



Storage Networking for Network Professionals

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

Storage Networking for Network Professionals



- Objectives:
 - Identify what networking is and how it relates to storage
 - Examine storage *networking* technology similarities and differences
 - Clarify and explain the language of networking: Terminology
- Topics
 - Network Vs. Storage
 - Disk & SCSI
 - Storage and IT
 - DAS, NAS, SAN
 - Fundamental Network Concepts
 - OSI, IPS (TCP/IP) and Fibre Channel
 - Traditional FC SAN, Traditional NAS, iSCSI SAN, Full Offload
 - Semantic "ANALyst"
 - Futures



Network Vs. Storage



Network Considerations	Storage Considerations
Information movement	Information repository
Data over Distance	Data over Time
Configuring and segmenting network topologies	Formatting and partitioning hard disks
Hubs, bridges, switches, routers, gateways	JBOD, RAID, Tape Controllers
Client/Server Applications	Initiator / Target Functions
Email, Web Browsing, File Mgt Network Interface	Backup, Archive, Mirror, Block Mgt, I/O Interfaces



Network Vs. Storage



Network Considerations	Storage Considerations
Speed, Solid-State	Capacity, Mechanical Movement
Media – Copper, Fiber, Wireless	Media – Disk, Tape, Paper
Bits in Frames, CRC error checks every frame	Bytes in Blocks, Parity check every byte
Buffers and Link transfer rates – Kbps, Mbps, Gbps	Cache and Disk/Tape I/O interface transfer rates - MBps
Interconnectivity network and device latencies - bottlenecks	Disk/Tape seek times and latencies – ms, us, ns
Applications – HTTP, NFS, SMTP Transport Protocols – (TCP/IP/Ethernet) , Fibre Channel	Logical Device Protocols – SCSI, ESCON, IDE/ATA Physical Protocols – Parallel SCSI



Data Units of Measure – Really Big Numbers



Unit of Measure	Acronym	Size
Byte	Byte	8 bits
Block	Blk	512 Bytes
Kilobyte	KB	1,024 Bytes
Megabyte	MB	1 Million Bytes
Gigabyte	GB	1,000 MB
Terabyte	TB	1,000 GB
Petabyte	PB	1,000 TB
Exabyte	EB	1 Million TB
Zetabyte	ZB	1,000 EB
Yottabyte	YB	1,000 ZB

How long would it take someone to download a one-yottabyte file using a 28.8 baud modem?



11 trillion years.



Data Units of Measure – Really Small Numbers




Unit of Measure	Power of 10	Number
milli	3	One thousandth
micro	6	One millionth
nano	9	One billionth
pico	12	One trillionth
femto	15	One quadrillionth
atto	18	One quintillionth
zepto	21	One sextillionth
yocto	24	One septillionth

1



Disk Technology

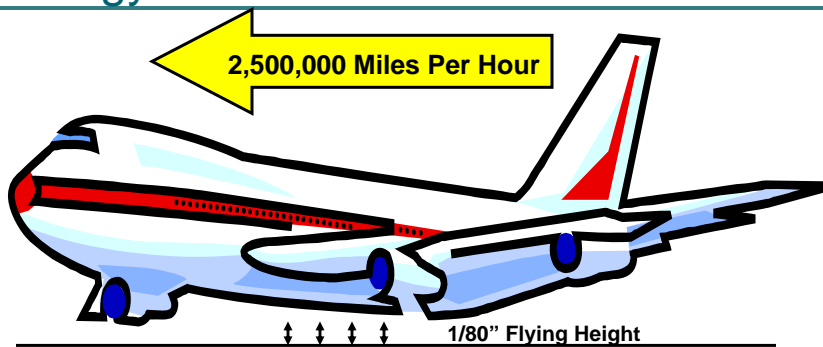


- Performance may depend on
- Revolutions per minute, Seek time,
 - Latency. Data transfer time
 - Caching mechanisms

Platter
Track
Sector
Head
Actuator Arm
Cylinder




Disk Head Flying Height Analogy




2005 head flying height – 1/2 micro-inch (“partial contact”)



Disk Logical Unit



Disk Logical Unit




Set of blocks accessible
by logical address
based upon physical
information



RAID Striping



- Redundant Array of Independent (Inexpensive?) Disks
- From large “monolithic” sequential disks
- Better performance though distributing I/O across multiple smaller disks in parallel



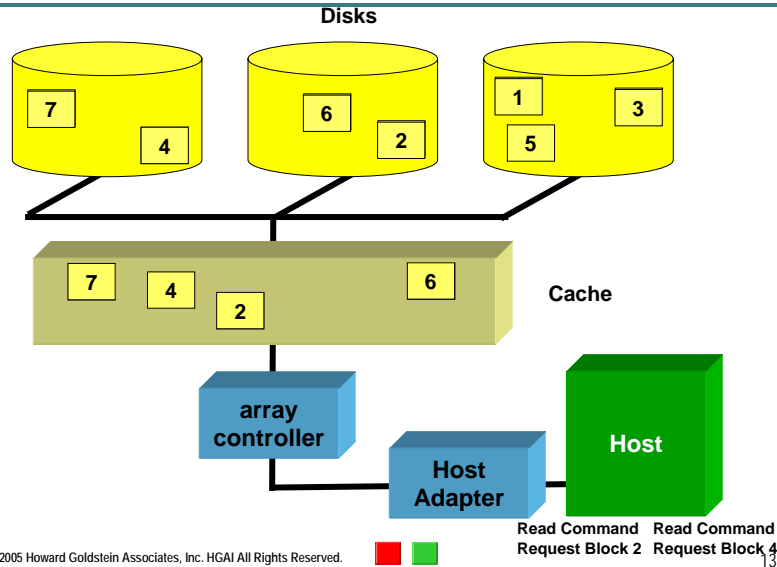
Cache Types




- Read Caching
 - Recently Used Cache: Data hot spots
 - Read-Ahead Cache: Sequential pre-fetch
- Write Caching
 - Write-Through Cache: Direct to non-volatile memory. Optimized for integrity
 - Write-Back (Write-Behind) Cache: Faster response to applications. Optimized for performance



Read Cache – Improves I/O Throughput



Write Cache – Improves I/O Throughput



Cache mechanics

- Memory (volatile and non-volatile)
- Index for locating cached data
- Data loader
- Activity indicator / discard processor
- “Dirty” black indicator (for write-back caches)
- Write-back processor
- Caching locations
 - Host system software and memory
 - Device or subsystem caching controller
 - Host I/O caching controller



Tagged Command Queuing

Incoming

I/O seq	Block
10	8
9	2
8	3
7	6
6	2
5	1
4	7
3	4
2	8
1	1

Re-Ordered

I/O seq	Block
10	8
2	8
4	7
7	6
3	4
8	3
9	2
6	2
5	1
1	1

	SATA-II	SCSI
Command Queuing	32	256


Device re-ordering to minimize seek time



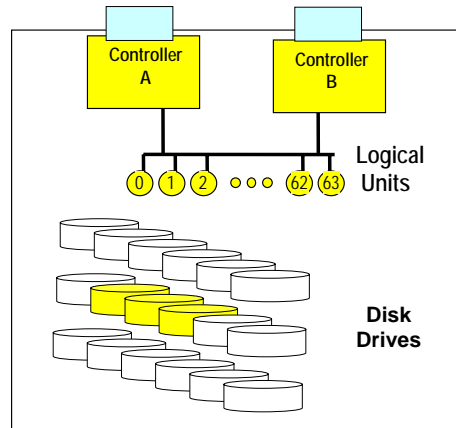
JBOD & RAID



Just a Bunch of Disks



Redundant Array of Independent Disks



Disks in the Data Center




Attribute	Desktop Disk	High-End Disk
Interconnect	ATA to SATA	SCSI, FC_AL, SAS
Retail Cost/GB	\$0.45 - \$0.75	\$2.50+
Performance	100 IO per Second	180+ IO per Second
Latency	12 msec	6-8 msec
Capacity	Up to 45GB	Up to 300GB
Reliability (MTBF)	500K Hours	12M Hours



RAID



- RAID 0 – Non Redundant Individual Access Striping (Performance)
- RAID 1 – Mirrored Volumes (Redundancy)
- RAID 0 + 1 – Striping Mirrored Array (Performance)
- RAID 3 – Parallel Access Array
- RAID 4 – Block Level Striping with Parity Disk
- RAID 5 – Striping with Distributed Parity
- RAID 1 + 0 – Mirrored Striping Array




Redundant Array of Independent Disks - RAID



- Disks are grouped together to emulate larger disks
- RAID delivers capacity and management benefits
- RAID delivers performance benefits
- RAID delivers reliability and availability benefits



RAID 0 - Data Striping



Minimum of two disks required

High I/O rate


No data protection (parity)

Optimized for fast access to data

Used for any application requiring high bandwidth (video production and editing, image editing)



RAID 1 - Disk Mirroring



Minimum of two disks required - "Mirror"


Safe but expensive

Highest overhead of all RAID types

Used for any application requiring high availability (accounting, payroll, financial)



RAID 1/0 - Disk Mirroring




Minimum of four drives required

High I/O performance with high data reliability

Used with a database server requiring high performance and fault tolerance



RAID 3 Data Striping with Single Parity Disk




Minimum of three disks required

Very high Read and Write data transfer rate

Best application for fast transfers of large sequential blocks (live video streaming)



RAID 5 - Data Striping & Parity on All Disks



Minimum of three disks required

Better Read performance than Write

Most versatile RAID level


Used with multiple applications (WWW, intranet, file, application, database, and news servers)



SCSI Anatomy



- SCSI devices can be Initiators or Targets
- SCSI IDs and Logical Unit Numbers (LUNs) identify individual devices and partitioned devices within devices
- SCSI Nexus - ITLQ



Basic SCSI Components ITLQ



- Host adapter:
 - A device that connects the host bus to the SCSI bus.
- Initiator:
 - Mode of operation within a device that determines what task needs to be executed and which target will perform the task.
- Target:
 - mode of operation within a device that controls the entire process after being successfully selected by an initiator.
- Logical Unit:
 - A physical or virtual peripheral device addressable through a target or peripheral controller.
- Q:
 - Used with Command Queuing to identify command



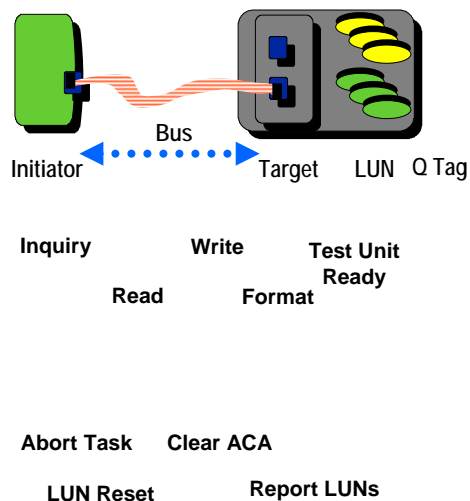
SCSI Operations: I/O & Task Management



I/O Operations

- Commands
 - Descriptor Block CDB
 - Operations Code
 - Length, Linking
 - Q - Task ID, Task Attribute
- Blocks
 - Data
- Status
 - Auto Sense

Task Management



SCSI ID: Parallel Bus Adapter Port Priority



SCSI ID	Priority*		SCSI ID	Priority*
7	1		15	9
6	2		14	10
5	3		13	11
4	4		12	12
3	5		11	13
2	6		10	14
1	7		9	15
0	8		8	16

* Where 1 is highest and 16 is lowest



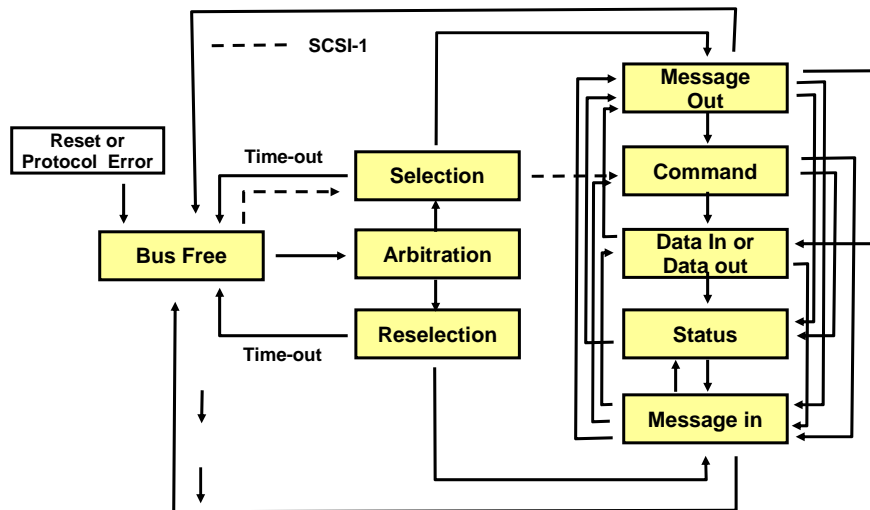
SCSI Bus Control Signals



Signal	Status	Description
BSY	Busy	Bus is in use
SEL	Select	Initiator or target is being selected/reselected
C/D	Control/Data	Data or control info is being sent on the bus
I/O	Input/Output	Sent to indicate data direction
MSG	Message	Sent to indicate in message phase
REQ	Request	Requests data transfer
ACK	Acknowledge	Acknowledges data transfer
ATN	Attention	Indicates message for initiator
RST	Reset	Hard reset clears all devices
DB	Data Bus	Indicates data bits and parity



Bus Phase Sequence Diagram



SCSI Commands



- **Mandatory Commands**
 - Test Unit Ready
 - Request Sense
 - Inquiry
 - Send Diagnostic
- **Optional Commands**
 - Receive Diagnostic Results
 - Log Select
 - Log Sense
 - Read Buffer
 - Write Buffer
 - Copy
 - Copy and Verify
 - Compare
 - Report Luns
- **Direct Access Commands**
 - Format Unit
 - Lock – Unlock Cache
 - Mode Select
 - Mode Sense
 - Power Conditions
 - Read Capacity
 - Reserve / Release
 - Seek
- **Sequential Access Commands**
 - Load / Unload
 - Prevent / Allow Media Removal
 - Rewind
 - Write File Marks



SCSI Command Queuing



- Definition – Queuing of I/O processes: Allows a target to accept multiple I/O processes for execution at a later time
- Two methods for implementation of queuing
 - Untagged
 - **Allows a target to accept one I/O process:**
 - From each initiator
 - From each logical unit
 - Tagged
 - **Allows a target to accept multiple I/O processes:**
 - From each initiator
 - From each logical unit
 - This is done by establishing an I_T_L_Q nexus for each command



SCSI Command Queue Types



- Simple Queue Tag
- Ordered Queue Tag
- Head of Queue Tag



IT Customer Environment



- Survey of 211 U.S. sites running mission critical database applications - Pain Points
 1. Managing Disk Space (40%)
 2. Backup (38%)
 3. Running out of Disk Space (31%)
- Customer Messages
 - Budgets are flat, spending is tactical
 - Only top pain points get funding
 - “Reduce my spending or pain this year, or Go Away”
 - Reluctance to buy from startups (dot-com hangover effect)



Storage Provisioning – The Problem



- Physical disks have fixed capacity and performance
 - So do the RAID sets built from physical disks
- Application capacity and performance requirements grow
 - Steadily, Independently, Exponentially
- Penalties for not meeting application needs are severe
 - Run out of capacity – risk application crashing
 - Run out of performance – risk alienating users



Storage Provisioning – The Problem



- Provisioning is Complex and Painful
 - For many customers, adding storage to an application takes a week
 - Only trained, top-tier system administrators can provision storage
- Lots of ongoing industry effort to simplifying provisioning
 - Holy Grail: make storage provisioning as simple as memory provisioning in servers – “If you don’t have enough, add more”



Steps in Storage Provisioning



1. Find an application that’s running out of space or performance
2. Determine how to add storage (extend a LUN, make a new LUN)
3. Determine how much storage to add
4. Determine where in pool to obtain additional storage from (taking performance & availability issues into account)
5. Allocate storage from pool
6. Set protection (LUN masking) of storage server so application can see it
7. Set switch zoning if necessary so application can see storage server
8. Add new storage to application LUN or form new LUN
9. Configure O/S file system to handle bigger/new LUN
10. Configure application to utilize bigger/new LUN
11. Configure Backup system to back up bigger/new LUN
12. Replenish pool, if necessary (buy disks, build new RAID sets)

How many GUI’s can you (do you want to) learn?

What happens if (when) you make a mistake?

Do you dare do this while the application is running?



What is Storage Virtualization?



- An **abstraction of storage** that separates
 - Host view
 - Storage system implementation
- Makes invisible to host:
 - physical pathing
 - device characteristics
 - physical data location
- Provides Location and Implementation Transparency
- Dynamic
 - Enables transparent “on the fly” reconfiguration
 - Allow data location to change transparently to host environment
- *There are a lot of different types and approaches and degrees of storage virtualization*



What do Users Want from Storage?



Application aspects of storage

- **Capacity**
 - **Application requirements**
 - **Growth potential**
- **Performance**
 - **Throughput**
 - **Responsiveness**
- **Availability**
 - **Failure resistance**
 - **Recovery time**

Physical aspects of storage

- Capacity
 - Disk or tape size
 - Number of disks / string
 - Number of tape devices
- Performance
 - Disk latency & seek time
 - Cache size & hit rate
 - Media data rate
- Availability
 - MTBF
 - Path redundancy



What is Storage Virtualization?



- **Physical disks**
 - Fixed size
 - Bounded performance
 - Do break (occasionally)

**Block-level
Virtualization**

- **Virtual disks**
 - As big, small or as or as many as users need
 - As fast as users need
 - Can be “very reliable”, or not
 - Can grow(!), shrink or morph



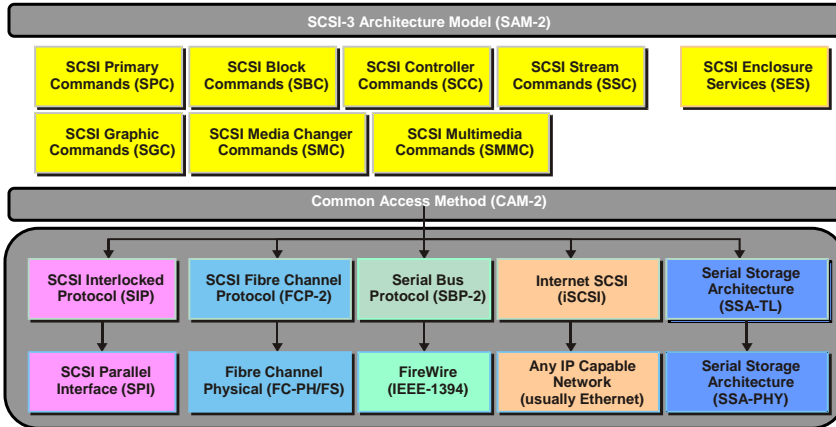
Storage Provisioning Revisited



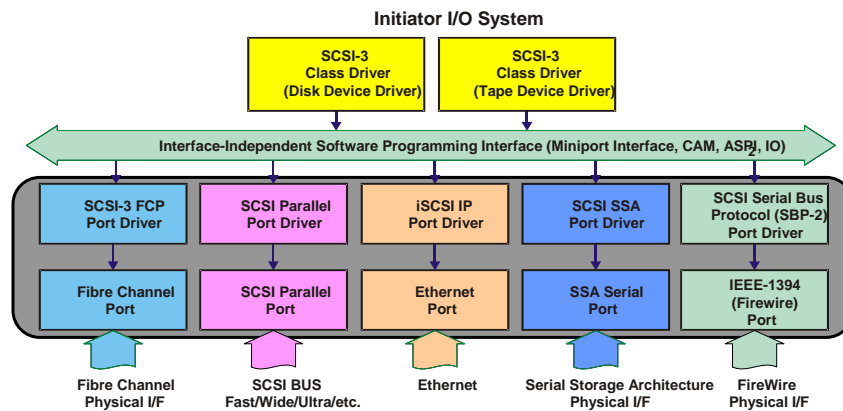
- Virtualization doesn't solve the provisioning problem
 - “Just another GUI to deal with”
- Virtualization enables automatic provisioning because of its flexible physical-to-logical mapping
 - All you need is some magic software that can speak to all those GUI's/APIs, understand application needs, and plan storage usage
 - Good luck finding it!
- Some partial, proprietary provisioning solutions exist
- Open solutions inhibited by lack of interoperability
 - Everyone's API is different and secret
 - API swaps don't scale
- What can be done? See SMI-S



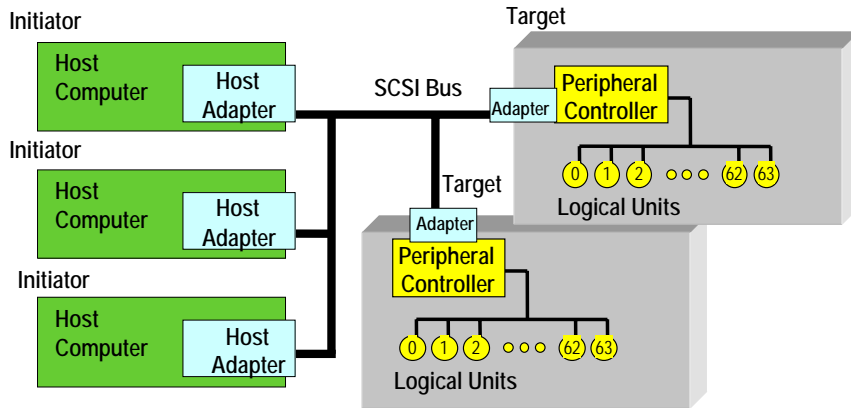
SCSI-3 Layered Structure



SCSI-3 Physical Interfaces



SCSI



Narrow Bus - 8 Device Adapter SCSI IDs
 Wide Bus - 16 Device Adapter SCSI IDs
 Half Duplex



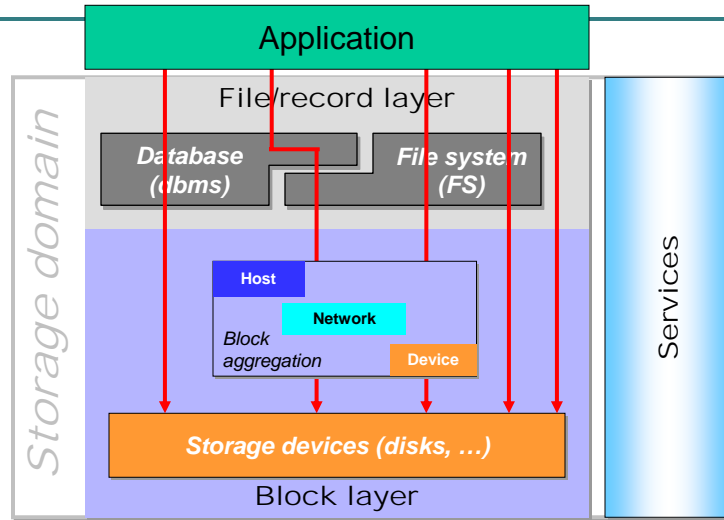
Protocol Optimization: SCSI Bus Speeds, Width, Distance



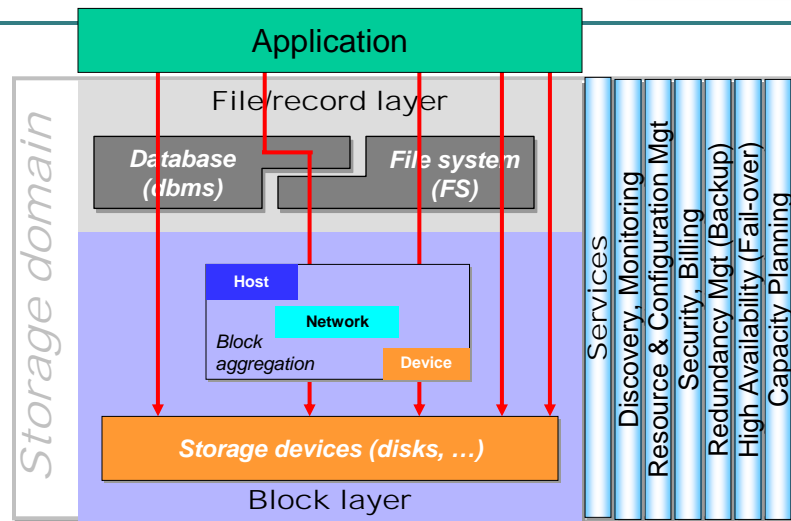
Single Ended	Fast-10	SCSI 2	Narrow 10MB/s	3 meters
			Wide 20MB/s	
	Fast-20	SCSI 3	Narrow 20MB/s	1.5 - 3 meters
			Wide 40MB/s	
High Voltage Differential	Fast-10	SCSI 2	Narrow 10MB/s	25 meters
			Wide 20MB/s	
Ultra SCSI	Fast-20	SCSI 3	Narrow 20MB/s	
			Wide 40MB/s	
Low Voltage Differential	Fast-40	SCSI 3	Narrow 40MB/s	12 - 25 meters
			Wide 80MB/s	
	Ultra 2	SPI-2	Narrow 80MB/s	
			Wide 160MB/s	



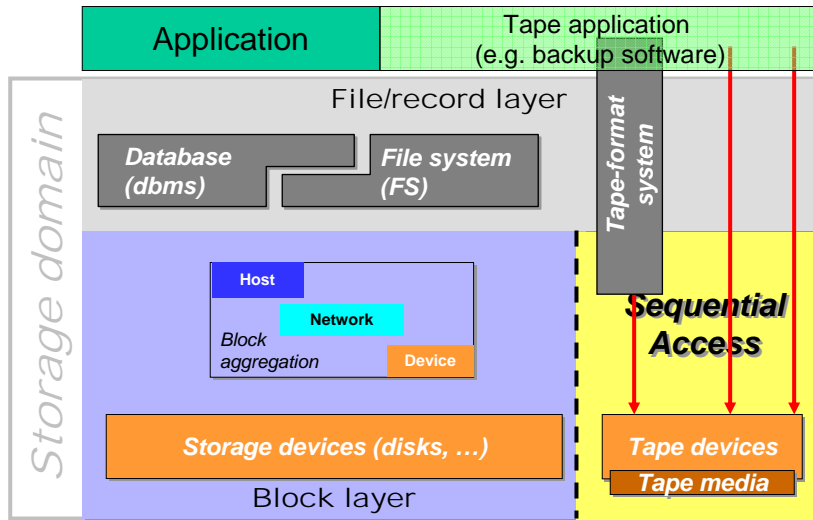
SNIA Shared Storage Model



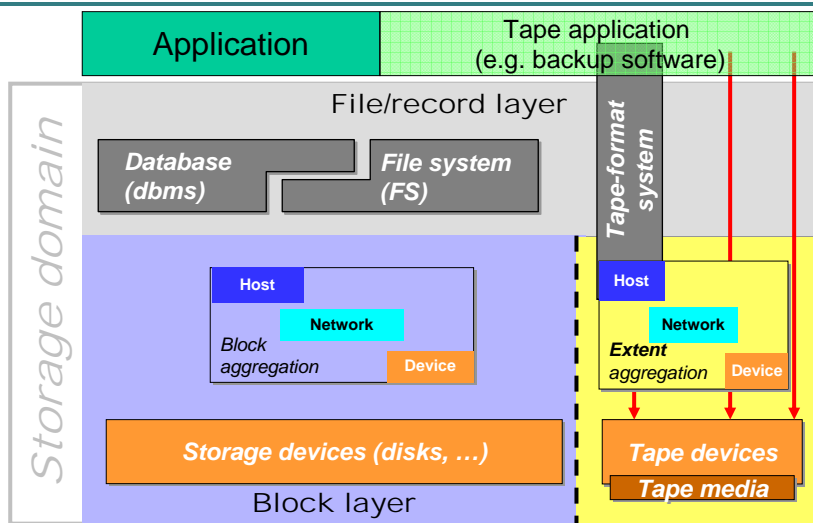
SNIA Shared Storage Model



SNIA Shared Storage Model



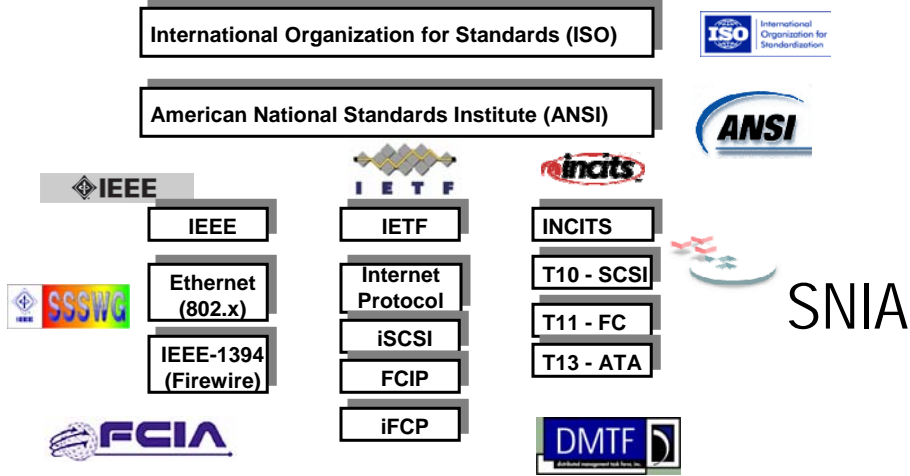
SNIA Shared Storage Model



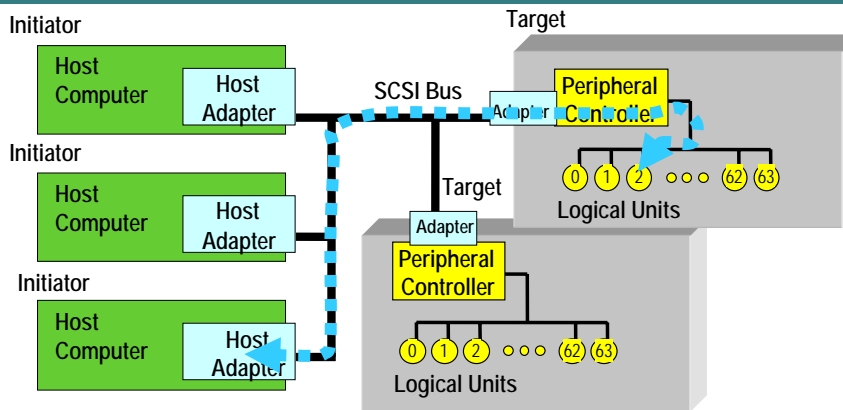
Standards



- Which organizations define which standards?

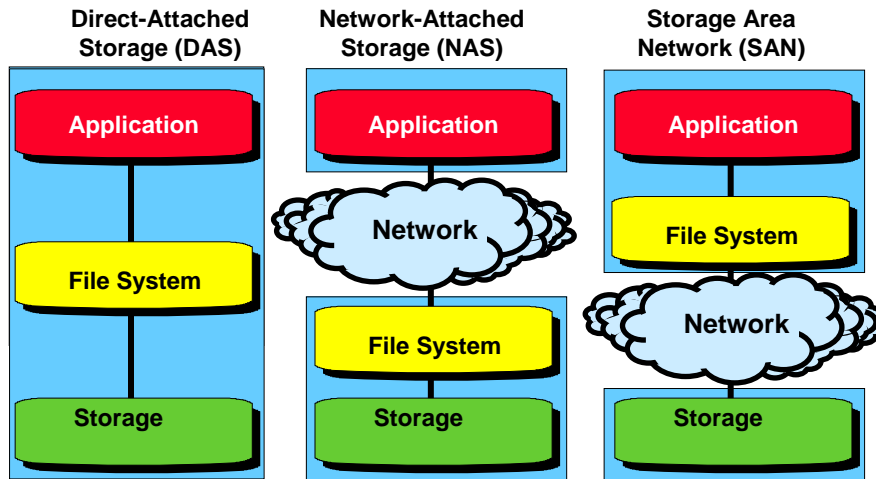


Virtual SCSI Cables on SCSI Bus

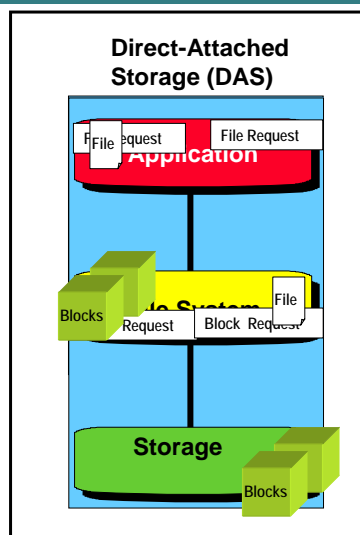


SCSI Bus: Arbitrate ID, Select ID, Identify LUN

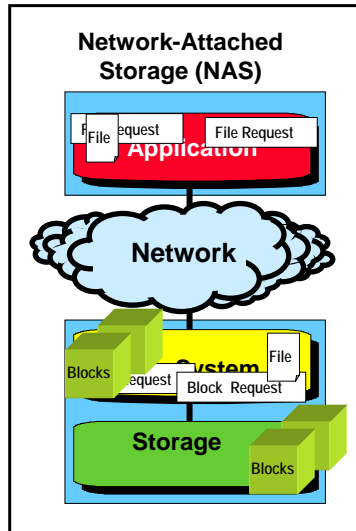
Basic Storage Network Technology



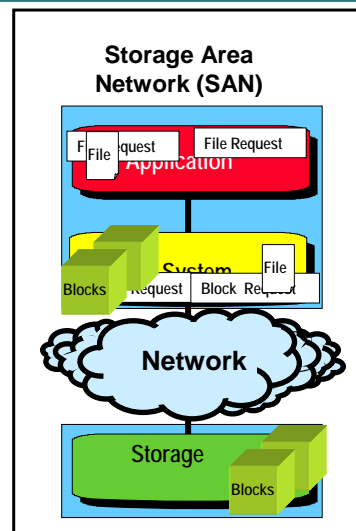
DAS



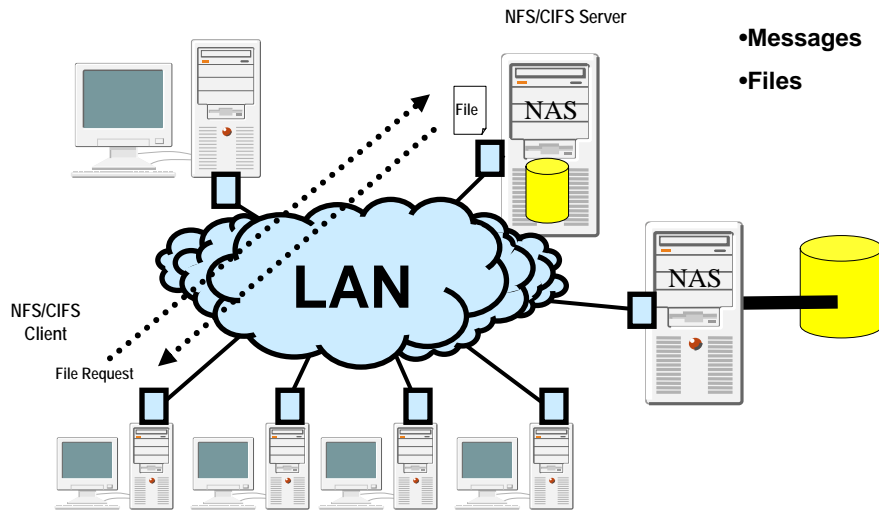
NAS



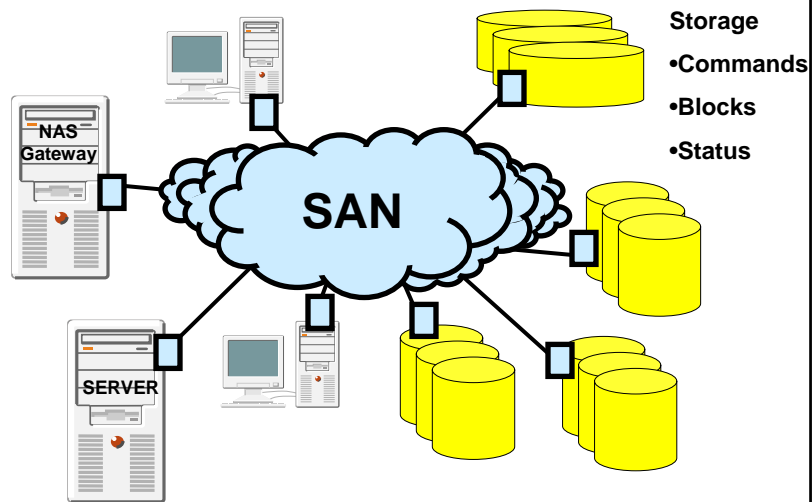
SAN




Network Attached Storage



Storage Area Network



Storage Area Network – “External SAN”



© Copyright 2005 Howard Goldstein Associates, Inc. HGAI All Rights Reserved.




59

LAN/SAN Benefits



- Greater Asset Utilization
- Expanded connectivity
- Port consolidation
- Extended distance
- Higher bandwidth
- Potential performance improvement

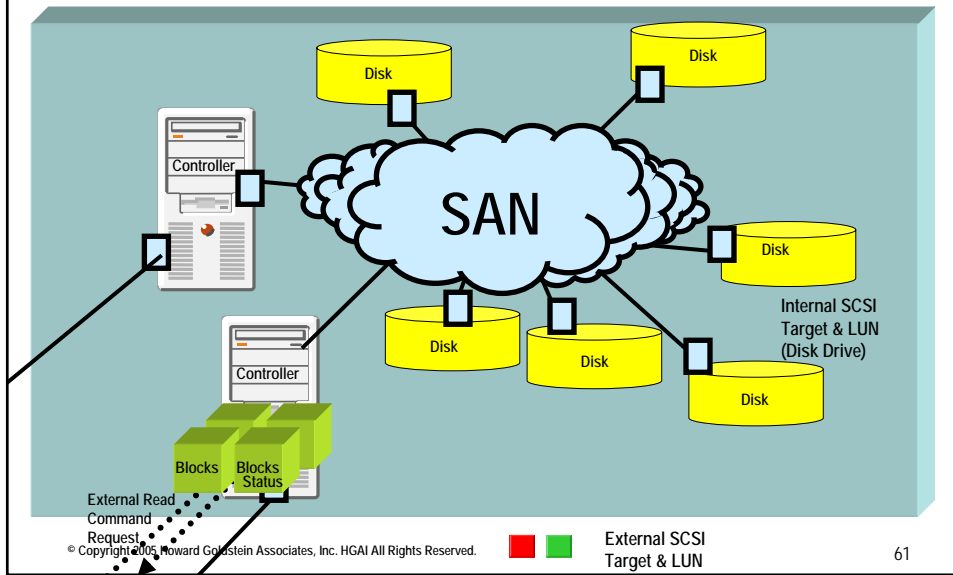


© Copyright 2005 Howard Goldstein Associates, Inc. HGAI All Rights Reserved.

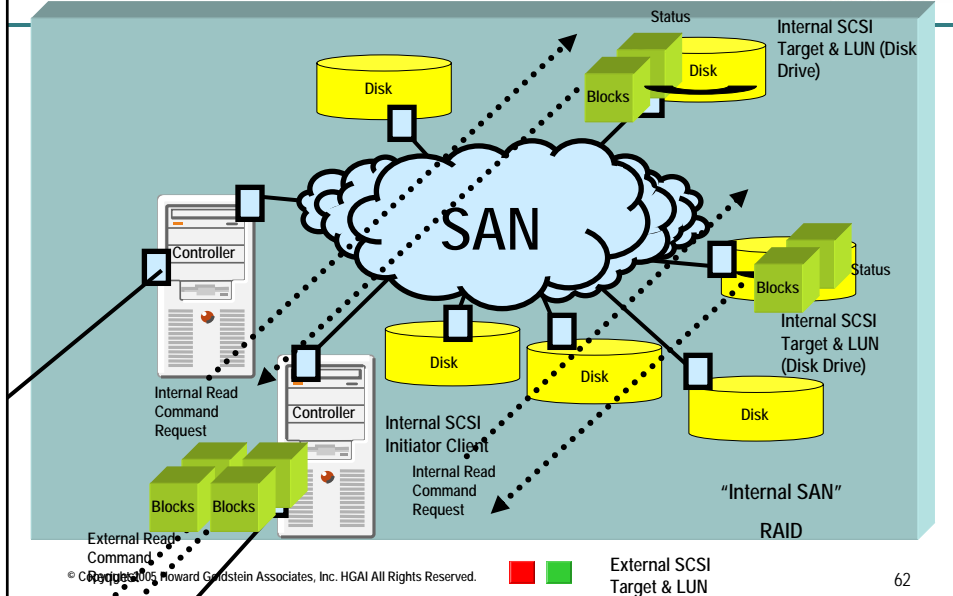


60

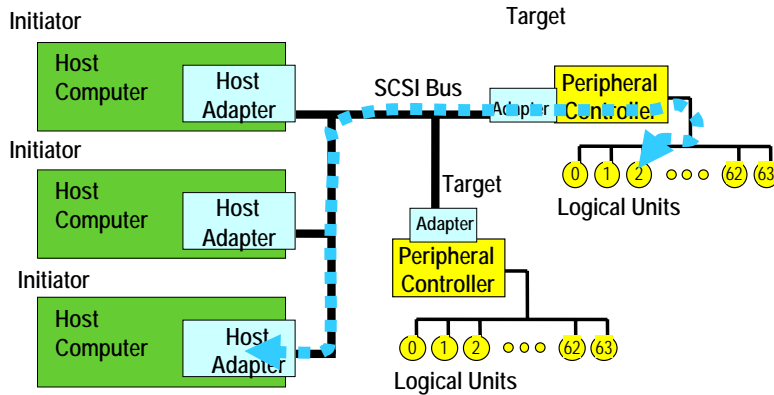
“External SAN” I/O Request to Storage Controller



“External SAN” I/O Request to Storage Controller



Virtual SCSI Cables on SCSI Bus



SCSI Bus: Arbitrate ID, Select ID, Identify LUN



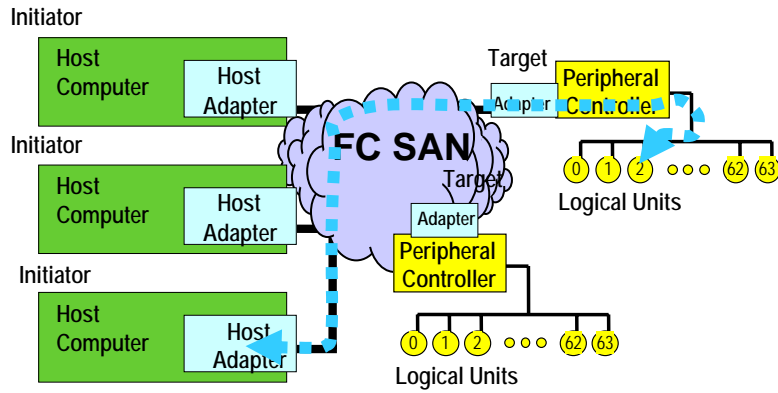
What's the difference between Fibre Channel and the English Channel?



Only one facilitates communication between 2 ports!



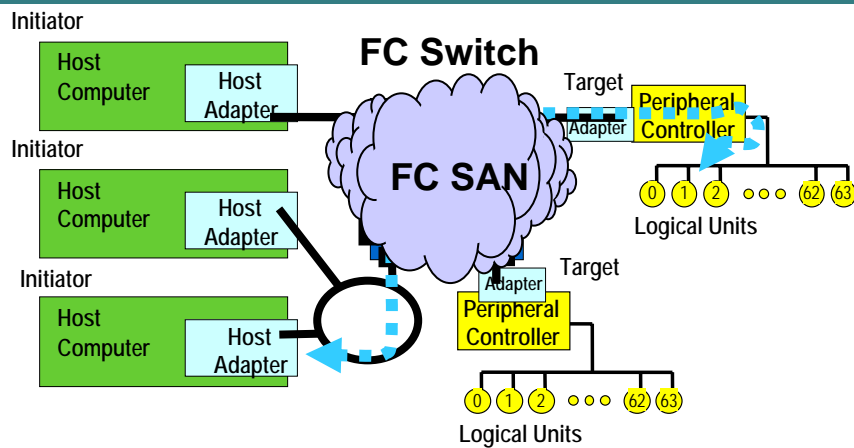
Virtual SCSI Cables Fibre Channel



FC: Fabric Login, N_Port Login, FC-4 Login



Virtual SCSI Cables Fibre Channel



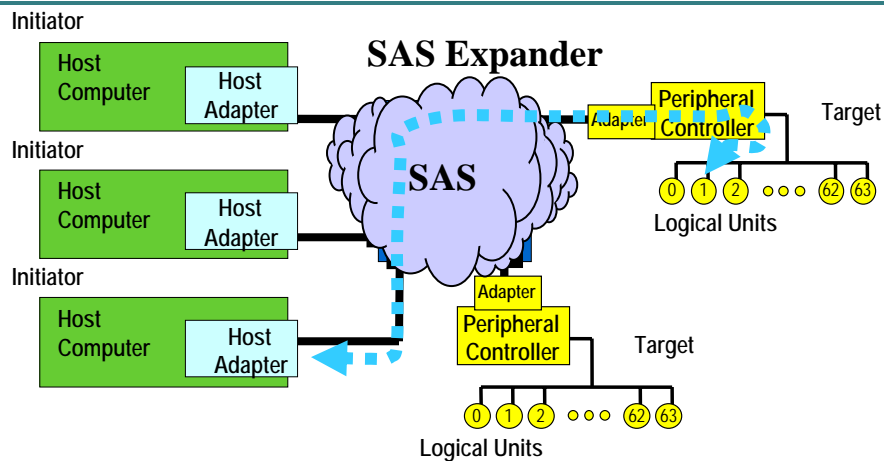
FC Loop: Arbitrate,
Open Loop Circuit

FC Fabric: Switch
Routing

FC Point-to-Point:
Direct Routing



Virtual SCSI Cables SAS



Serial Attached SCSI Bus: Open, Information, Close



Fundamental Network Concepts



- Circuit Switch vs. Packet Switch
- Connection-Oriented vs. Connectionless
- Layered Protocols
- Flow Control
- Payload Integrity
- Names vs. Addresses
- Segmentation, Convergence & Reassembly
- Hardware Offload

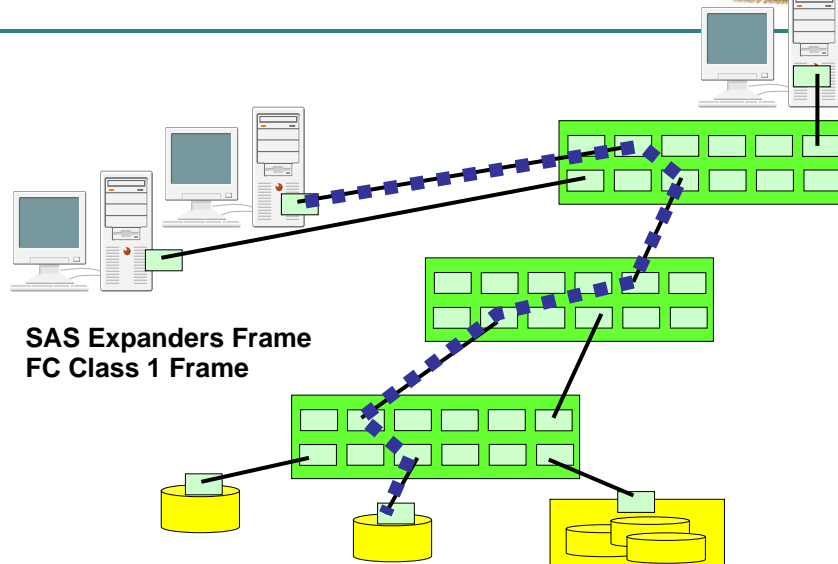


Circuit Switch – Packet Switch

- Fixed path whether moving or not versus route packets on the current best path as they arrive




Circuit Switch



Packet (Frame) Switch

HGAI

Howard Goldstein Associates, Inc.



© Copyright 2005 Howard Goldstein Associates, Inc. HGAI All Rights Reserved.




71

Connection-Oriented vs. Connectionless

HGAI

Howard Goldstein Associates, Inc.

- Not path related!
- Connection-oriented requires a state exchange to track communications
- Connectionless takes things as they come



Telephone


Telegraph

© Copyright 2005 Howard Goldstein Associates, Inc. HGAI All Rights Reserved.



72

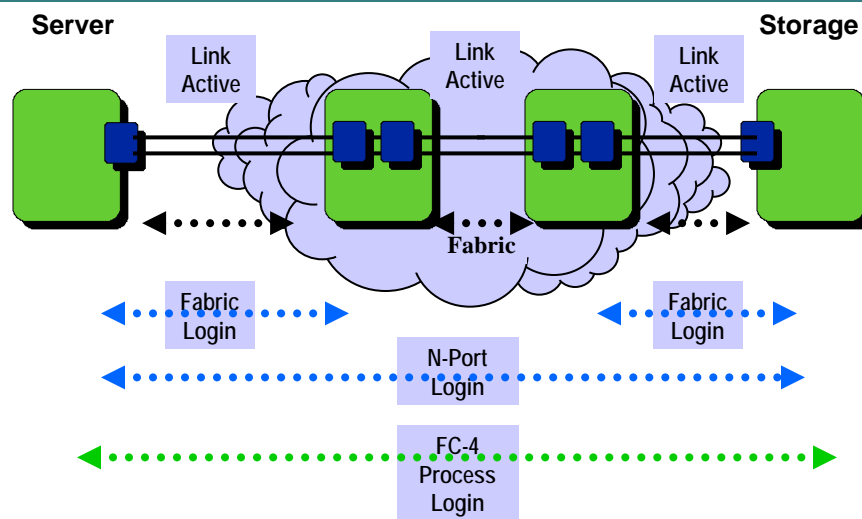
Connection-Oriented vs. Connectionless



- Internet Protocol Suite
 - TCP connections start with negotiation setup between hosts
 - Similar to telephone connection
 - Call, Answer, Talk, Hang up
 - IP Connectionless – Handle packets as they come
- Fibre Channel Class 2, 3 Switch Routing




Connection-Oriented FC Sessions - Login



Layered Protocols



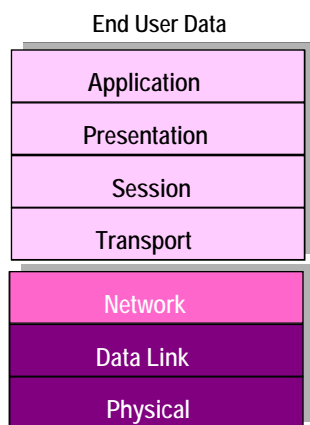
- Protocol Data Units (PDUs)



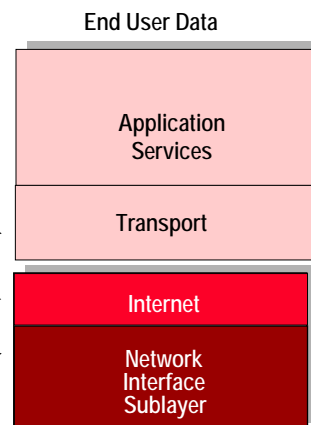
OSI & IPS (TCP/IP)



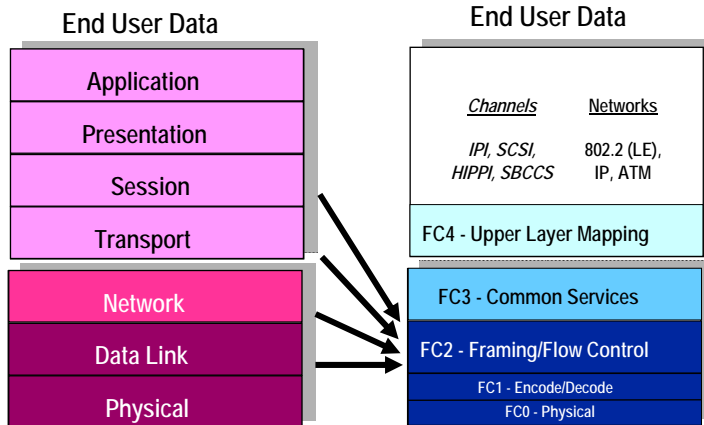
Open Systems Interconnection



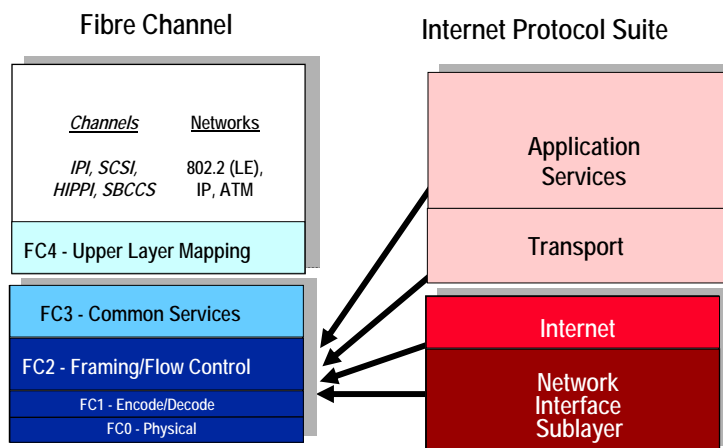
Internet Protocol Suite



OSI & Fibre Channel Architecture



Fibre Channel & IPS



Flow Control

- Internet Protocol Suite
 - TCP Sliding Window
 - Advertise Larger Windows
 - Congestion Control
- Fibre Channel
 - Credit-Based Fibre Channel
 - Advertise Credit

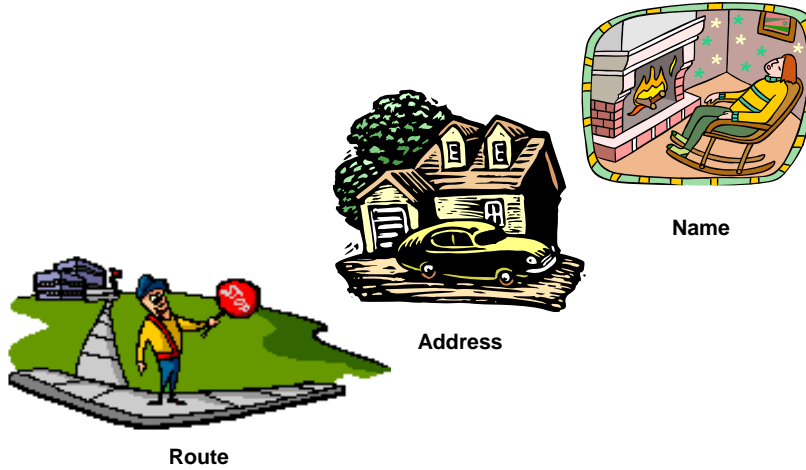


Payload Integrity

- Do you always want payload integrity?
- Not necessarily with Video and Voice
- Performance vs. Quality tradeoffs
 - A pixel in a online-rented movie frame
 - A medical imaging application just prior to heart surgery where you want to be sure the speck of grey isn't line noise!



Name, Address & Route





Name, Address & Route

- Addressing and Identification schemes

- Address is used to locate

- Names or identifiers uniquely represent something or someone



Name, Address & Route



- Name is a specific identification
 - Machine, User, Application - Unique target for data
 - Internet
 - Fully Qualified Domain Names WWW.HGAI.COM
 - Media Access Control MAC ID – Physical Port Identifier
 - Storage
 - World Wide Node Names WWNN (Fibre Channel)
 - World Wide Port Names WWPN (Fibre Channel)
 - Enterprise Unique Identifier EUI (iSCSI)
 - iSCSI Qualified Name IQN (iSCSI)
- Address identifies where the target is located
 - Hardware, Network, Port, Socket
 - Port ID FCID (Fibre Channel)
 - IP Address
- Route identifies path to get to the destination



Transport Addressing Ports



- Ports are used to deliver information to relevant application services
- Well-known and “Not-so-well-known” Transient

Port	Name	Description
20	FTP-3DATA	File Transfer (Data Channel)
21	FTP	File Transfer (Control Channel)
22	TELNET	Telnet
25	SMTP	Simple Mail Transfer
53	DNS	Domain Name Services
3260	ISCSI	Internet SCSI

- Comparable to Well known FCID Port ID Addresses
 - FFFFC Name Server
 - FFFFFD Fabric Controller



Hierarchical, Routable Addressing



- IP addresses can be divided into three parts
 - Network Address (NETID)
 - Subnetwork Address using (Subnetwork Mask 255.255.255.0)
 - Host IP Address
 - Example 172.16.2.6
 - Network 172.16.0.0, Subnetwork 172.16.2.0, Host IP Address 172.16.2.6
- FC addresses can be divided into three parts
 - Domain (Domain ID)
 - Area (Area ID)
 - Device (Device ID)
 - Example 0102EF
 - Domain 01 (Switch 01), Area 02 (Port 02 on the Switch), Device EF (ALPA of Loop Device)



SATA Address Comparison – 1 Gram



- **SATA Point-to-Point $2^0 = 1$**



- **SATA II “Port Multiplier” $2^4 = 15$**



SAS Address Comparison – 1 Gram




- **SAS “Expanders” $2^{14} = 128 \times 128 = 16,384$**



FC Address Comparison – 1 Gram



- **All Fibre Channel $2^{24} = 16,777,214$**



IP Address Comparison – 1 Gram



- All IPv4 – $2^{32} =$
4,294,967,292

- All IPv6 – $2^{128} =$
340,282,366,920,938,463
,463,374,607,431,768,21
1,456



X 56



Address Assignment



- DHCP Mechanism for allocating and assigning IP addresses to clients
 - Automatic: Host requesting an address are provided with a permanent IP address
 - Dynamic: Hosts requesting an address are provided with a temporary address
 - Manual: Host IP addresses are manually configured and DHCP just delivers these assignments
- Fibre Channel
 - Loop Initialization LISM, LIFA, FIPA, LIHA, LISA, LIRP, LILP
 - Select Master, Fabric Assigned, Hard Assigned, Soft Assigned, Report Position, List Position
 - Fabric Login Controller
 - Point-to-Point - Lowest World Wide Port Name (WWPN) Assigns



Name Service



- Internet
 - Domain Name Services for IP
 - Fully Qualified Domain Name FQDN to IP Address
 - Running on distributed servers
- Storage
 - Name Services for Fibre Channel
 - WWPN to Fibre Channel Identifier FCID, Fabric service running in the Switch
 - iSNS Storage Name Service (iSNS) for iSCSI



Routing in Routers, Routing in Switches




- Routing
- Switch frame “steering” over one “appearance of media”
- Router packet “forwarding” across many “appearances of media”
 - A Day in the Life starts at home:
 - Walk to the bus stop and get on the bus
 - Take the bus to the dock
 - Walk on to the Ferry boat, sail across the river
 - Get off the boat and hop in a cab.
 - Drive to the office and get out of the cab.
 - Run to the elevator and hop in
 - Jump off the elevator and walk to your desk



Segmentation & Hardware Offload



- Segmentation, Convergence and Reassembly



- Hardware Offload for Performance Acceleration

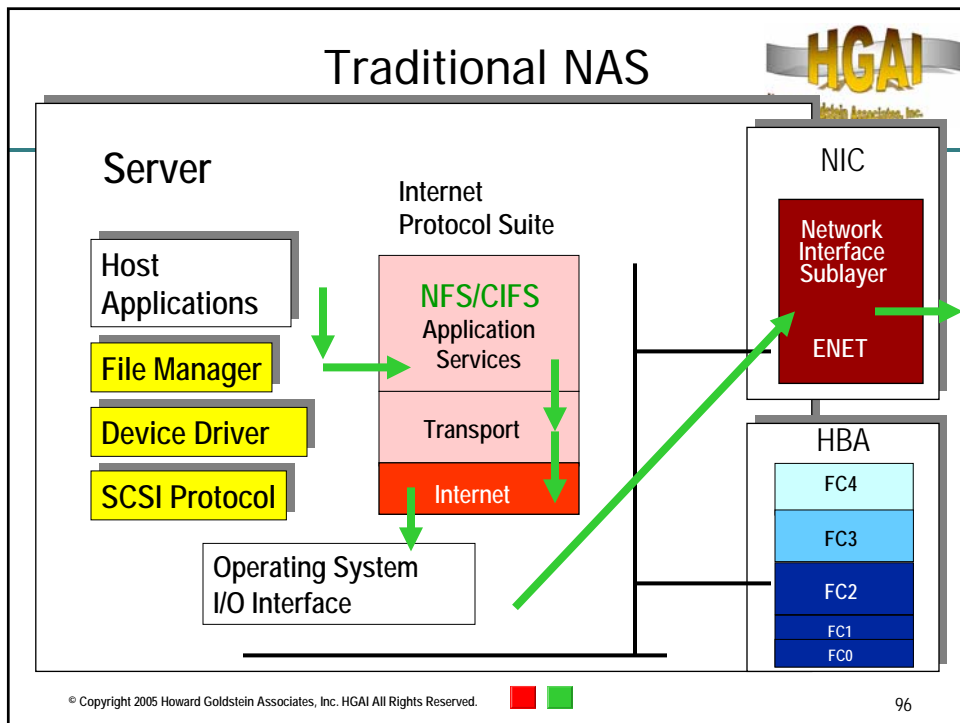
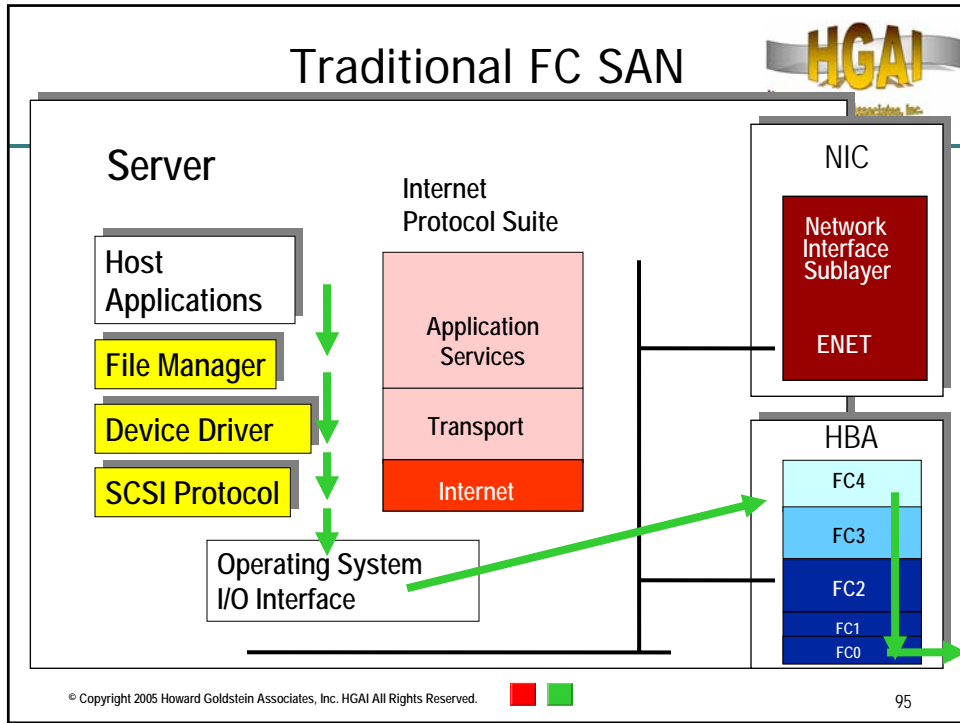


Server to Storage Information Flow Requirements



- Applications have no way to qualify information flow requirements to storage
- Operating Systems have few interfaces to request special network services such as QoS
- In many ways, the Operating Systems are the most primitive network requesters in computer systems
- Traditional approach requires as fast as possible interfaces
 - Low Latency, High Bandwidth

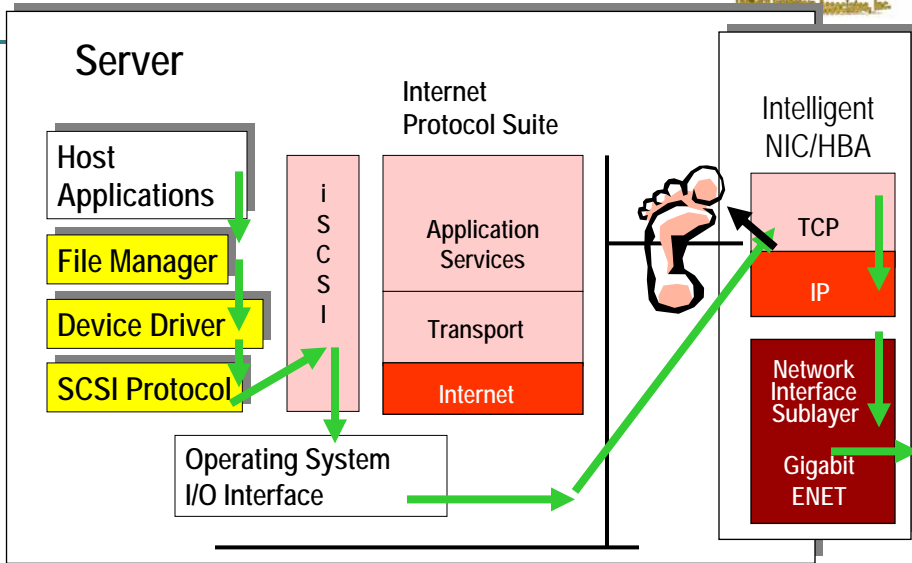




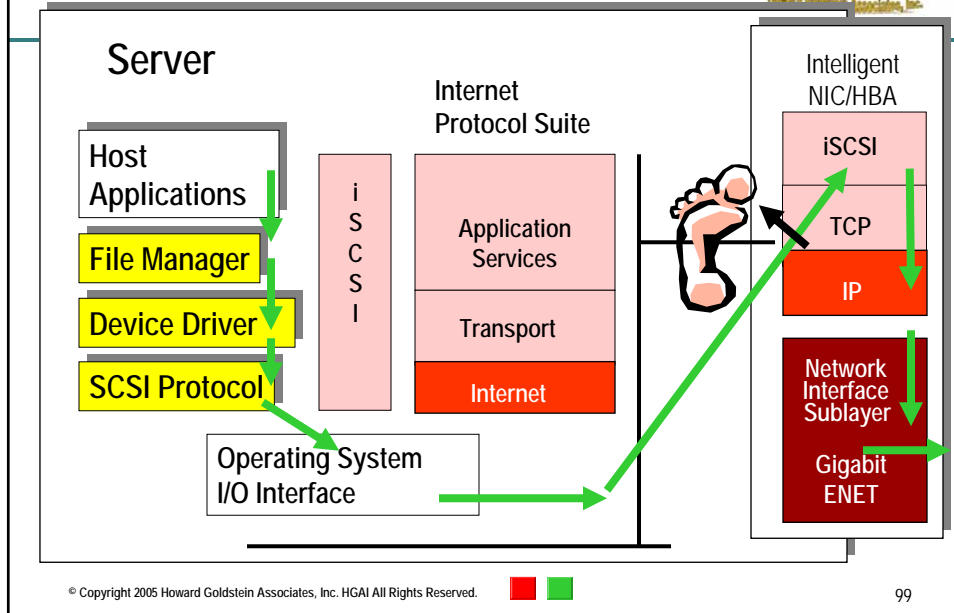
What is iSCSI?



iSCSI – TCP/IP Offload




iSCSI – Full Offload



What is a Physical Transport Network?



- A group of connected systems - at least two!
 - Host, Node, Computer, Server, Workstation, Printer, Router, Storage Device
- Network Interface Cards (NICs), Host Bus Adapters (HBAs), Storage Adapters (SAs)
- A medium or the “appearance of” a connection medium
 - Unshielded Twisted Pair(s) (UTP), Coaxial Cable(s), Multi-Mode Fiber(s), Leased Line(s), RF
- A connection protocol or method
 - Ethernet, Token Ring, PPP, Frame Relay, ATM, Fibre Channel, SONET/SDH, DWDM



Network Components



- Network Topologies
 - Data Networks
 - Point-to-Point, Bus, Token Ring, Fabric
 - Storage Networks
 - Bus, Arbitrated Loop, Point-to-Point, Fabric
 - Hybrids
- Network Types
 - Local Area, Metropolitan Area, Wide Area, Storage Area
- Backbone Networks
- Network Interconnectivity
 - Repeaters, Hubs, Bridges, Switches, Routers, Gateways



Storage Network Choices



“Very Small

Integrated Device Electronics (IDE)
Small Computer Systems Interface (SCSI) Bus
Enterprise System Connectivity (ESCON)
Serial ATA, Serial Attached SCSI

“Very Skinny

Fibre Channel



Storage over IP

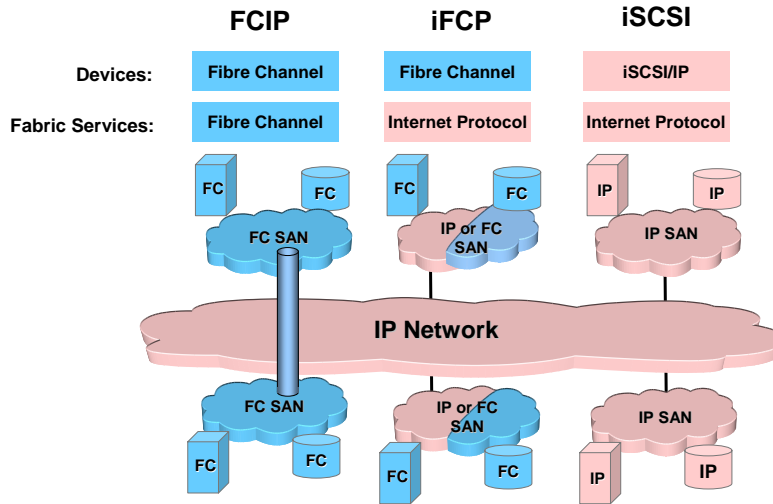
- iSCSI/TCP/IP/GE
 - internet SCSI (iSCSI)
 - Transmission Control Protocol (TCP)
 - Internet Protocol (IP)
 - Gigabit Ethernet (GE)
- internet Fibre Channel Protocol (iFCP) & internet Storage Name Service (SNS)
- Fibre Channel Internet Protocol (FCIP)



Infiniband



IP Storage




Beware of Semantics



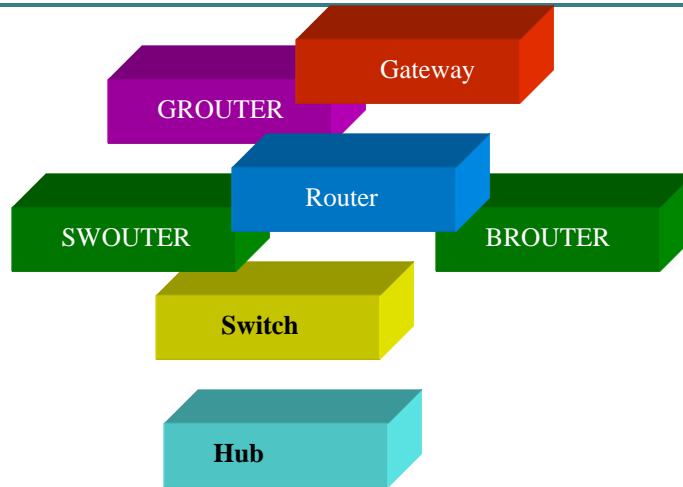
Network terminology misuse is rampant

- By the Storage Network Industry
- By Vendors
- By Media
- By Educators

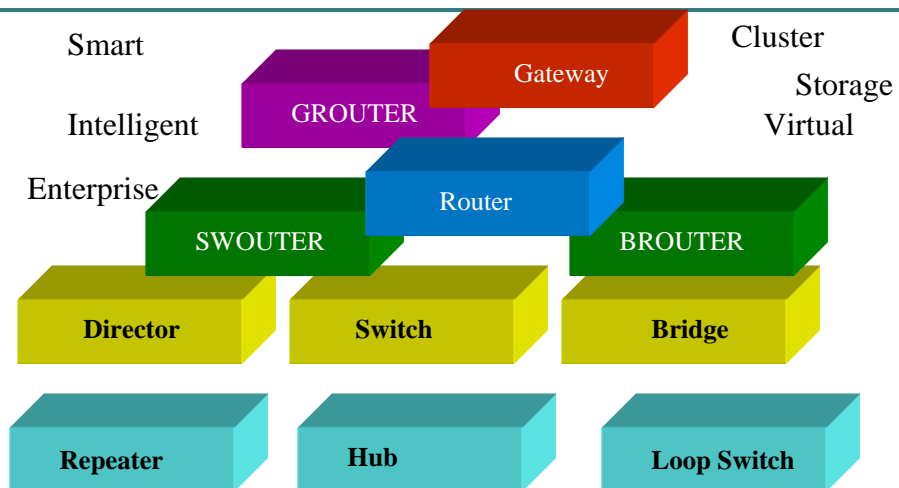
Confusing the customer!



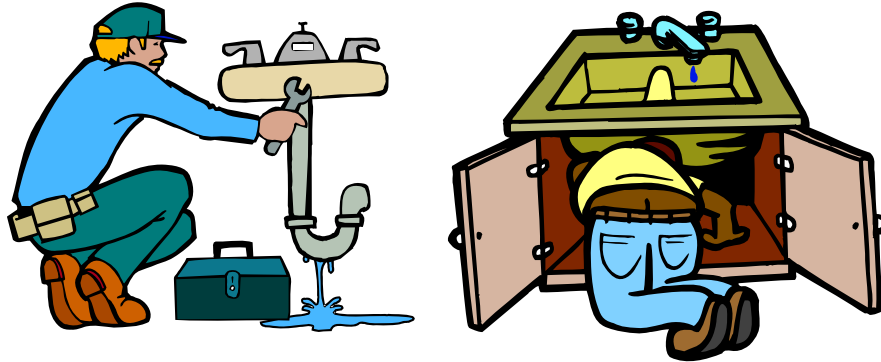
Beware of Semantics



Beware of Semantics



Plumbing is Beautiful




© Copyright 2005 Howard Goldstein Associates, Inc. HGAI All Rights Reserved.



107

Storage Networking Technology Great Expectation Waves




© Copyright 2005 Howard Goldstein Associates, Inc. HGAI All Rights Reserved.

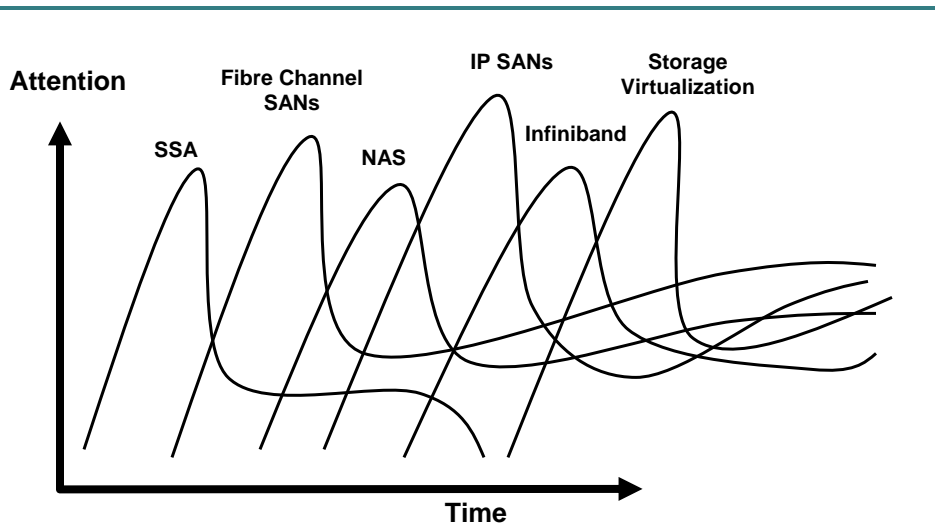


108

Storage Networking Great Expectation Wave



Storage Networking Technology Great Expectation Waves




Storage Networking Futures



- Upper layer protocol changes
 - Lazy and Hyper reads and writes
 - System SAN awareness, Dynamic discovery
- Higher bandwidth rates for aggregation and higher speed processors for servers, storage and adapters
- External SANs with Fibre Channel and iSCSI/TCP/IP/GE
- Internal SANs with SAS and SATA
- NAS & SAN integration



Storage Networking for Network Professionals




This concludes
Storage Networking for Network
Professionals



Thanks for attending and hope to see you again

HGAI offers training!
WWW.HGAI.COM



Good Luck!

