

# iSCSI SCSI over the Internet

Thomas Jerry Scott

Will It Really Happen?

# Historical Background

- u SCSI Drives and Drive Chains
  - External Drive Capability
- u RAID Level 0
  - Disk Mirroring
- u RAID Level 5
  - Added Security with Hot Swapping
- u RAID Level 10
  - Faster Writes, Full Redundancy

# Network Attached Storage

- u User Desktops connect to storage via Network.
  - Often TCP/IP
  - TCP/IP Settings must be properly setup
  - Good for items like on-line backup
- u The NAS “Good News”
  - Everyone has TCP/IP, so connections easy
- u The Knock on NAS
  - Slower than directly attached storage

# Fiber Channel Disk Systems

- u Move disks out of server into a “Disk Farm”
  - Facilitates adding drives to farm as storage needs arise
  - Requires special “Fiber Channel Disk Controllers” often called HBA’s
    - u HBA’s are expensive, typically in the \$1K range per HBA
  - Disk farms use standard SCSI disk drives

# Storage Attached Networks

- u Normally built on Fiber Channel physical disk environments
- u Redundancy benefits using RAID approaches
- u A single SAN can provide storage for different servers
- u More disk storage space can be added to different servers by adding drives to the SAN

# Storage Area Networks -- 2

- u SANs provide many benefits
  - Unify storage for many servers
  - Easier backups for the enterprise
  - Facilitate storage growth
  - Easier for many servers and desktops to share the same data
  - Provide volumes that may exceed physical disk sizes
  - Typical speeds are 1 and 2 Gigabits/sec

# iSCSI Motivation

- u Modern computing environments
  - Have many servers, often over 100
- u If all servers have internal disk farms
  - How to back up the 100 Servers?
- u Moving the storage to a disk farm
  - Can facilitate the backup problem
  - Several, instead of 100, high speed backups on the SAN

# iSCSI Motivation -- 2

- u Demand for better storage management is the catalyst for change
- u Demand for iSCSI created by network administrators unwilling to standardize on NAS or FC SAN.
- u iSCSI hopes to create a standard way of managing storage, using TCP/IP networks
  - Deliver the ease and cost effectiveness of NAS
  - With the reliability, performance and scalability of a Fiber Channel SAN.



# iSCSI and NAS

- u NAS devices are storage devices
  - That include NICs
  - Accessible via network connections
- u iSCSI devices are different in that they are accessed via an iSCSI HBA
- u To the computer, the iSCSI HBA looks just like any other SCSI HBA
- u To the network, the iSCSI HBA appears to be a NIC

# iSCSI Basics

- u iSCSI runs on a TCP / IP / Ethernet network
- u IP packets are routed
  - IP routing can provides resiliency if a intermediate node fails
- u TCP provides the reliability for iSCSI
- u 1Gbps switched Ethernet Networks provide the necessary speed for many applications

# iSCSI Basics -- 2

- u iSCSI transmits native SCSI commands and data over TCP/IP
- u Host and physical storage system can be separated by IP Networks
- u iSCSI transfers and stores SCSI commands and data at any iSCSI enabled storage location
- u Host and Storage could be connected via a LAN, a WAN, or the Internet

# iSCSI Basics -- 2

- u When an iSCSI server needs to move some data to storage,
  - It transfers the data to its HBA, where it becomes standard SCSI data.
  - The data is then enclosed in a TCP/IP packet and is sent out via the Ethernet network.
  - Once it gets to the iSCSI storage device, the TCP/IP packet information is stripped off
  - The data is then moved to the device's internal SCSI controller
  - Which in turn transfers it to disk.

# iSCSI and TCP/IP

- u Standard SCSI commands and data are encapsulated as a serial string of bytes preceded by an iSCSI header
- u The TCP/IP layer breaks this into packets to transfer the data between an iSCSI enabled host and the remote storage system.
- u If a read request for data has been sent, the data blocks are retrieved from the physical drives, encapsulated back into iSCSI and returned to the requesting host via TCP/IP.

# iSCSI and Ethernet Networks

- u Ethernet is the most common LAN today
  - Network administrators understand and use Ethernet on a regular basis today
- u iSCSI does not require unique equipment or cabling
  - Uses regular Ethernet cabling and switches
- u 1Gb Ethernet switches and NICs are common
  - 1Gb Ethernet cards are in the \$125 cost range now and will get lower as adoption happens

# The 10Gbps Ethernet Standard

- u The 10Gigabit Ethernet 802.3ae standard
  - Approved by IEEE, June, 2002
  - Preserves the 802.3 Ethernet frame format
  - Preserves the minimum and maximum frame size of the current Ethernet standard
  - Supports full duplex operation only
  - Specifies an optional media independent interface (XGMII)
  - Supports a speed of 10 Gbps at the MAC interface
  - Defines two families of PHYs - LAN and WAN



# iSCSI Graphical View

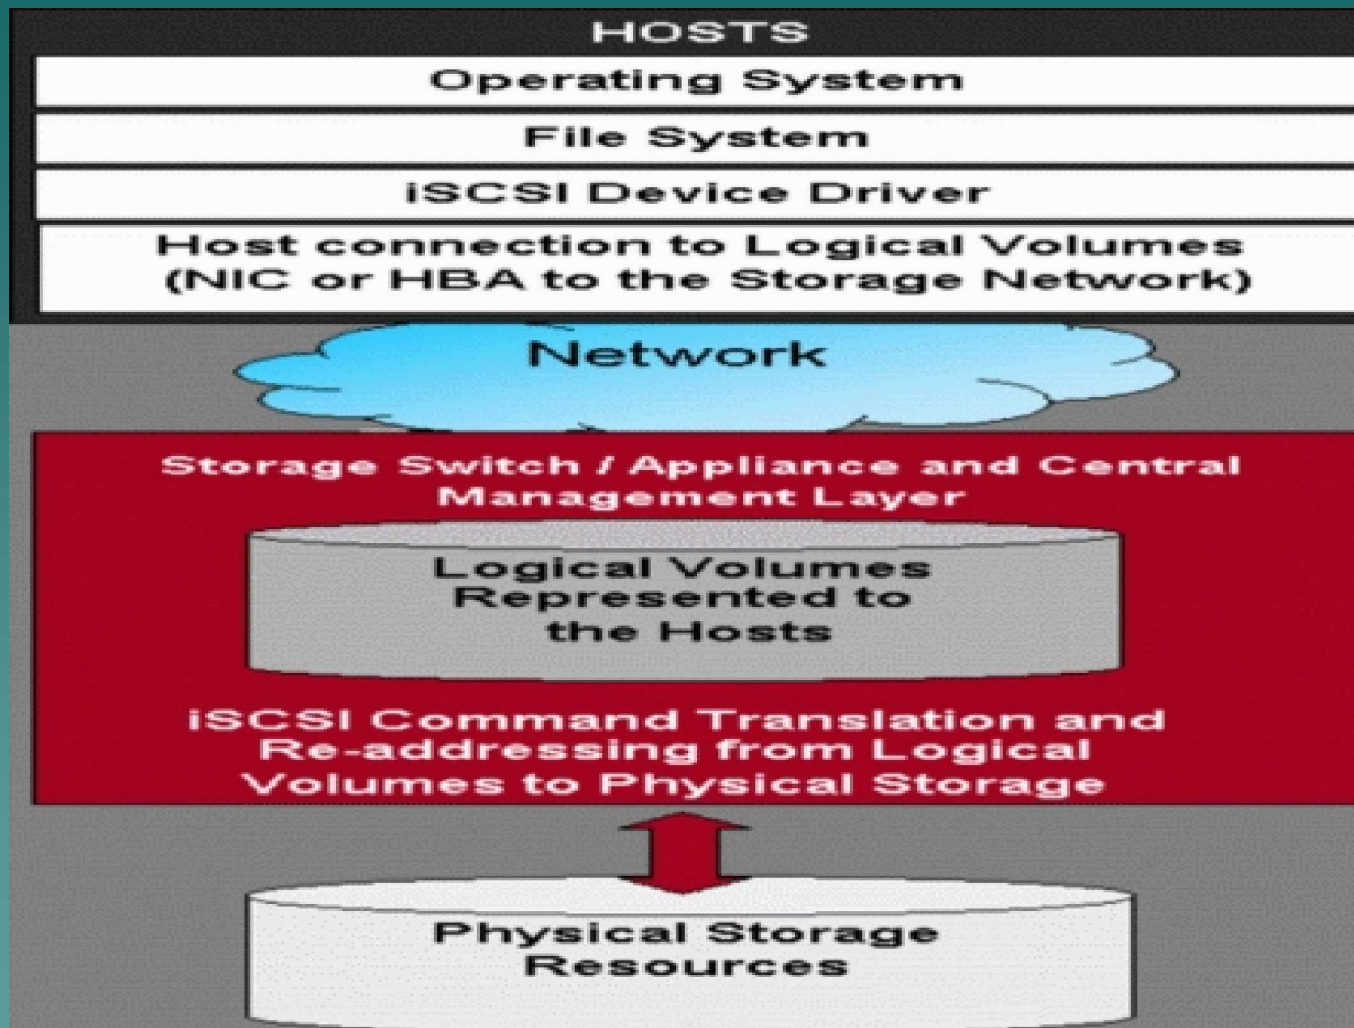


Figure 1



# Questions for iSCSI Adopters

- u How will we upgrade to 1Gbps networks?
- u How much will this cost?
- u Where's the storage management?
- u Where's volume creation and administration of policies?
- u Where's access control, efficient storage utilization, and load balancing?
- u What about legacy storage support and a simple user interface?

# Virtualization Software

- u Virtualization is available on high-end raid systems
- u Using virtualization, the network administrator can view all networked storage devices on a single screen
- u Storage resources can be gathered into a single storage resource pool
- u A network administrator can now define new logical volumes from the pooled resources

# Logical iSCSI Volumes

- u Logical volumes
  - Can take almost any shape and size
  - Can be independent of physical barriers
    - u Such as raid enclosures, physical disks, protocols, and distance
- u These logical volumes can include
  - Small volumes of a few MBs
  - Large volumes of several TB
  - Mirrored volumes or striped volumes

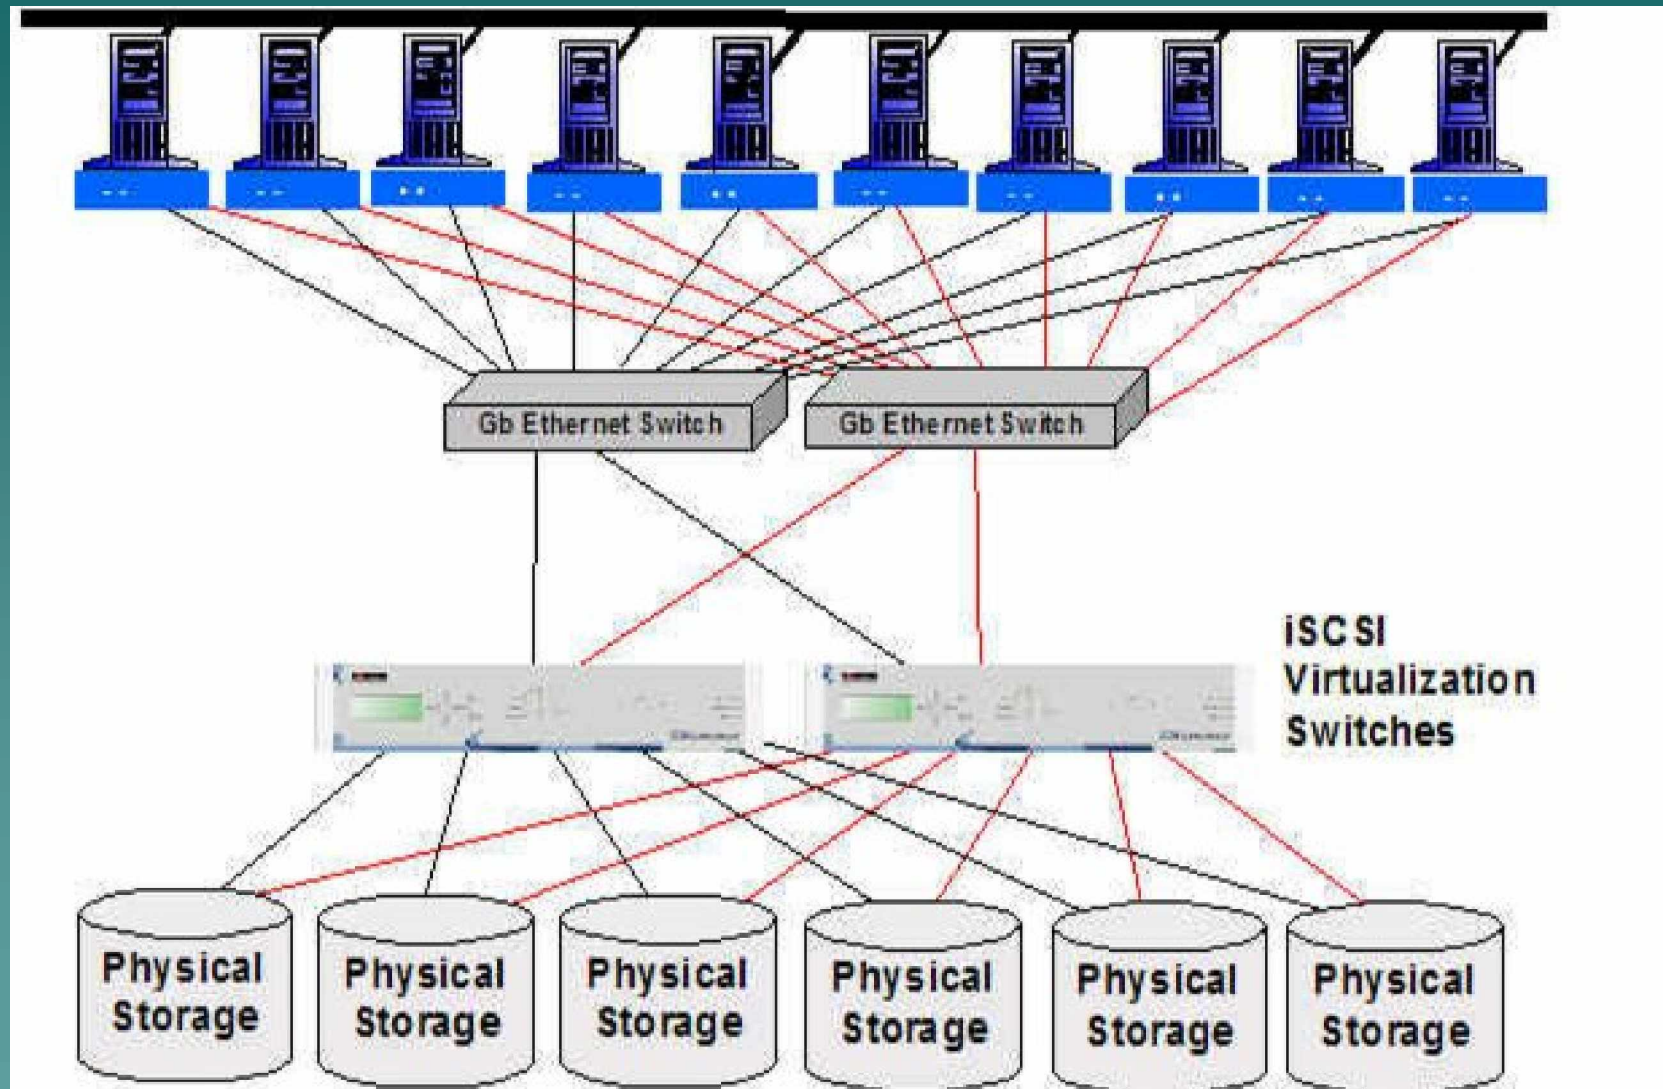
# Administering iSCSI Systems

- u From a central and simple interface, network administrators can
  - Administer common policies and services across the entire pool of logical volumes
- u This can be independent of the vendor brand, type, and protocol represented by each physically attached storage system.

# Administering iSCSI Systems -2

- u The network administrator must
  - Assign the logical volumes
  - Then apply secure access permissions to the hosts wanting to access the volumes
- u Without secure access permission
  - Hosts are not aware of and cannot access the logical virtualized volumes.

# iSCSI Virtualization

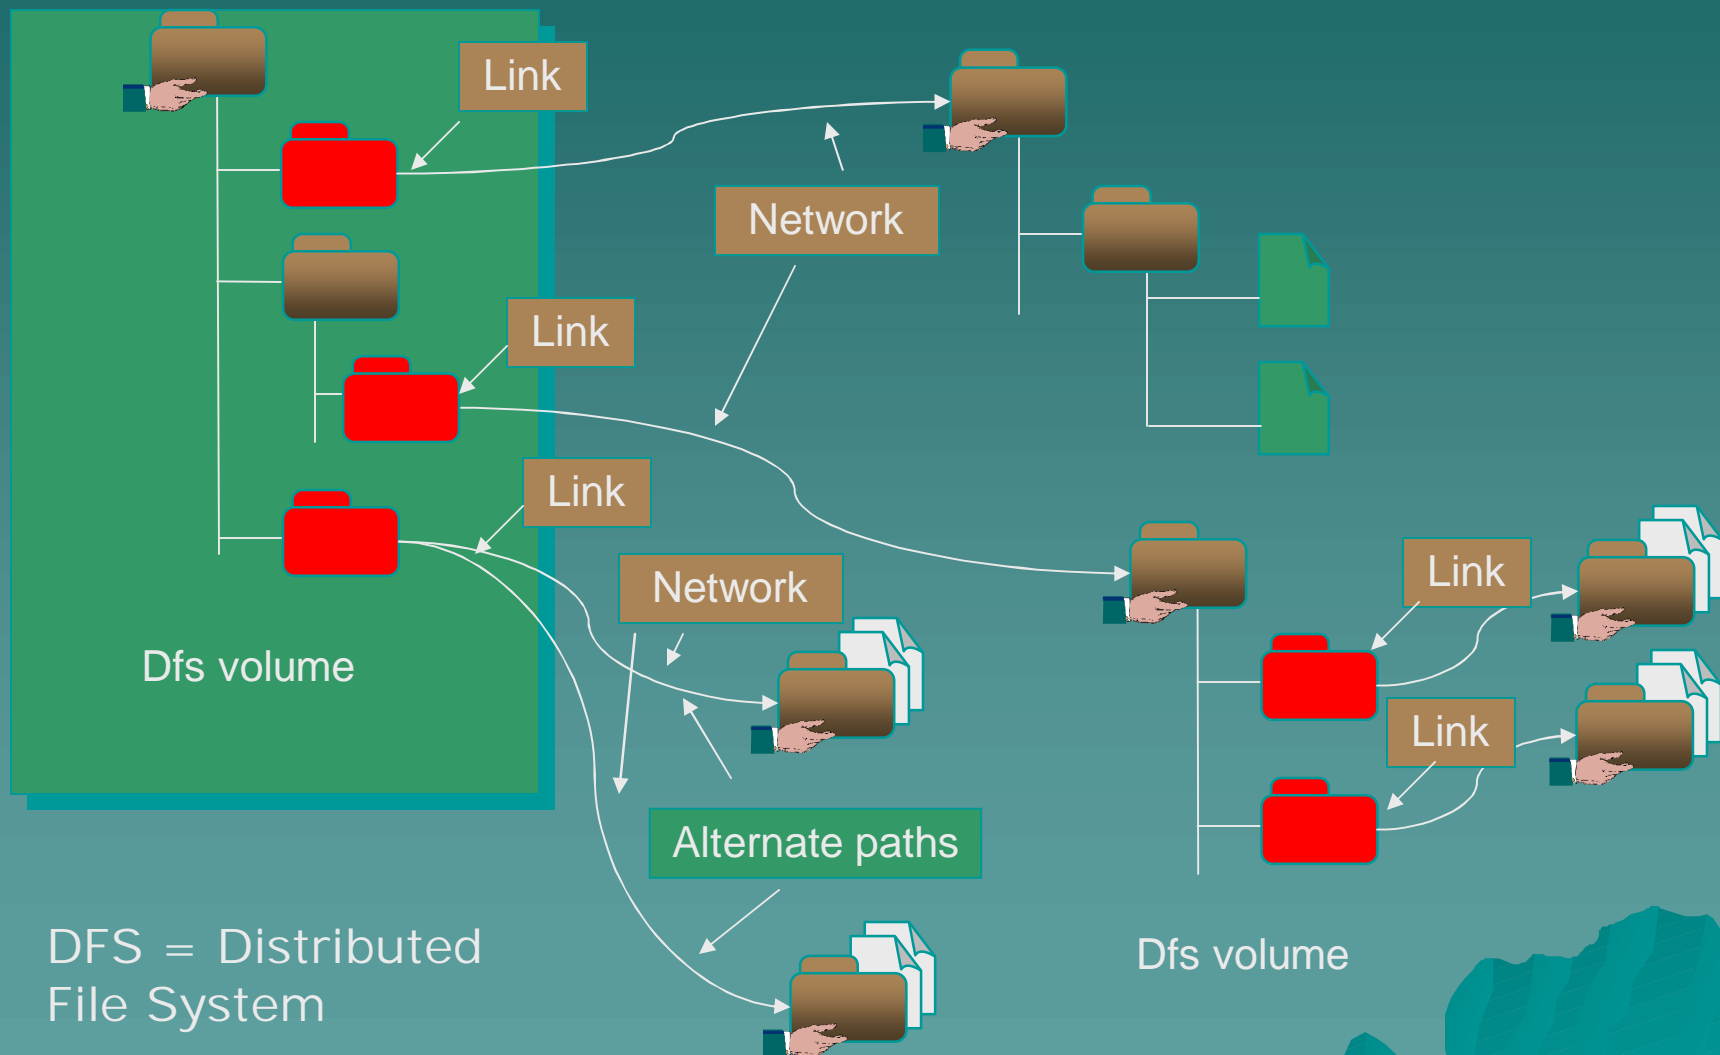


# iSCSI in Operating Systems

- u iSCSI support is becoming part of every popular server platform
  - Microsoft, Linux, Solaris, HP, and IBM
  - Microsoft's 2003 Storage Server Initiative
- u iSCSI 1Gb Ethernet and Virtualization
  - Combine to deliver a standard way to manage small to enterprise class storage networks.



# MicroSoft DFS Capability in W2000



DFS = Distributed File System



# iSCSI Cost Issues

- u Using storage routers IT managers will be able to hook more servers into their storage networks
  - Take advantage of existing IP networks
- u IT administrators can use Gigabit Ethernet adaptors instead of Fibre Channel HBAs
  - \$125 for Gb, vs \$1K for HBA
- u More skilled people in Ethernet than in Fibre Channel is another savings.

# iSCSI Performance Concerns

- u 2 Gbps Fibre Channel networks are faster than 1 Gbps iSCSI networks!
  - Faster Line Speed
  - Less overhead than TCP/IP packets
- u Still, many currently used client-server applications currently run well over 10/100 Mb Switched Ethernets
  - Hospital data processing
  - On Line data entry systems
  - Many others
- u 1Gbps iSCSI solutions will make these much faster and easier to manage

# When Not to Use iSCSI

- u IT managers shouldn't use iSCSI
  - In applications that require very high throughput, such as graphics intensive applications
  - In applications that need to process large amounts of transactions per second
- u Improving iSCSI performance
  - 10Gbps Ethernet
  - 10Gbps Ethernet to the WAN
  - TCP/IP offloading iSCSI adapters..

# The Future for iSCSI

- u Ethernet Issues
  - 1 Gig Ethernet available now
  - 10 Gig Ethernet on the way
  - 10 Gig Ethernet standard adopted
- u Storage Management driving the need
  - The need for storage is increasing
  - Better management is always necessary
  - Industry moving in iSCSI direction
- u 2004 and beyond .....