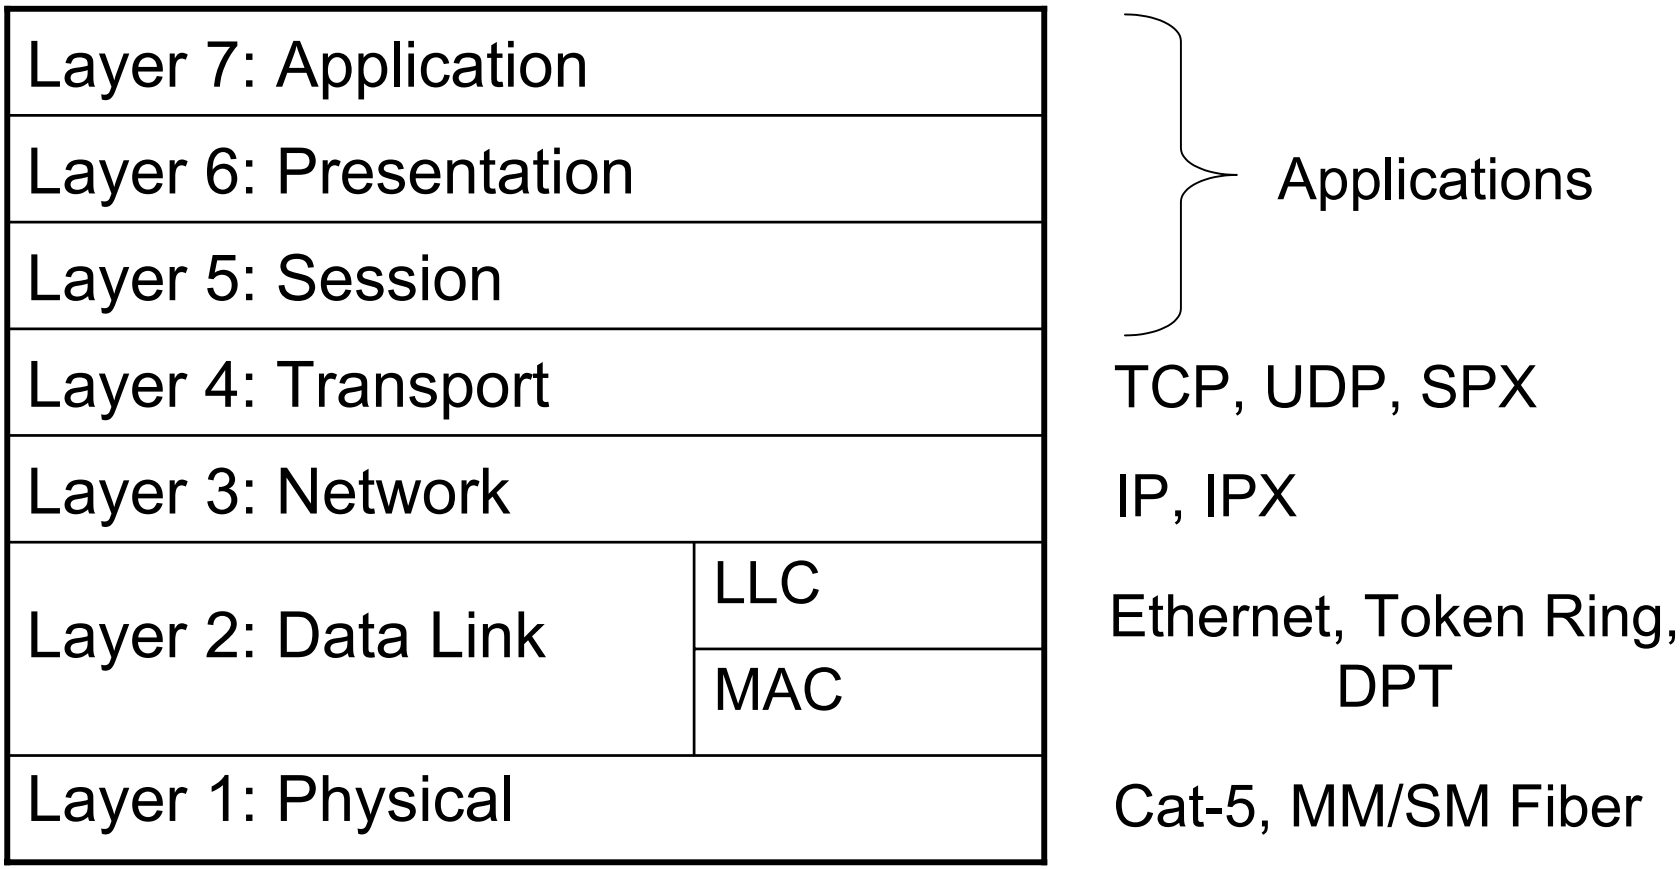

Department of Defense DOD

A less common but valuable layered model

Process / Application	Applications and upper protocols
Host to Host	TCP, UDP, SPX
Internet	IP, IPX
Network Access	Cat-5, Ethernet, MM/SM Fiber

Storage Layered Models

Storage Layered Models

- As various Storage Technologies were developed, corresponding models were created for the specific protocols or topologies
- Storage Layered models, while serving their specific purpose, do not capture the entire communication process from application to storage itself.
- Missing pieces not captured in existing storage models:
  - File systems
  - Virtualization
  - Encapsulation

SCSI Architecture Model SAM-2

- Excellent representation of the most fundamental piece of storage – SCSI.
- Clearly diagrams the two references to SCSI:
  - SCSI protocols and commands
  - SCSI interfaces.
- Does not capture processes outside the SCSI layer.

SCSI Architecture Model SAM-2

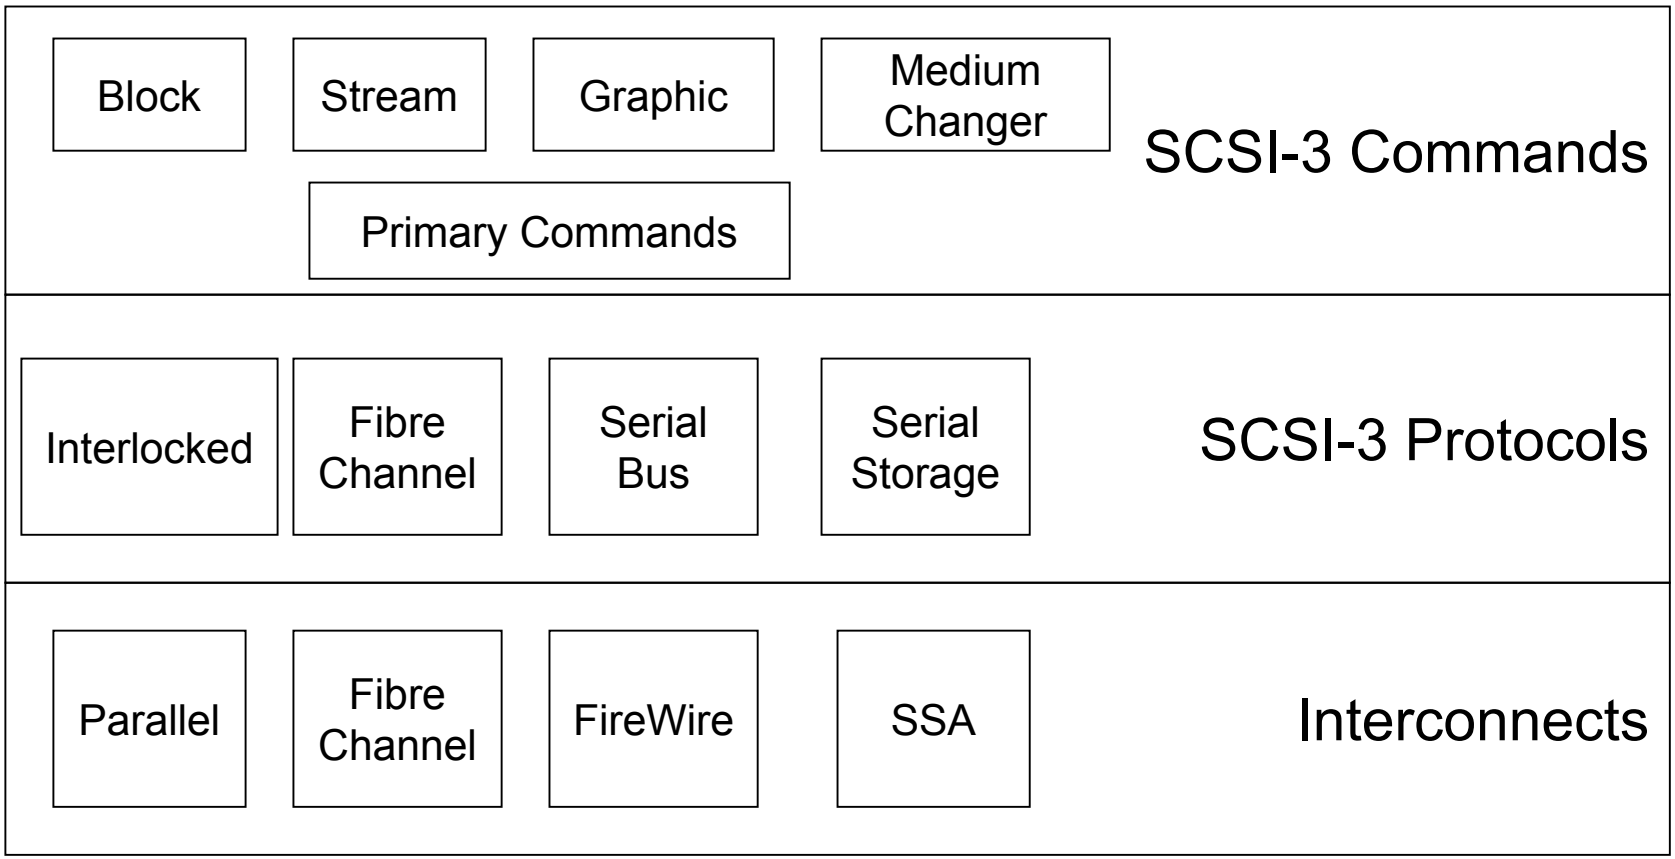

Fibre Channel

- Fibre Channel has proven to be the most widely deployed infrastructure of Storage Networks.
- The Fibre Channel model clearly diagrams the two references to Fibre Channel:
  - Fibre Channel transport medium – the physical layering, framing, and encoding
  - Fibre Channel Protocol (FCP) – serialized SCSI protocol
- The Fibre Channel model is very similar to the OSI model in structure, but does not give a clear understanding of the processes beyond Fibre Channel itself.

Fibre Channel

FC4: Upper Layer Protocols	SCSI-3 (FCP)
FC3: Common Services	TBD
FC2: Data Delivery	Framing, Classes of Service
FC1: Ordered Sets / Encoding	8b/10b encoding
FC0: Physical Interface	MM fiber, Copper

Mapping Storage Protocols to Networking Models

- Mapping Storage protocols to traditional layered models allows for a greater understanding of a protocol's role and function
- This mapping allows for easier fault isolation
- There are gray areas of layer mapping such as routing/switching of Fibre Channel
- Mapping Storage Protocols to traditional models leaves many unfilled holes in the models

SCSI Protocol

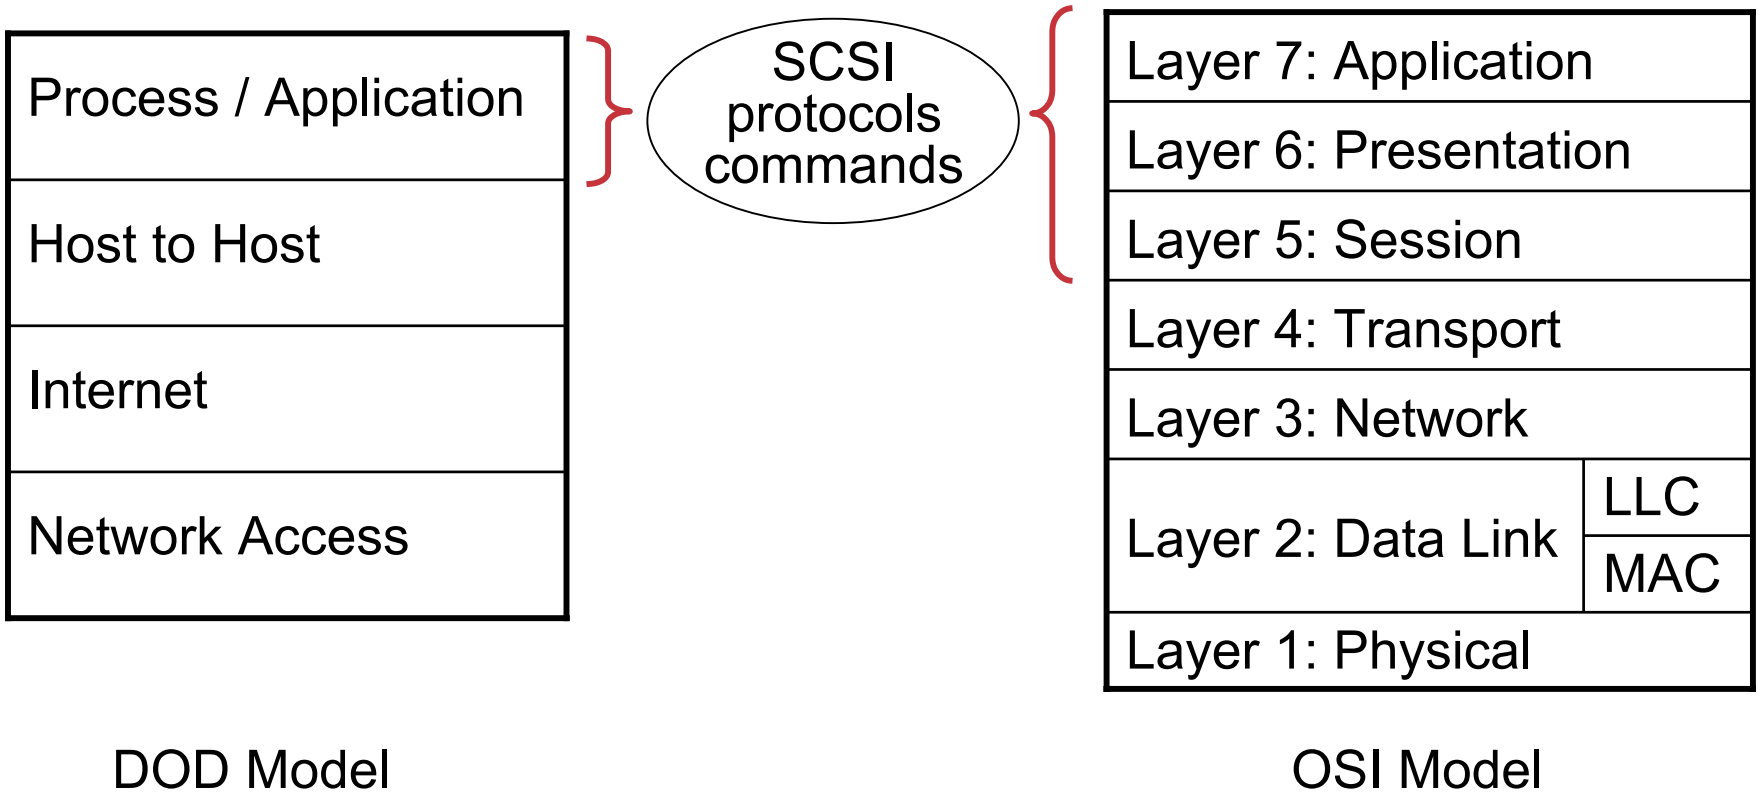

SCSI Protocol

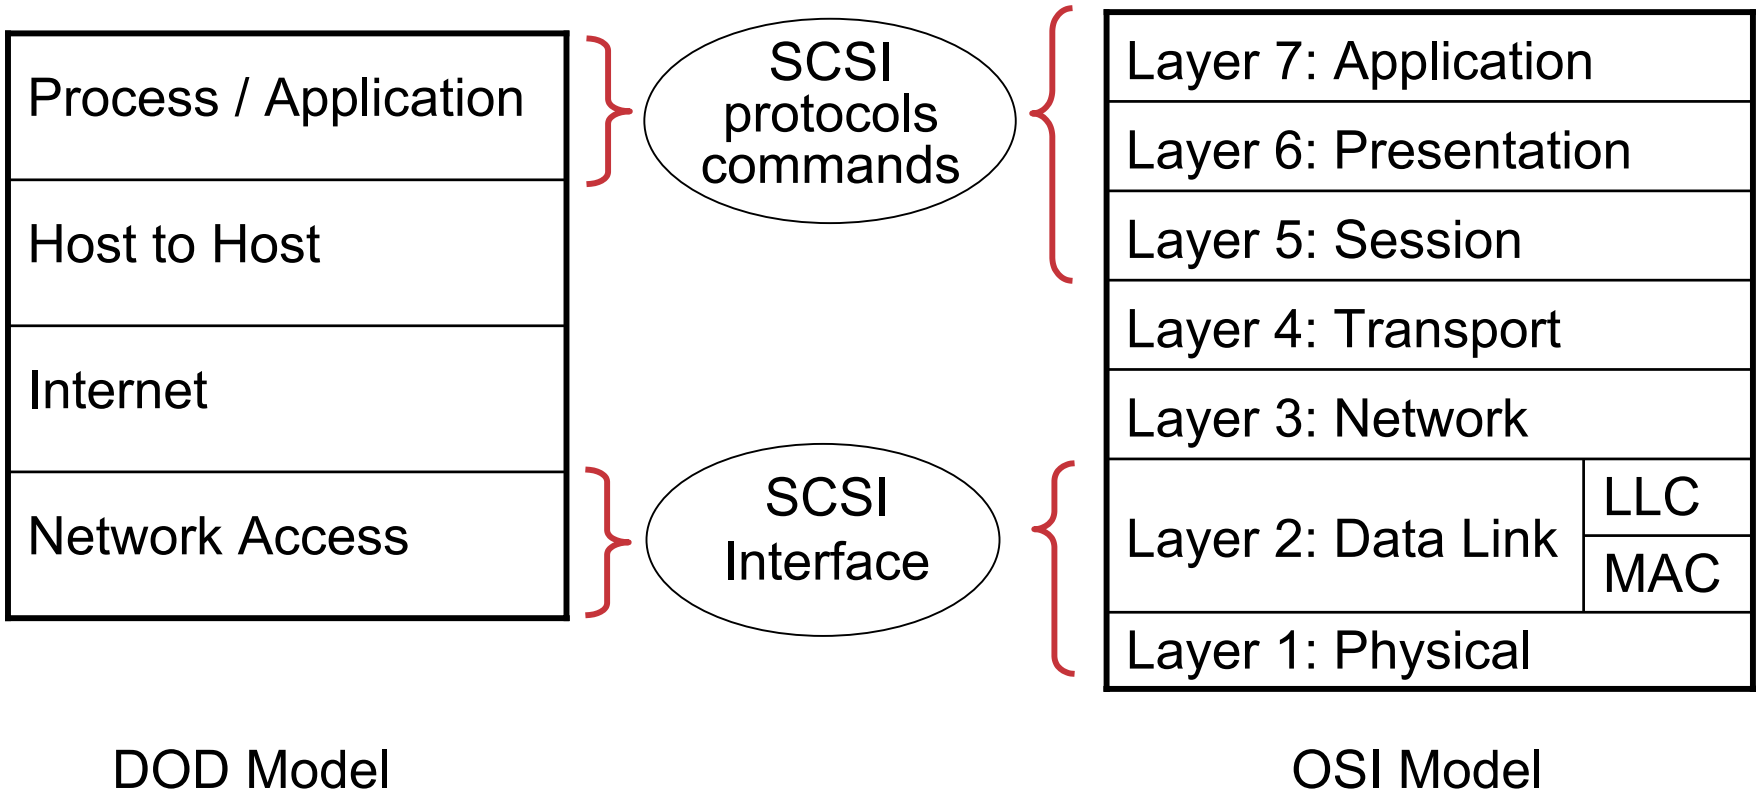

Fibre Channel

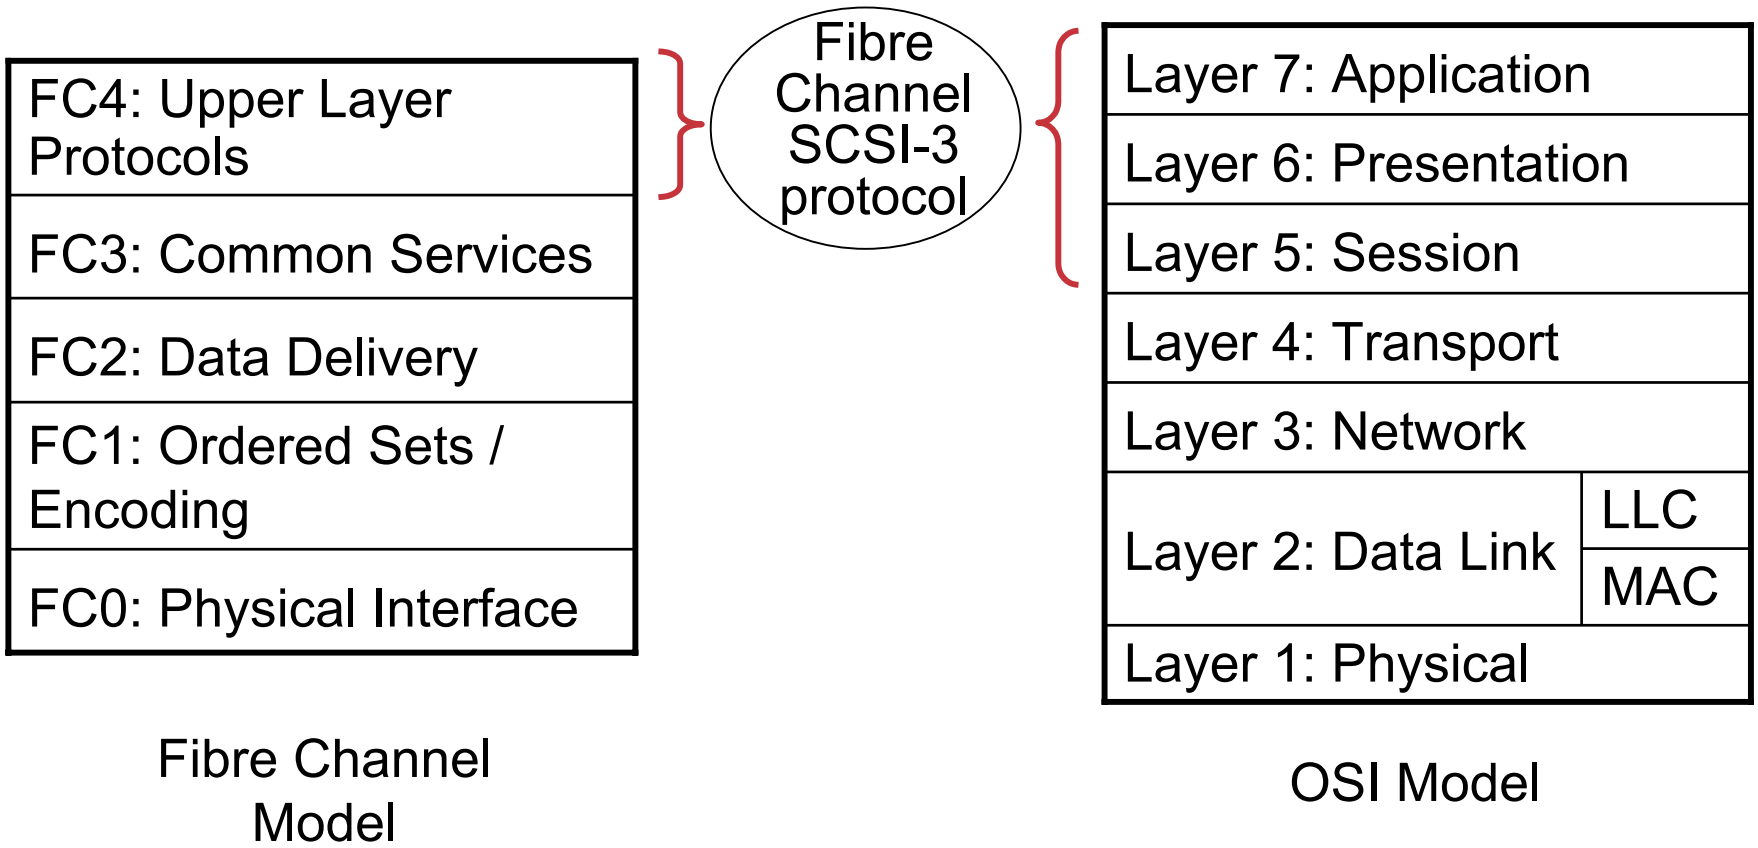

Fibre Channel

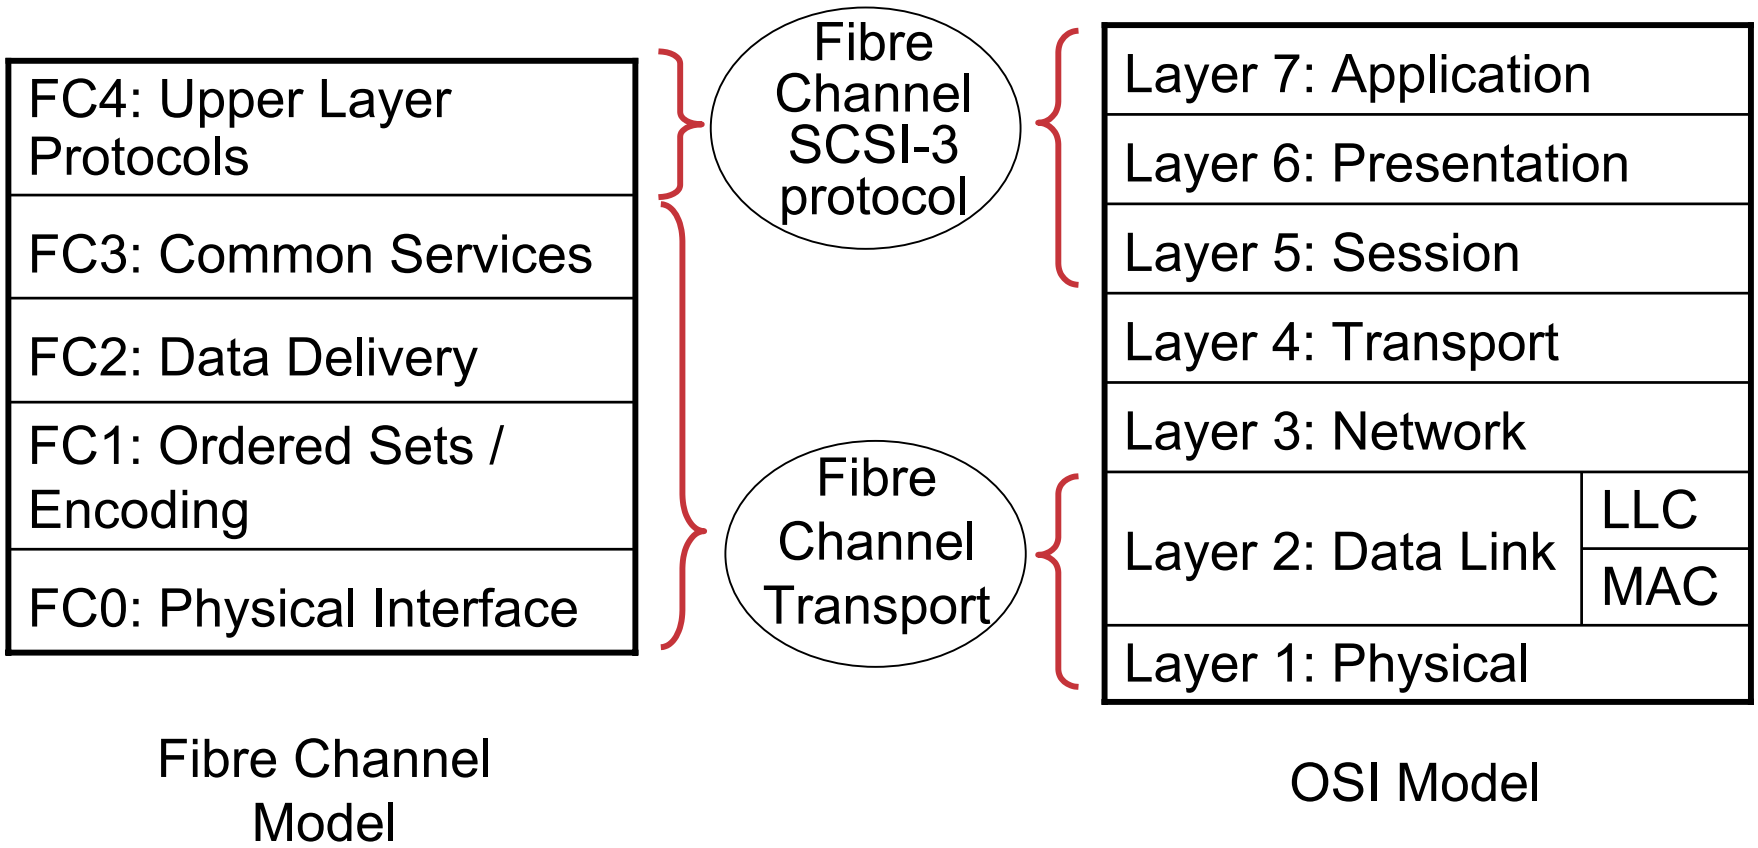

The Answer:

Storage Infrastructure Model

Storage Infrastructure Model SIM

- An overall model is needed to map the functions and layers necessary in a Storage Solution

- SIM is able to incorporate older technologies as well as newer technologies:
  - Internal
  - Direct attached Arrays
  - FC SANs
  - iSCSI
  - NAS

Storage Infrastructure Model SIM

- SIM is able to capture functions outside of storage specific technologies:
  - File systems
  - Virtualization
  - Application integration
- SIM was developed jointly between the Field Interoperability and Test (FIT) lab and field personnel:
  - Tested rigorously against all major platforms and theoretical limits
  - Tested in real-life fault isolation scenarios
  - Used among engineers to quickly isolate and demonstrate theories and concepts

Storage Infrastructure Model SIM

SIM contains four standard layers

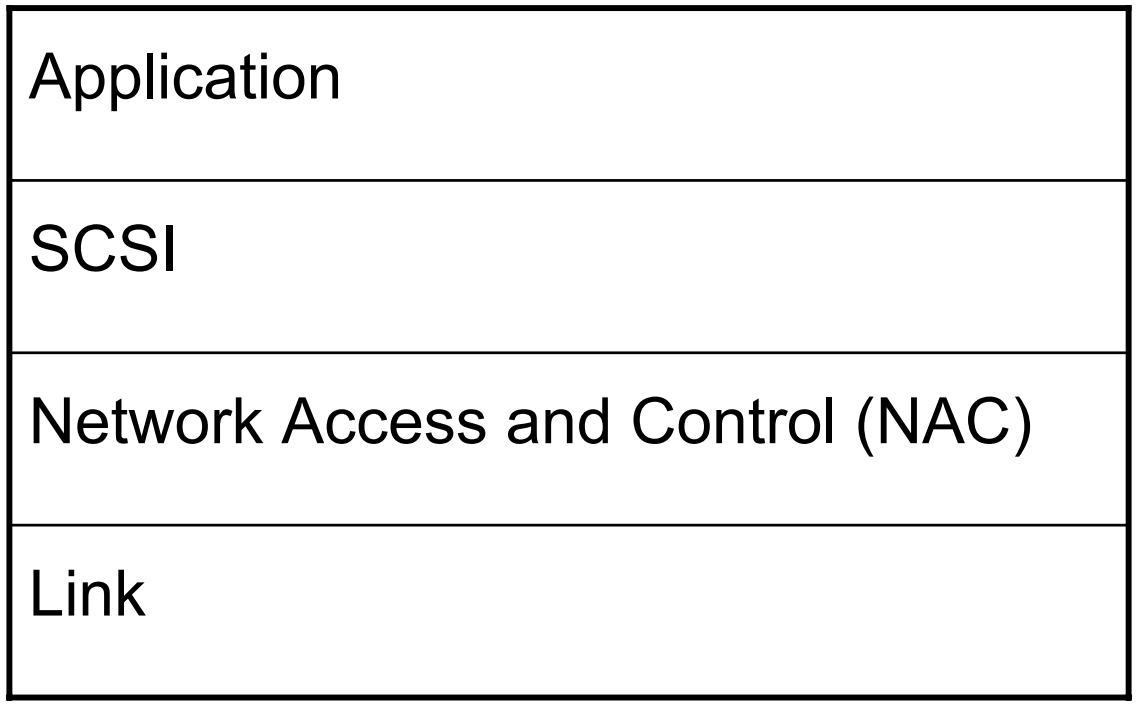

Storage Infrastructure Model SIM

SIM also has two additional layers

Storage Infrastructure Model SIM

Guidelines

- The layers of SIM are not to be as rigid in form or placement as OSI, DOD, or FC models
- Layers may be removed, inserted, and subdivided as necessary
- SIM's purpose is to enable the conceptualization of a Storage Solution, therefore it must be adaptable for various possibilities
- Unlike other models where one model representation is sufficient, it is necessary to present SIM in full context from Storage to Host

Storage Infrastructure Model SIM

Link Layer

- Physical layer comprising of physical transport medium, electrical, and fiber optic components

- Analogous to the following layers from traditional models:
  - FC0 from the FC model
  - Layer 1 (Physical) of OSI
  - Part of Interconnect from SAM-2
  - Part of the Network Access from DOD

Storage Infrastructure Model SIM

Link Layer

- Example components:
  - MM fiber
  - SM fiber
  - Category 5
  - GBICs (long wave and short wave)
  - Fiber repeaters
  - DWDM solutions and multiplexers
  - Modulation schemes

Storage Infrastructure Model SIM

Network Access and Control (NAC) Layer

- NAC is the widest layer of SIM
- The NAC layer contains the protocols and technologies necessary to transport storage communications on the Link layer, maintaining data integrity, and managing complex communications requirements
- NAC is analogous to the following layers from traditional models:
  - Layer 2, 3, and 4 from OSI
  - FC1, FC2, and FC3 from FC
  - Internet and Host to Host layers from DOD
  - Remaining parts of Interconnet from SAM-2

Storage Infrastructure Model SIM

Network Access and Control (NAC) Layer

- It is necessary to subdivide the NAC layer into its individual components upon each implementation

- Example Components:
  - Ethernet
  - Fibre Channel Transport
  - TCP/IP
  - SCSI Interface (beyond physical cabling captured in Link layer)

Storage Infrastructure Model SIM

SCSI Layer

- SCSI-3 Command/Protocol
- The basic elements of storage communications
- Necessary to understand that all block level I/O exists with basic SCSI-3 command set
- Encompasses both SCSI-3 Protocol and SCSI-3 Command layers of the SAM-2 model

Storage Infrastructure Model SIM

Application Layer

- The Application layer is the final process on the host that is initiating storage communications or receiving storage communications from the storage itself

- Examples:
  - Database
  - File Access (via File System)
  - Storage Control mechanism (fdisk, I/O write utility)

Storage Infrastructure Model SIM

File System Layer:

- Is the applicable file system a host may be using for read/write access to the SCSI layer
- Is not always necessary as some applications write raw blocks to via the SCSI layer
- Is subdividable as in the use of network attached storage file systems and various cluster file systems

Storage Infrastructure Model SIM

Virtualization Layer

- Virtualization Layer represents SCSI access in a modified format to either the host system or to the storage subsystem.
- Examples of Virtualization solutions:
  - RAID levels
  - Data mirroring (RAID 1 or others)
  - Storage partitioning
  - Storage expansion

Storage Infrastructure Model SIM

Virtualization Layer

- Examples of Virtualization formats:
  - array based existing in the storage subsystem itself
  - appliance based existing between the storage and host
  - host based which may reside at a hardware layer or software layer
- Virtualization may occur multiple times within a solution allowing for virtualization on top of virtualization

SIM Topology Examples

Storage Infrastructure Model SIM

Internal storage – database writing raw SCSI blocks

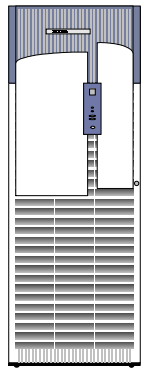

Application (database)
SCSI
NAC (SCSI parallel bus)
Link (ribbon cable)

SCSI
NAC (SCSI parallel bus)
Link (ribbon cable)

Storage

Host

Storage Infrastructure Model SIM

Internal storage – database writes via a file system

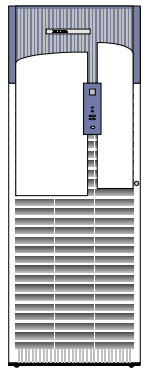

Application (database)
File System
SCSI
NAC (SCSI parallel bus)
Link (ribbon cable)

SCSI
NAC (SCSI parallel bus)
Link (ribbon cable)

Storage

Host

Storage Infrastructure Model SIM

External Array - JBOD

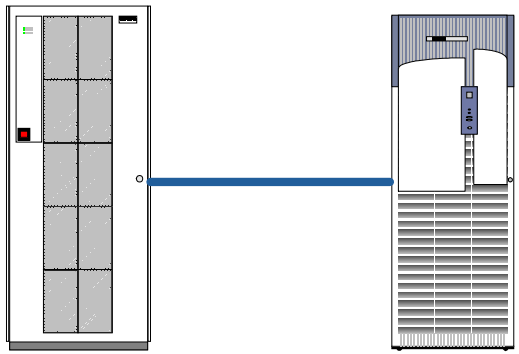

Application (database)
File System
SCSI
NAC (SCSI parallel bus)
Link (parallel cable)

SCSI
NAC (SCSI parallel bus)
Link (parallel cable)

Storage

Host

Storage Infrastructure Model SIM

External Array – JBOD with software RAID on host

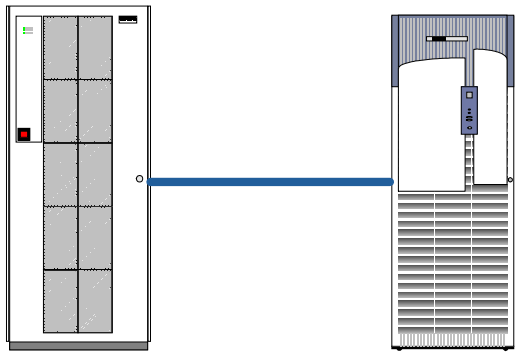

Application (database)
File System
Virtualization (RAID 5)
SCSI
NAC (FC Transport)
Link (MM Fiber)

SCSI
NAC (FC Transport)
Link (MM fiber)

Storage

Host

Storage Infrastructure Model SIM

External Array – JBOD with hardware RAID on host

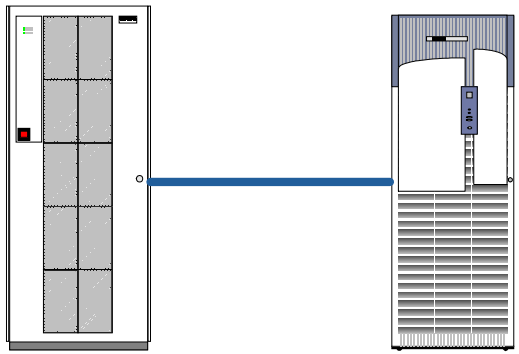

Application (database)
File System
SCSI
Virtualization (RAID 5)
NAC (FC Transport)
Link (MM Fiber)

SCSI
NAC (FC Transport)
Link (MM fiber)

Storage

Host

Storage Infrastructure Model SIM

External Array – Virtualization on Array

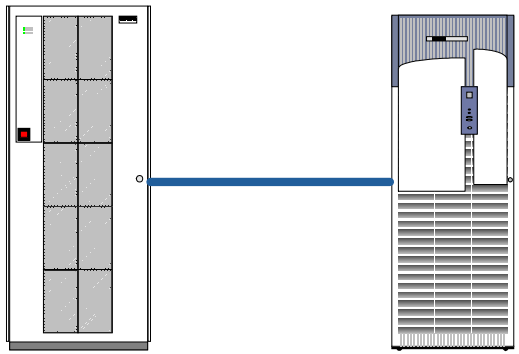

SCSI
Virtualization
NAC (FC Transport)
Link (MM fiber)

Application (database)
File System
SCSI
NAC (FC Transport)
Link (MM Fiber)

Storage

Host

Storage Infrastructure Model SIM

Tape Library Application

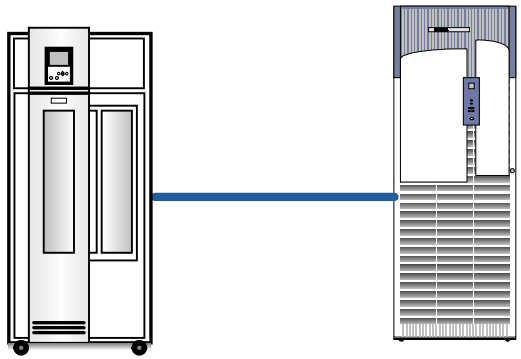

SCSI
NAC (FC Transport)
Link (MM fiber)

Application (backup)
SCSI
NAC (FC Transport)
Link (MM Fiber)

Tape

Host

Storage Infrastructure Model SIM

NAS Appliance

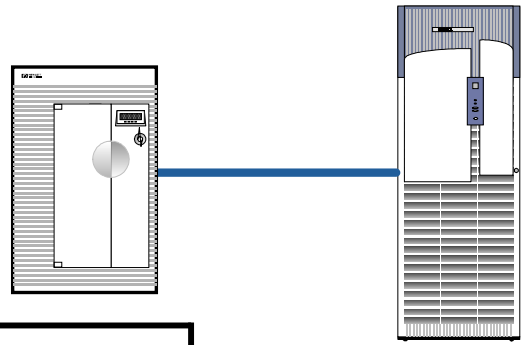

SCSI	
Virtualization	
File System	Local FS
	NFS or CIFS
NAC (TCP/IP over Ethernet)	
Link (Cat 5)	

Application
File System (NFS or CIFS)
NAC (TCP/IP over Ethernet)
Link (Cat 5)

NAS Appliance

Host

Storage Infrastructure Model SIM

External Array with Virtualization Appliance

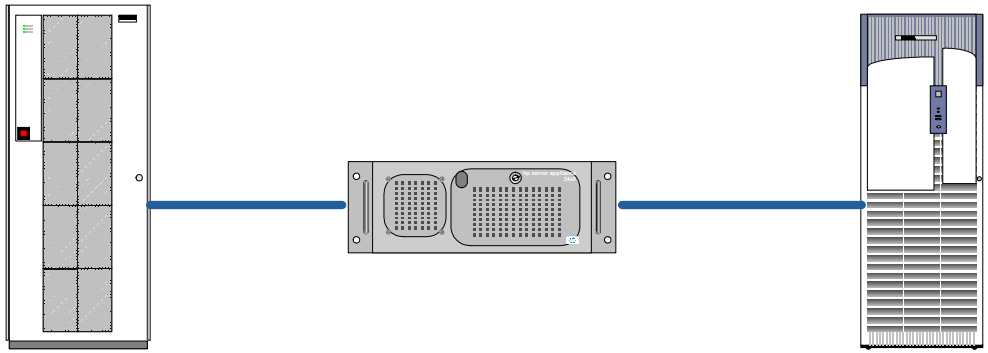

- Virtualization exists outside of host or storage
- Same concept applies whether an appliance is used or virtualization is embedded in the fabric itself

Storage Infrastructure Model SIM

External Array with Virtualization Appliance

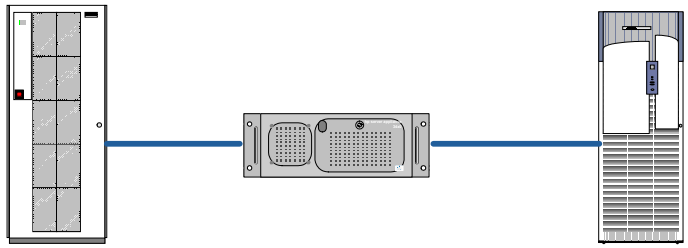

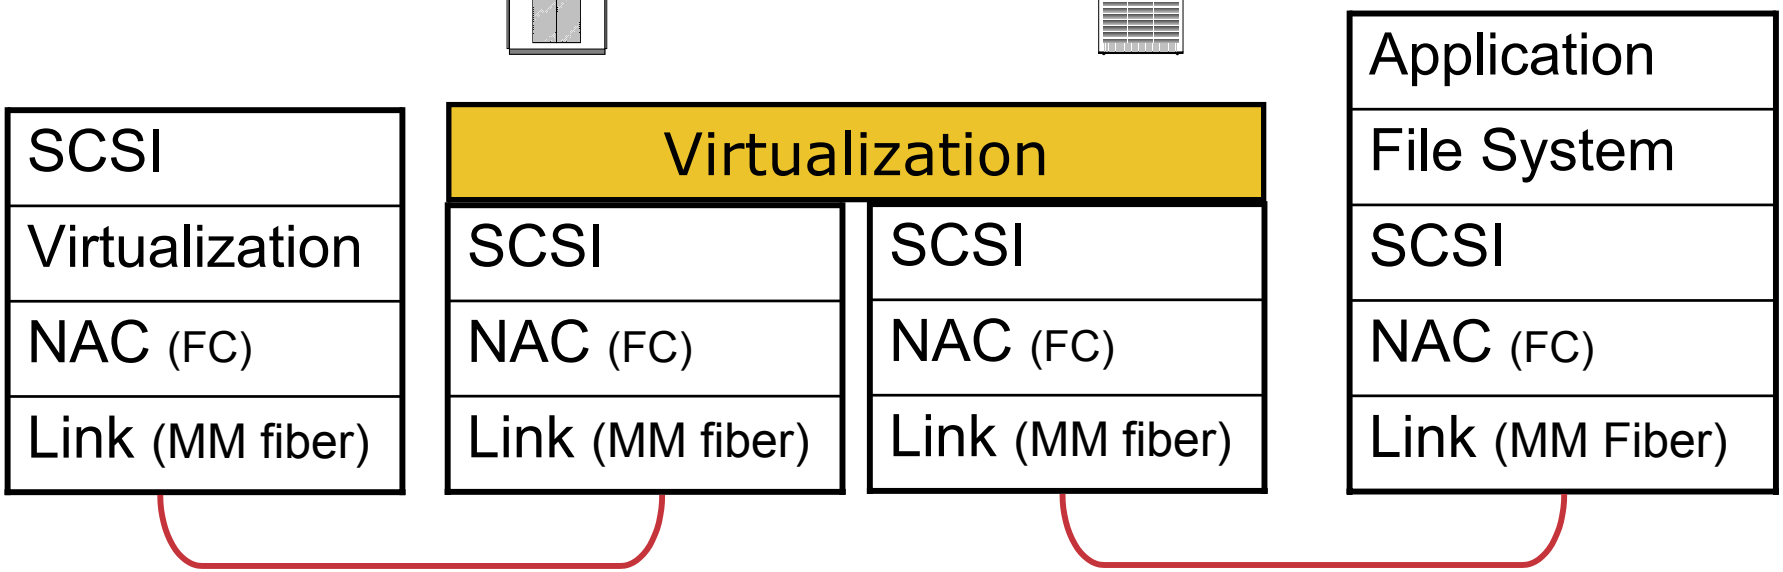

Storage Infrastructure Model SIM

FCIP

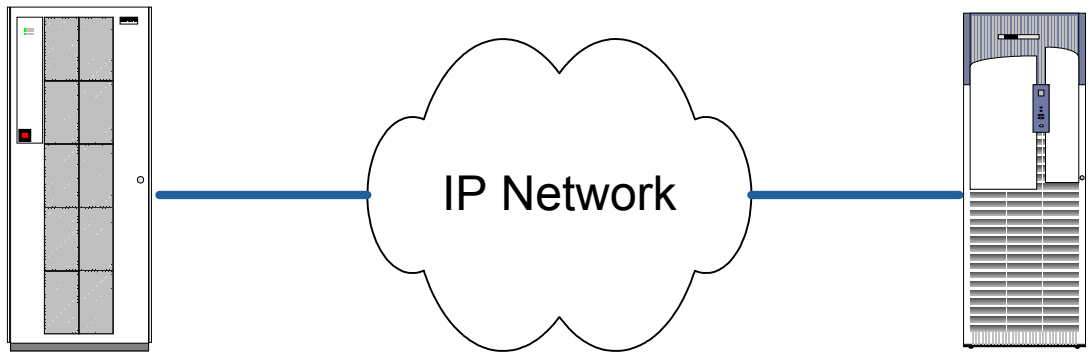

- Model must be inserted within the model for tunneling applications
- NAC and SCSI layers are the layers transferred in FCIP

Storage Infrastructure Model SIM

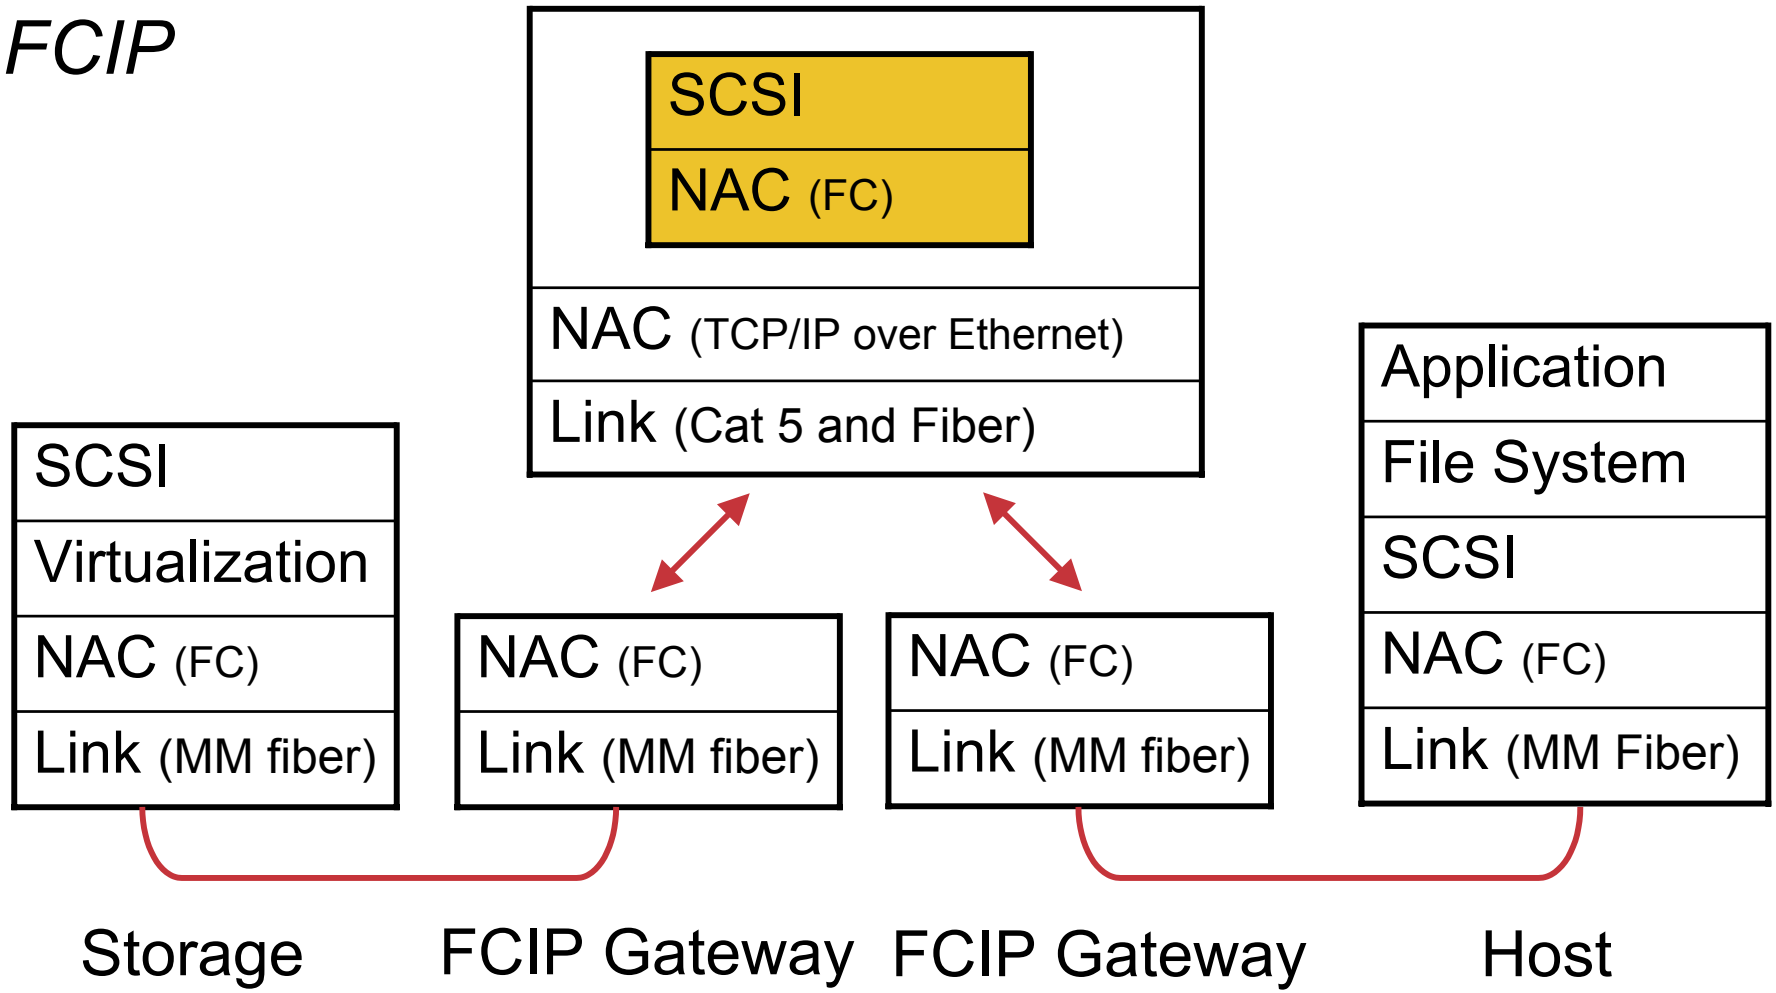

SIM Advantages

SIM and Administration

- A Layered model approach such as SIM should be used not only for theoretical conceptualization but also for standard administrative purposes
- Documentation – Diagram each server and storage configuration labeling the following:
  - FC WWNs
  - IP addresses
  - RAID groups
  - LUN mappings and security
  - Mirroring or advanced copying applications

SIM and Administration

- Once the existing configuration is documented according to the SIM model, future growth and expansion can be well planned by adjusting the model
- Areas of consolidation can be easily identified and planned by adjusting the model
- Knowledge transfer to new administrators or between systems administration and network administration is much easier and effective

Fault isolation

- In traditional networking, the fastest method of fault isolation is process of elimination according to the OSI model
- Specific network configurations can develop specific ‘if this then…’ methods to troubleshoot and isolate problems in the network
- SANs can use the SIM model in much the same way to isolate problems according to each layer

Fault Isolation Scenarios

Fault Isolation Scenario

Link Layer isolation

- Customer Configuration:
  - 30 Brocade 2800 switches in entire fabric
  - 24 W2K Data Center hosts clustered in groups of 4
  - DWDM solution connecting the two islands across a river
- Problem – Once a week, the W2K hosts were losing access to storage. Loss of access was not sudden; it was a slow process losing one LUN at a time over a 45 minute period.

Fault Isolation Scenario

Link Layer isolation

- Vindicated the Application layer as some LUNs were internal storage and were not affected. This was a SAN issue
- The logs did show SCSI reserve/release errors possibly indicating SCSI layer conflicts
- We were able to conclude the large number of SCSI reserve/release errors were due to the large number of PLOGIs in the fabric. The NAC layer was causing issues in the SCSI layer

Fault Isolation Scenario

Link Layer isolation

- By setting up a test lab and running a fiber analyzer, we noticed several things:
  - The occurrences were coinciding with the customer's weekly fabric reconfiguration
  - The fiber analyzer showed a timing conflict between PLOGIs on the HBAs and the SCSI reserve/release messages
  - The other hosts on the fabric (UNIX hosts) were not having issues
- The fabric registration process was timing out the W2K hosts and causing SCSI reserve/release errors

Fault Isolation Scenario

Link Layer isolation

- We concluded that the latency on the DWDM solution was too great for the threshold on the HBAs
- We increased the timeout thresholds on the HBAs to allow for the latency
- We also recommended the customer increase the time interval between the reconfiguration of Fabric A and Fabric B so as to not add additional disturbance to the configuration

Fault Isolation Scenario

NAC Layer isolation

- Customer Configuration:
  - 28 Brocade 2800 switches
  - 90 Sun Solaris hosts
  - 88 W2K hosts (mixed vendors)
  - 88 HP-UX 11i hosts
  - 24 XP 512 arrays
- Problem – twice in a three week period the entire SAN came down and the SUN hosts were being flooded with online/offline/failed errors

Fault Isolation Scenario

NAC Layer isolation

- Vindicated the Link layer by analyzing switch logs and looking for physical link errors – none were found.Fiber plant was shot with fiber analyzer for physical errors – none were found
- Vindicated the Virtualization layer (on the XP arrays) as no errors were noticed in the logs. Test hosts were isolated to direct attached to the XPs as a test group for the SAN itself
- The fiber analyzer did however point to two unusual situations:
  - a high amount of RSCNs were constantly generated
  - every time a RSCN was issued the SUN host would issue PLOGIs

Fault Isolation Scenario

NAC Layer isolation

- Customer informed us the RSCNs came from a test bed of servers on the SAN that were constantly being attached and detached from the SAN
- It was noticed the HBAs on the SUN hosts were sending PLOGIs every time either a RSCN 1 or RSCN 3 was being sent:
  - RSCN 1 is a non-global change in the fabric
  - RSCN 3 is a global change in the fabric requiring PLOGIs
- The solution was a modification in the HBA firmware to eliminate PLOGIs for RSCN 1 messages

Fault Isolation Scenario

Application Layer isolation

- Customer configuration:
  - 20 + W2K servers (Dell and HP)
  - 4 Brocade 2800 switches
  - mixed arrays of HP VA and EMC Clariion
  - virtualization appliance between hosts and storage

- Customer noticed a 50% or more SAN performance degradation on close to 1/3 of the servers

Fault Isolation Scenario

Application Layer isolation

- Vindicated the Link layer by looking for link errors on switches and ensuring devices were successfully registering with switches – also did physical inspections of cables and GBICs

- Vindicated the NAC layer by running ‘portperfshow’ on switches. Hosts with normal performance showed identical performance rates on switches. HBA performance and error status and failover software was also checked with no positive indications

Fault Isolation Scenario

Application Layer isolation

- Vindicated the Virtualization appliance by remapping good performance LUNs to bad performance hosts. The good LUNs were suddenly showing poor performance
- Vindicated the SCSI layer by looking in the application and event logs with no errors shown
- We discovered the slow performing hosts all had a legacy monitoring software loaded. By disabling the service for the monitoring software – the performance for the hosts in question dramatically improved

Summary

Summary and Conclusions

- Traditional Networking has long used layered models for better conceptualization, administration, and fault isolation of networks.
- Storage Networking, although evolving from a different concept than traditional networking, can gain significant advantages by employing the layered model approach
- The SIM model is an excellent model used to capture all aspects of a Storage Area Network
- The SIM model is a flexible model that has aided in the conceptualization, administration, and fault isolation of existing SANs

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