Technical Symposium

Introduction to Storage Area Networking

Joshua Kaya, Consulting Systems Engineer
Cisco Systems Canada, Toronto
jkaya@cisco.com

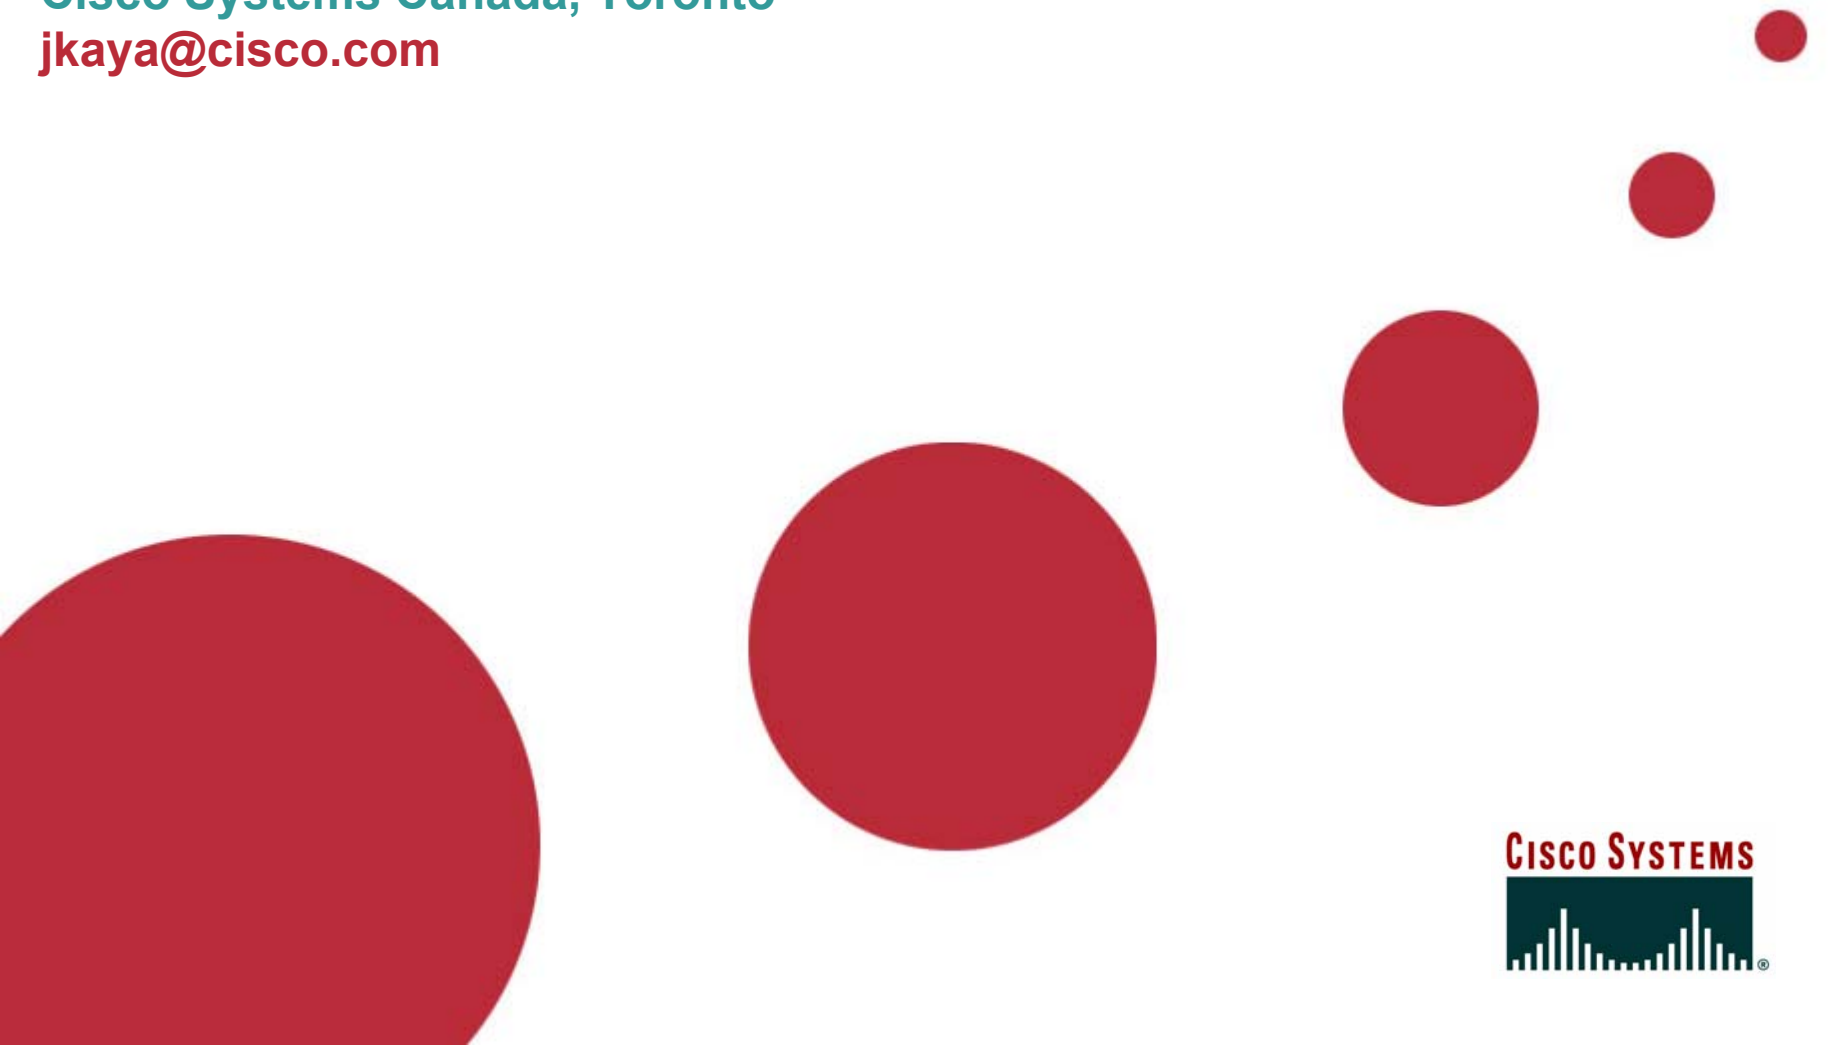

Cisco Systems

Agenda:

- Hard drive basics
- Storage Architectures:
  DAS, NAS and SAN
- SCSI
- Fibre Channel
- FCIP & iSCSI

Computing Origins

- **1500: Leonardo da Vinci designs a mechanical calculator**
- **1906: The vacuum tube is invented by American physicist Lee De Forest**
- **1941: Konrad Zuse introduced the first programmable computer called the Z3**
- **1956: The Air Force accepts the first UNIVAC Solid State Computer**

Computing Origins

- **1956: IBM introduced the 305 RAMAC (Random Access Method of Accounting and Control), which stored 5 million characters (~5MB) on *50* disks, each 24 inches in diameter**
- **1962: IBM introduced the model 1301 Advanced Disk File. The key advance of this disk drive was the creation of heads that floated, or flew, above the surface of the disk on an "air bearing", reducing the distance from the heads to the surface of the disks from 800 to 250 microinches**
- **1973: 1973, IBM introduced the model 3340 disk drive, which is commonly considered to be the father of the modern hard disk. This unit had two separate spindles, one permanent and the other removable, each with a capacity of 30 MB**

History of the Hard Drive: Last 15 years

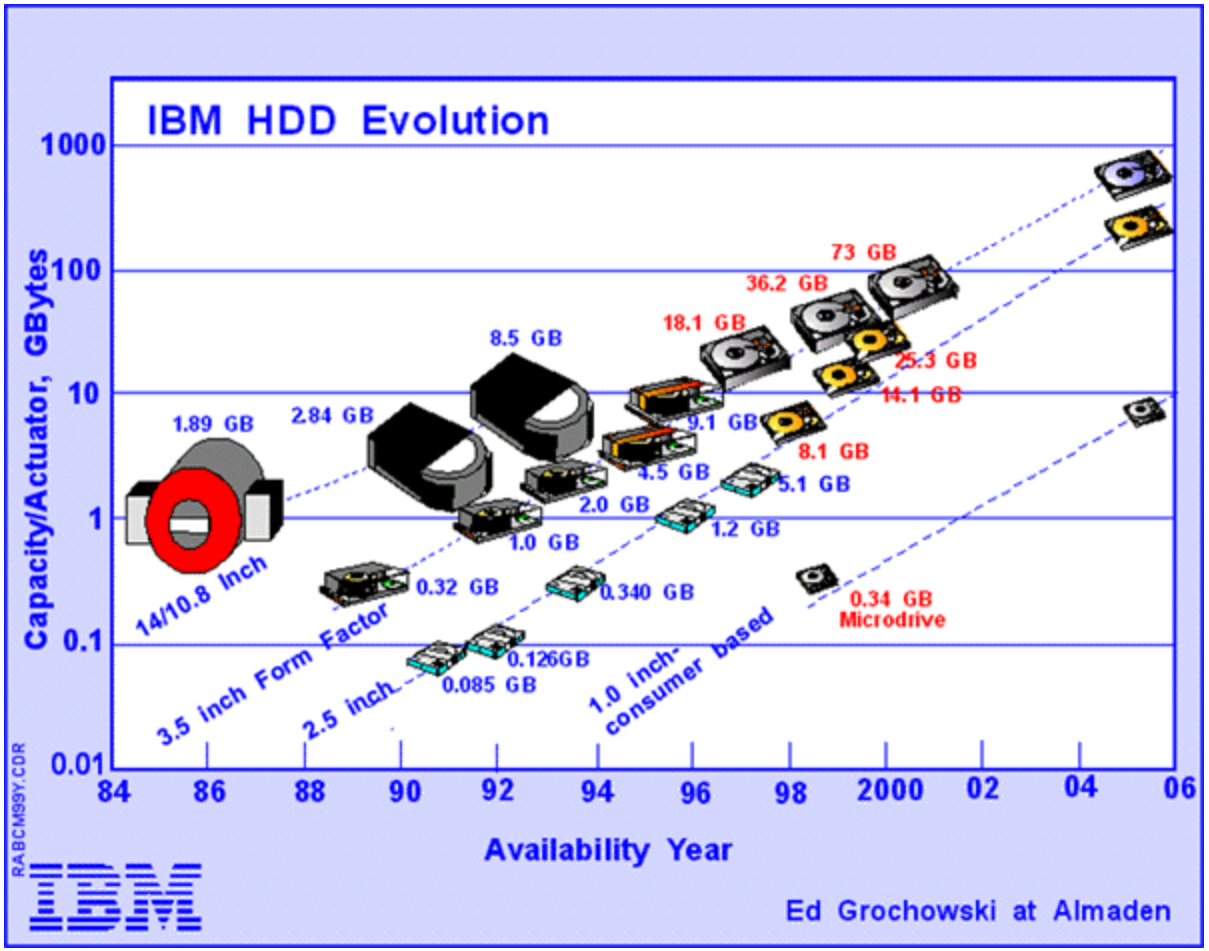

System/390 Announcement

IBM U.S. Marketing & Services press release distributed on September 5, 1990:

"In its most comprehensive announcement of products, features and functions in more than a quarter century, IBM today introduced System/390 -- the industry's most advanced integration of system architecture, design, technology and function.

…

Among the functional highlights are:

- Enterprise System Connection (ESCON) architecture that implements high speed fiber optic channels, allows customers to locate traditional "computer room" equipment anywhere within a 9-kilometer (5.6 miles) radius, and facilitates any-to-any connectivity"

Hard Drives Physical Characteristics

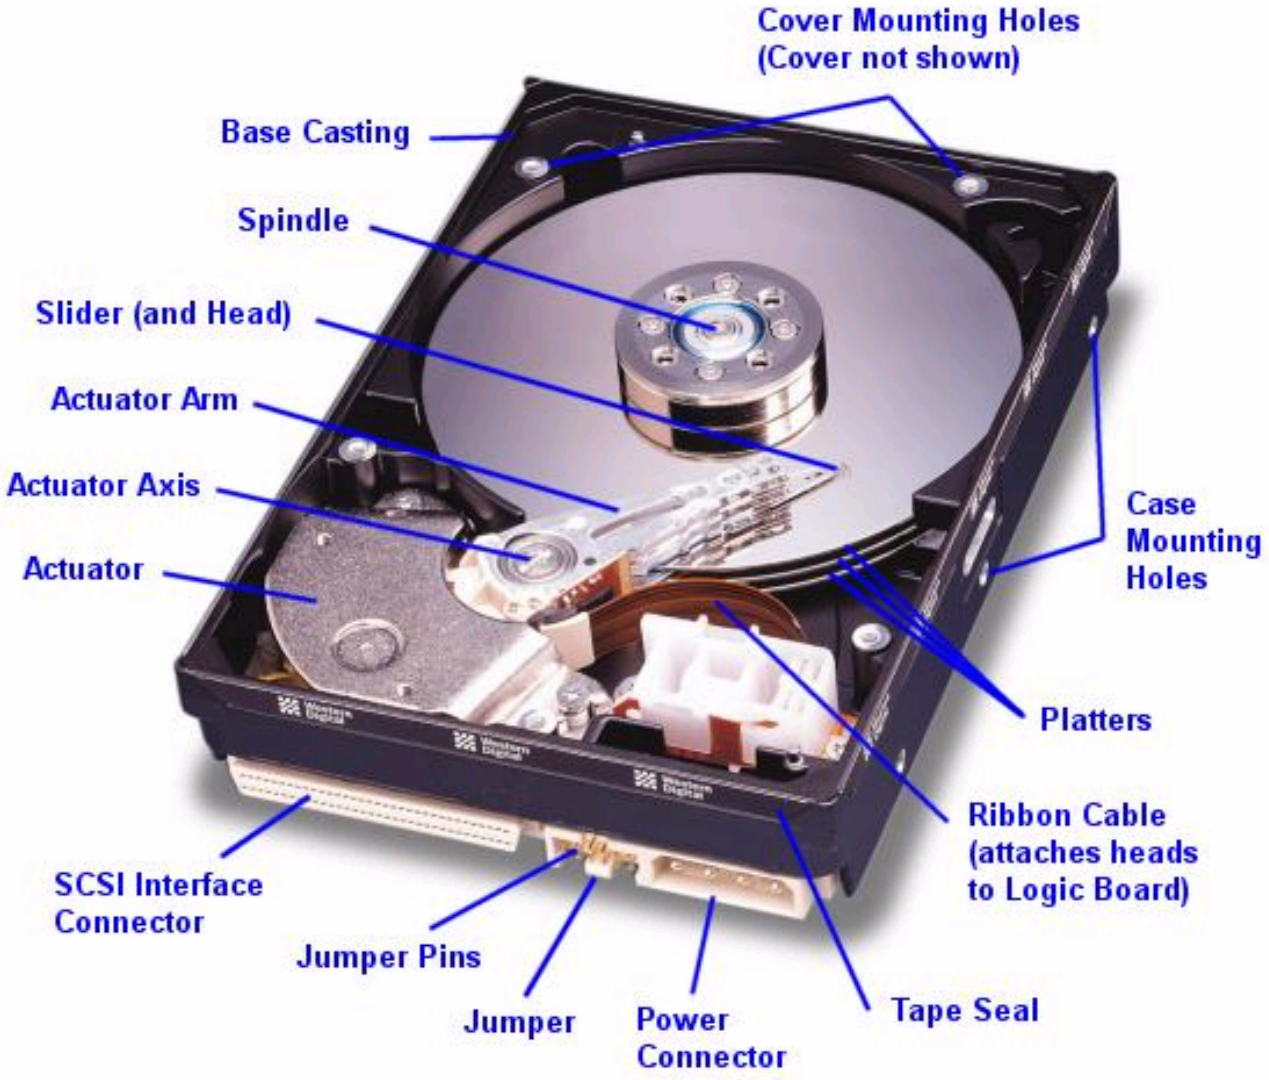

Hard Drive Essentials

- Hard drive platters are formatted into tracks, each of which contain a certain number of sectors.
- Low-level hard drive sectors are 512 bytes (few exceptions).
- Various size hard drives have varying number of platters, tracks and varying number of sectors per track.
- A particular combination of tracks and sectors is implemented at the factory during the low-level format.

Hard Drive Platter: Tracks and Sectors

- This single platter has 20 tracks and 258 sectors.
- Not all tracks have same amount of sectors: 16 in the outer tracks, 9 for the inner tracks.
- This characteristic is called Zoned Bit Recording (ZBR).
- Both sides of this platter are formatted similarly.
- Note that sectors on both sides are not necessarily aligned: this is called sector skew.

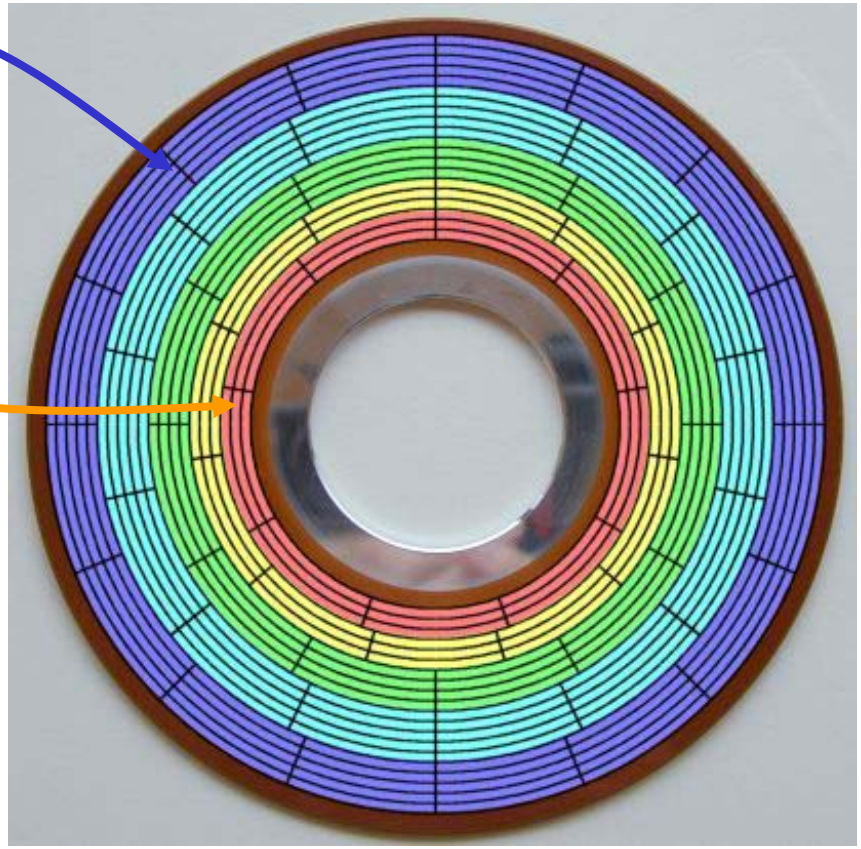

ZBR: Different methods

Zone	Tracks in Zone	Sectors Per Track	Data Transfer Rate (Mbits/s)
0	454	232	92.9
1	454	229	91.7
2	454	225	90.4
3	454	225	89.2
4	454	214	85.8
5	454	205	82.1
6	454	195	77.9
7	454	185	74.4
8	454	180	71.4
9	454	170	68.2
10	454	162	65.2
11	454	153	61.7
12	454	142	57.4
13	454	135	53.7
14	454	122	49.5

From Quantum Fireball Product Manual
© 1996 Quantum Corporation

Zone	Tracks in Zone	Sectors Per Track	Data Transfer Rate (Mbits/s)
0	624	792	372
1	1,424	780	366.4
2	1,680	760	357
3	1,616	740	347.6
4	2,752	720	338.2
5	2,880	680	319.4
6	1,904	660	310
7	2,384	630	295.9
8	3,328	600	281.8
9	4,432	540	253.6
10	4,528	480	225.5
11	2,192	440	206.7
12	1,600	420	197.3
13	1,168	400	187.9
14	18,15	370	173.8

From Deskstar 40GV and 75GXP Product Manual
© 2000 International Business Machines Corporation

Modern hard drives are low-level formatted at the factory – controllers don't know all this information

Logical Block Addressing (1)

- Addressing a single sector on a hard drive requires knowledge of the platter number, track number and sector number.
- Tracking these parameters is complex, and *varies between hard disk models*.
- Modern drives present each 512 byte sector as numbered memory location.
- This number is called the Logical Block Address (LBA).

Logical Block Addressing (2)

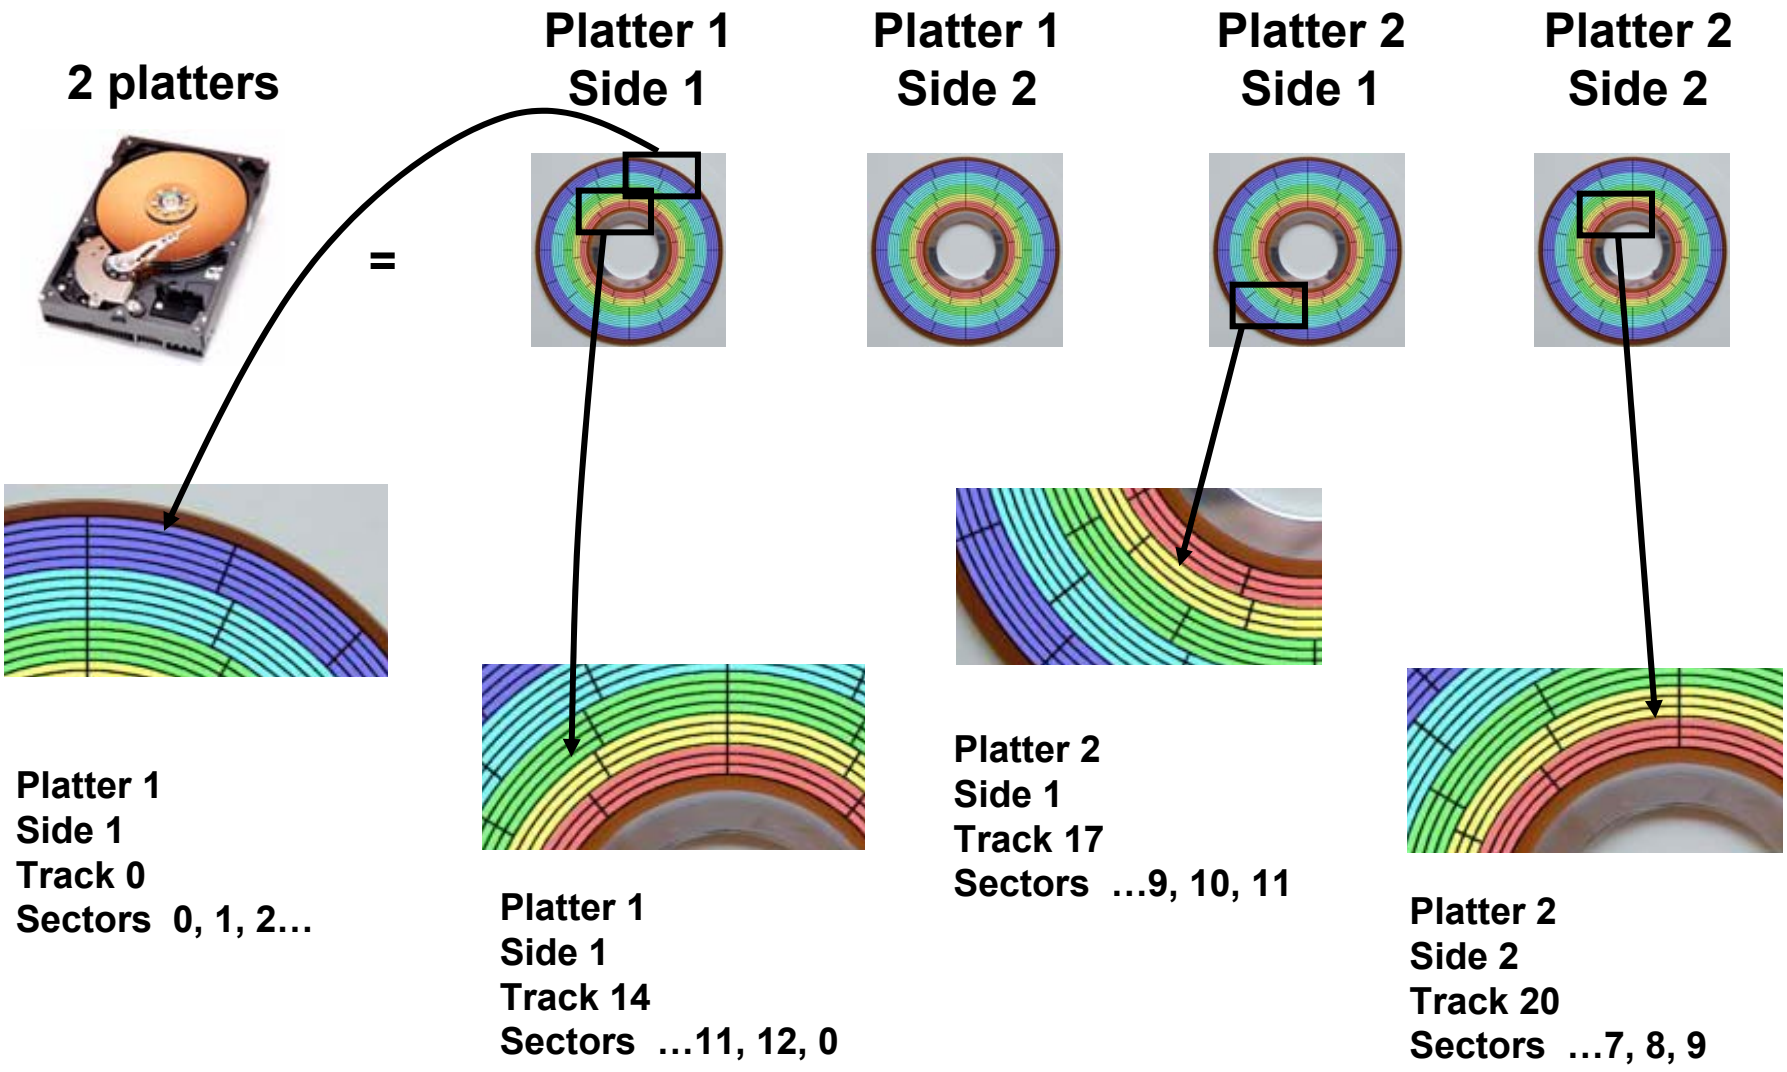

Logical Block Addressing (3)

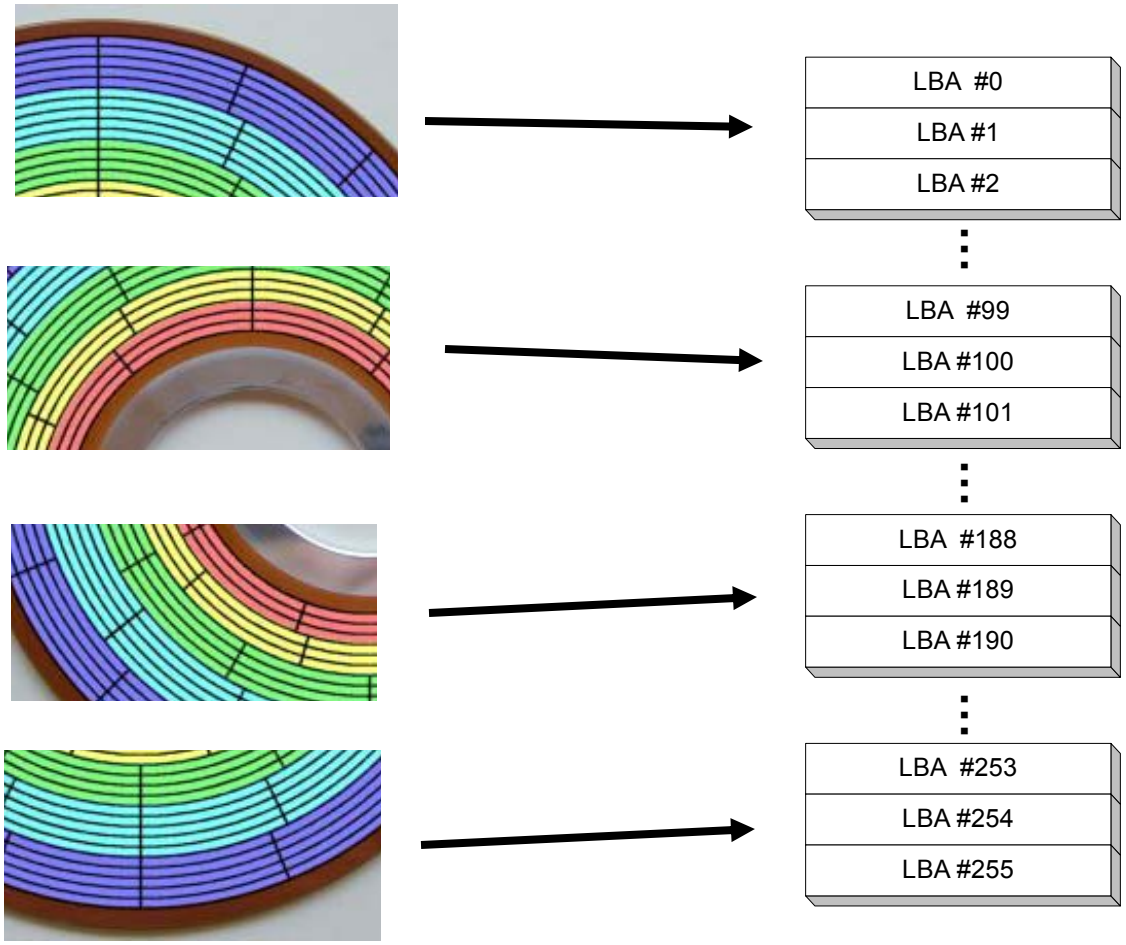

Total disk capacity:

256 Logical Blocs

X

512 bytes

---

131 072 bytes

What is Virtualization?

- Turning Physical Disks in to Logical Volumes, LUN's, Virtual Disk Management
- RAID Common Types today are 0, 1, 0+1, 1+0, 4, 5
- Snapshots, Checkpoints, 3rd Mirror Break Off
- Replication

RAID:

- Redundant Array of Independent Disks

- RAID feature can be implemented in two different manners:

  - Host based RAID (Software)

  - Controller based RAID (Hardware)

Host Based RAID

- **Veritas Volume Manager is an example of host based RAID which creates virtual structures and manage them on the host. Software RAID is layered on top of the operating system**

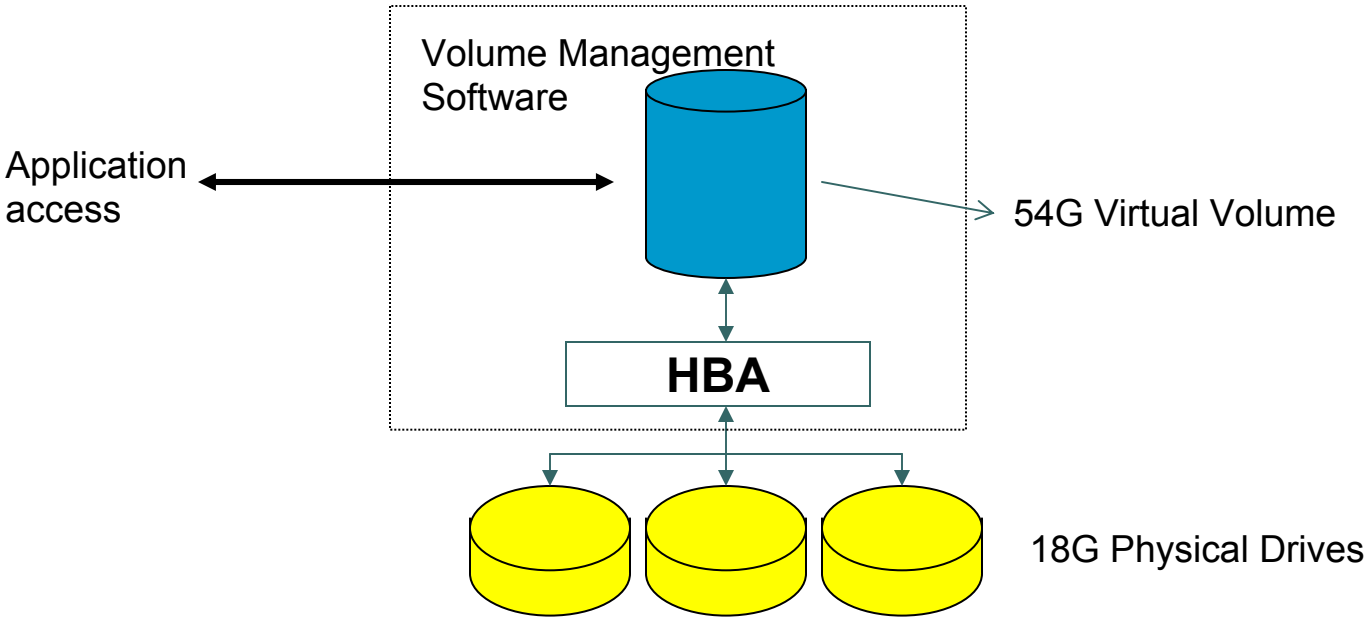

Sub-System Based RAID

- **RAID algorithms computed by ASICs**
- **Dedicated controller for RAID provides better performance.**

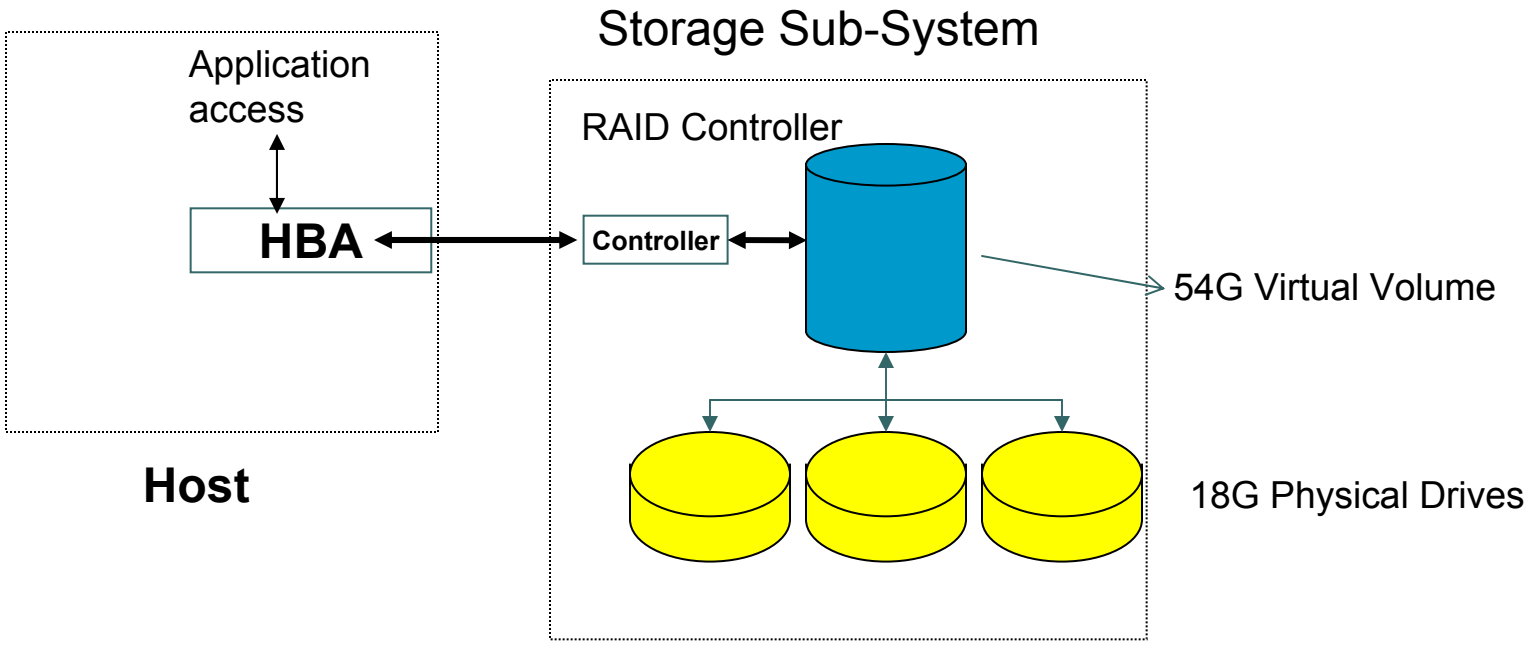

RAID levels

RAID Level	Description	Min Disks
0	Striping/Concatenation	2/1
1	Mirror	2
0+1	Striping/Concatenation then Mirror	4
1+0	Mirror then Striping/Concatenation	4
2	Hamming Code	N/A
3	Fix parity with concert I/O	N/A
4	Fix parity with Random I/O	3
5	Stripe with distributed parity with Random I/O	3 without log 4 with log

Parity: The quality or state of being equivalent

In datacom, parity calculations typically leverage the XOR hash/operation to provide error detection and correction.

1 0 = 1

0 0 = 0

1 1 0 = 0

11001100 10101010 = 01100110

1111 0000 1010 0101 = 1100

Storage Architectures

Storage Architectures

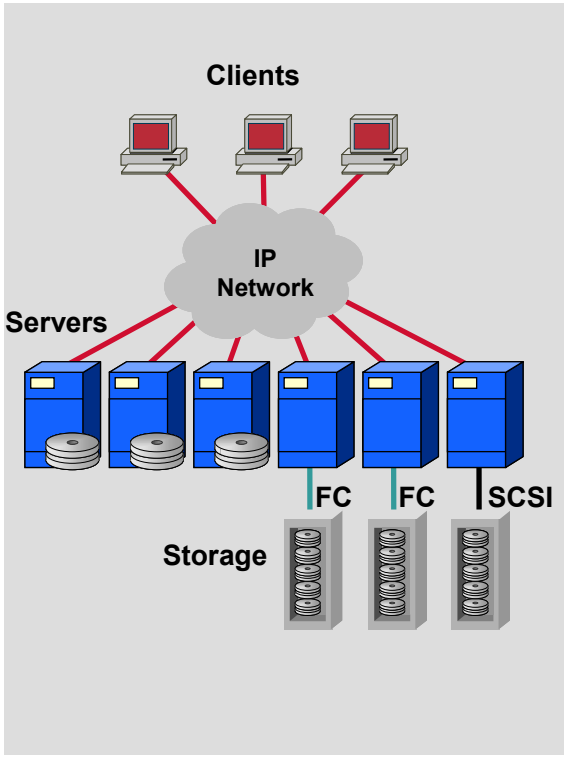

Direct-Attached Storage (DAS)

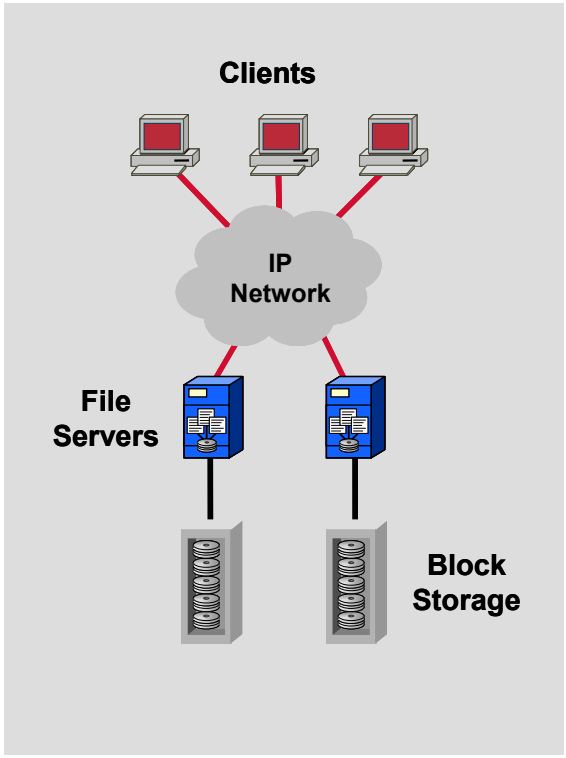

Network-Attached Storage (NAS)

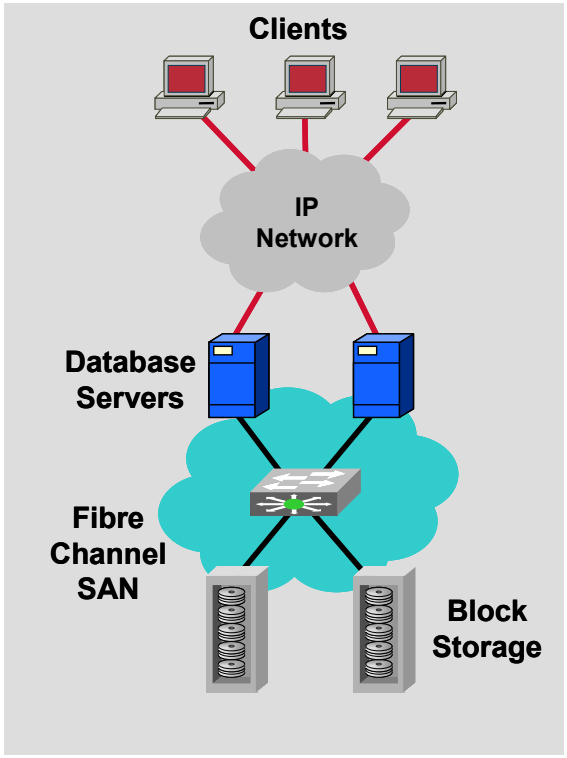

Storage Area Network (SAN)

What is DAS?

- DAS = Direct Attached Storage
- DAS solutions provide:
  - Low-cost, slow to medium speed storage for use in home computers and small businesses
  - Can provide higher-cost, high-performance storage, which can be used in solutions needing high-speed access

DAS Architecture

- DAS uses an I/O Channel architecture, which resides between a computer (initiator) and the device (target) used to store its data.
- Storage device is only accessible by attached host computer.
- Block level access to data.

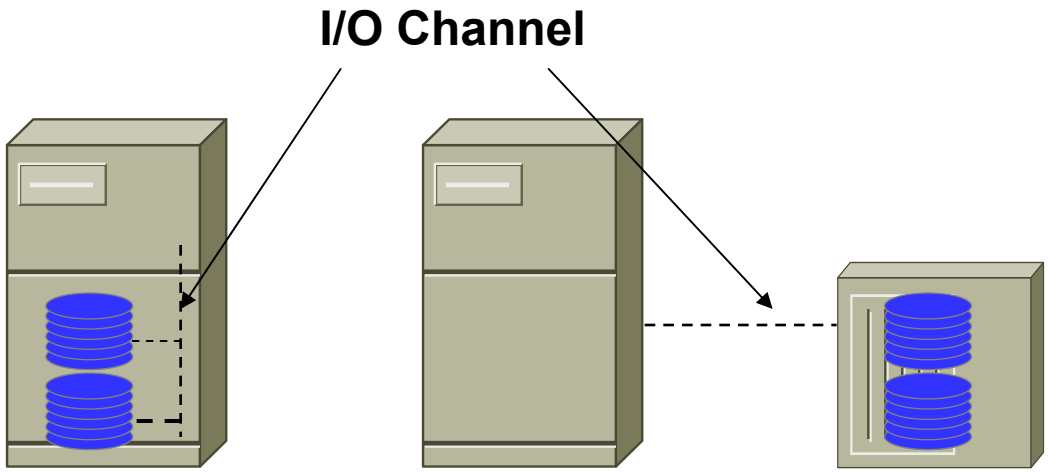

DAS Options

- The "typical" storage environment

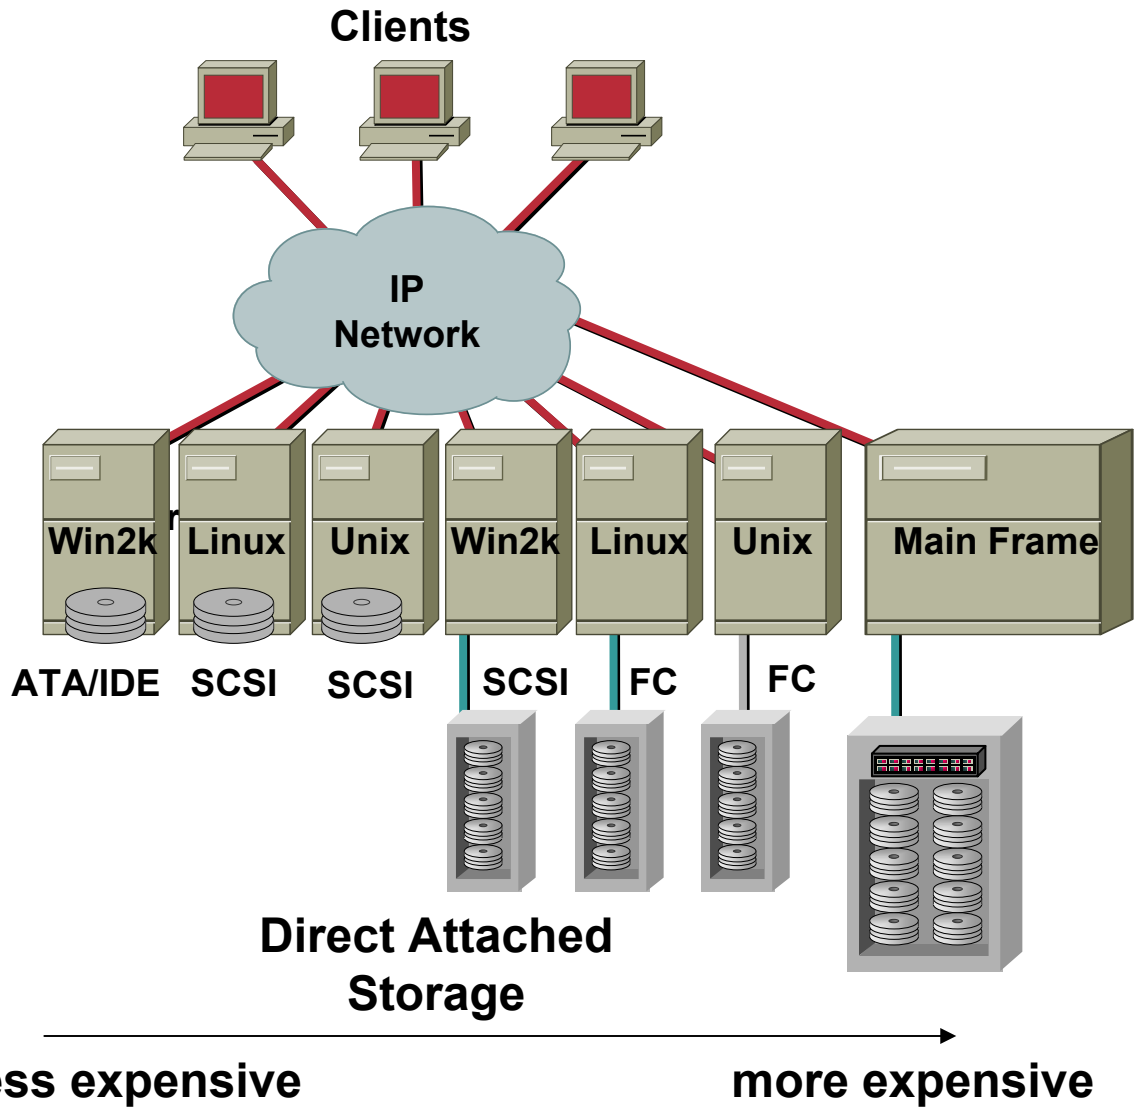

What is Network Attached Storage?

- NAS = Network Attached Storage
- NAS devices are network attached "appliances"
- NAS is the attachment of storage devices to the Local Area Network (LAN)

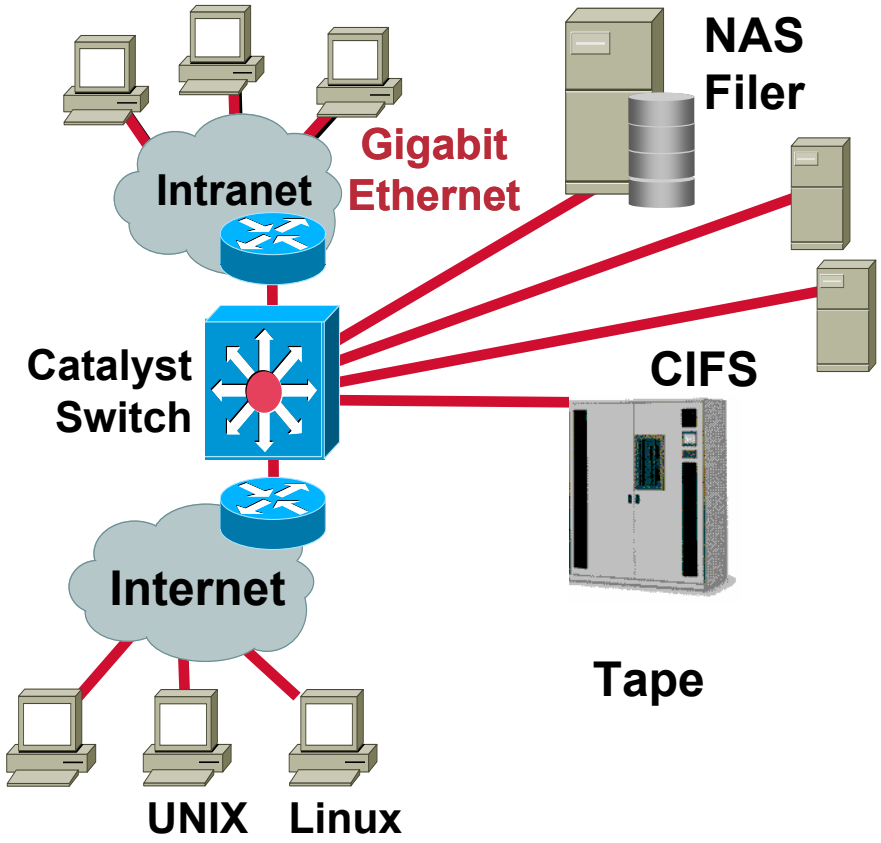

Example: DAS vs. NAS

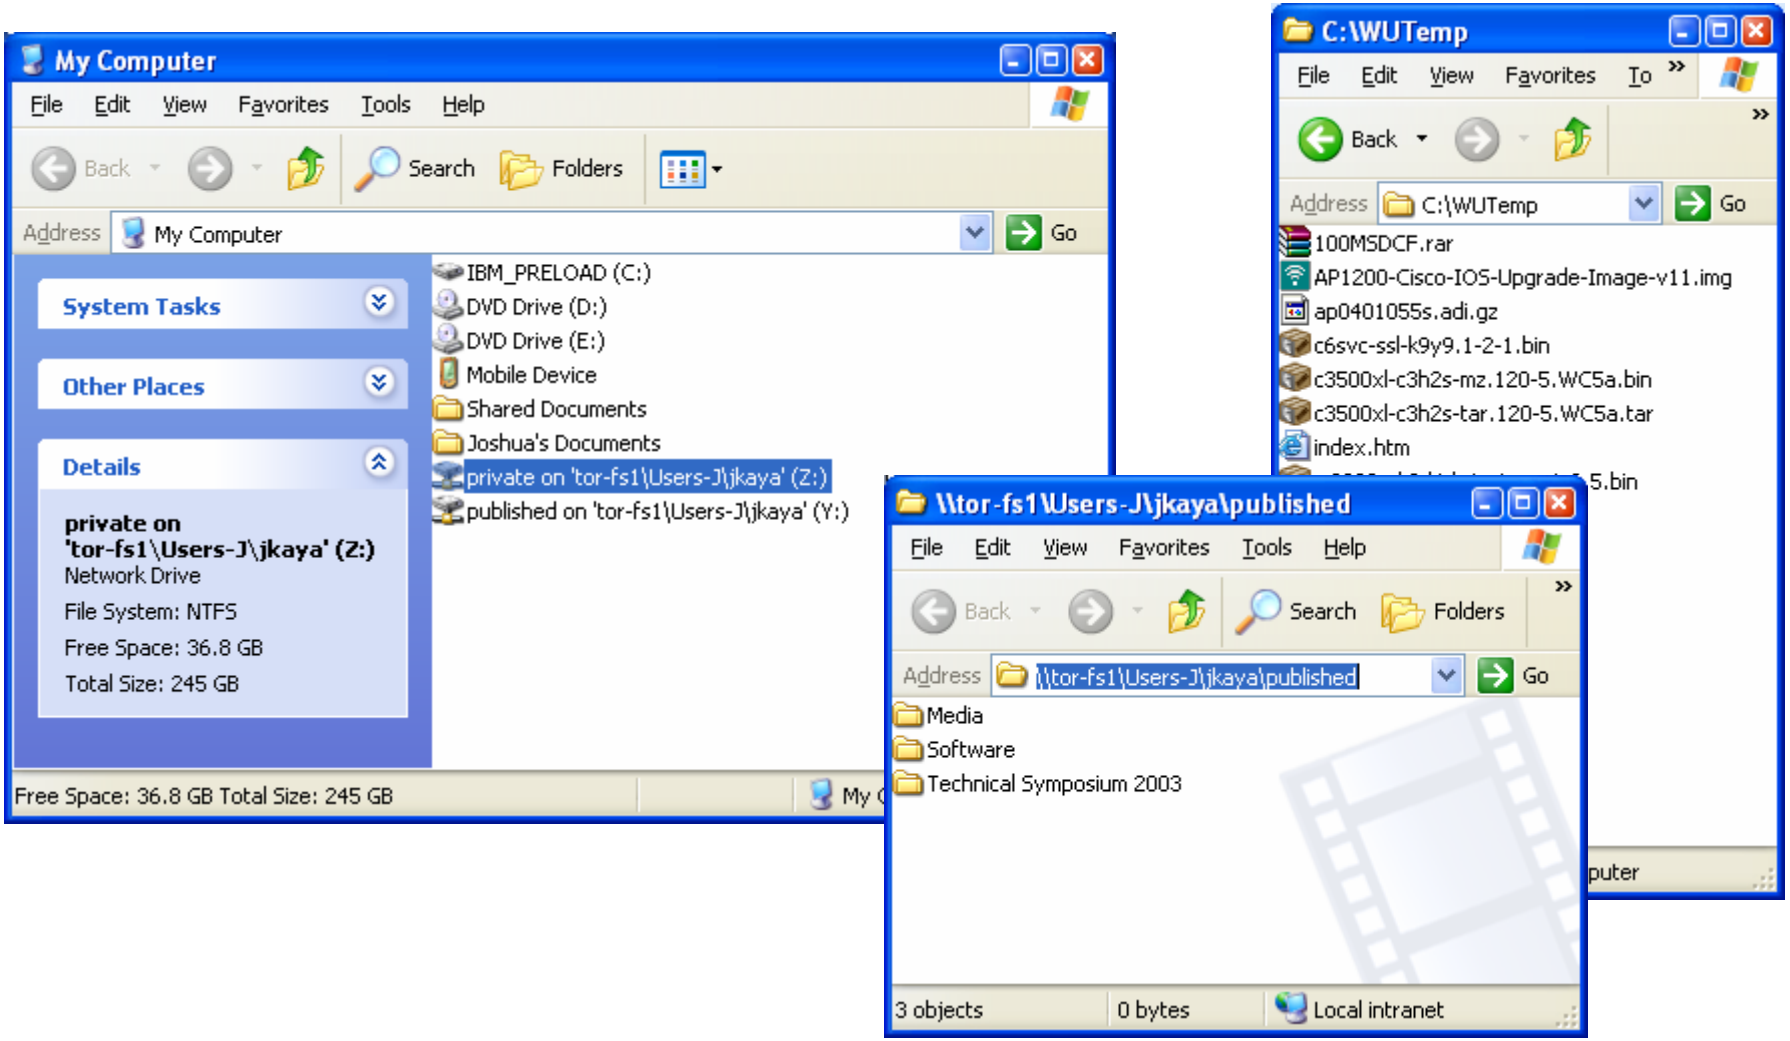

What is Fibre Channel?

Fibre Channel is a protocol for transporting data between devices.

Fibre Channel is the transport technology most commonly used for SANs today.

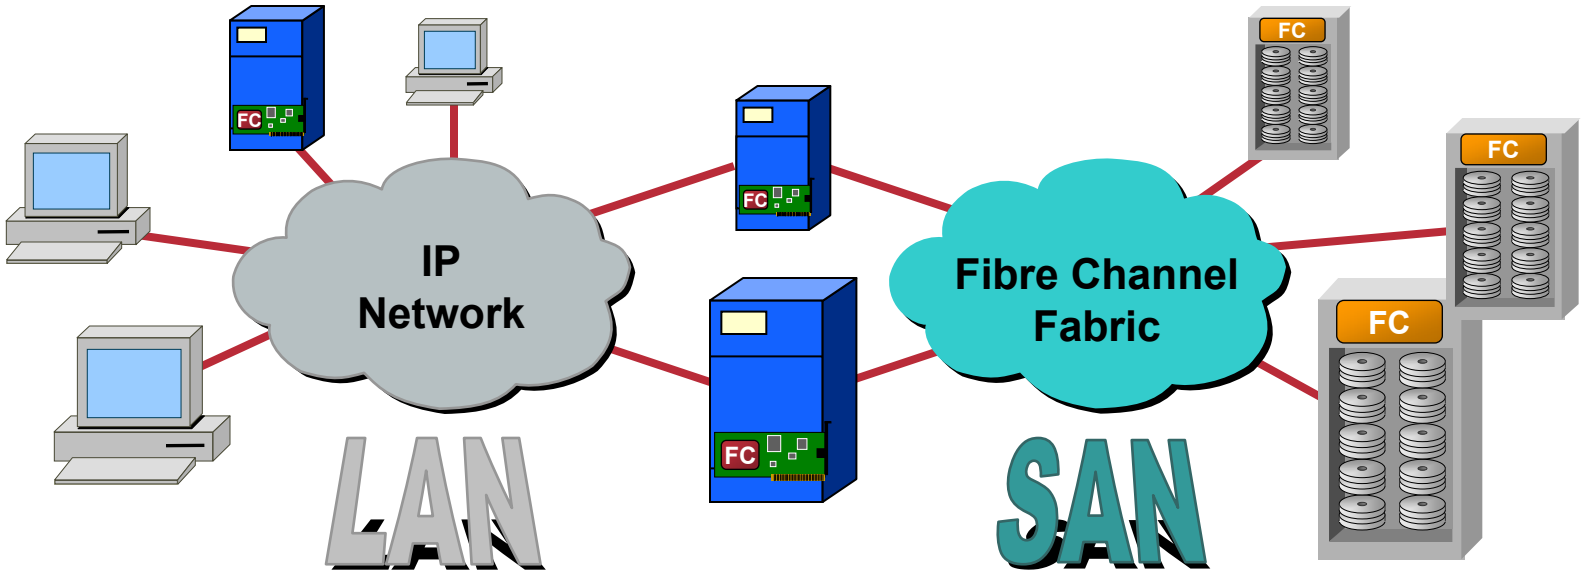

SAN Components

- **A SAN consists of:**
  - **Host systems with host bus adapters and drivers**
  - **An interconnection network with switches and hubs**
  - **Disk and tape storage subsystems**
  - **Block-level I/O protocols used to access the storage devices**

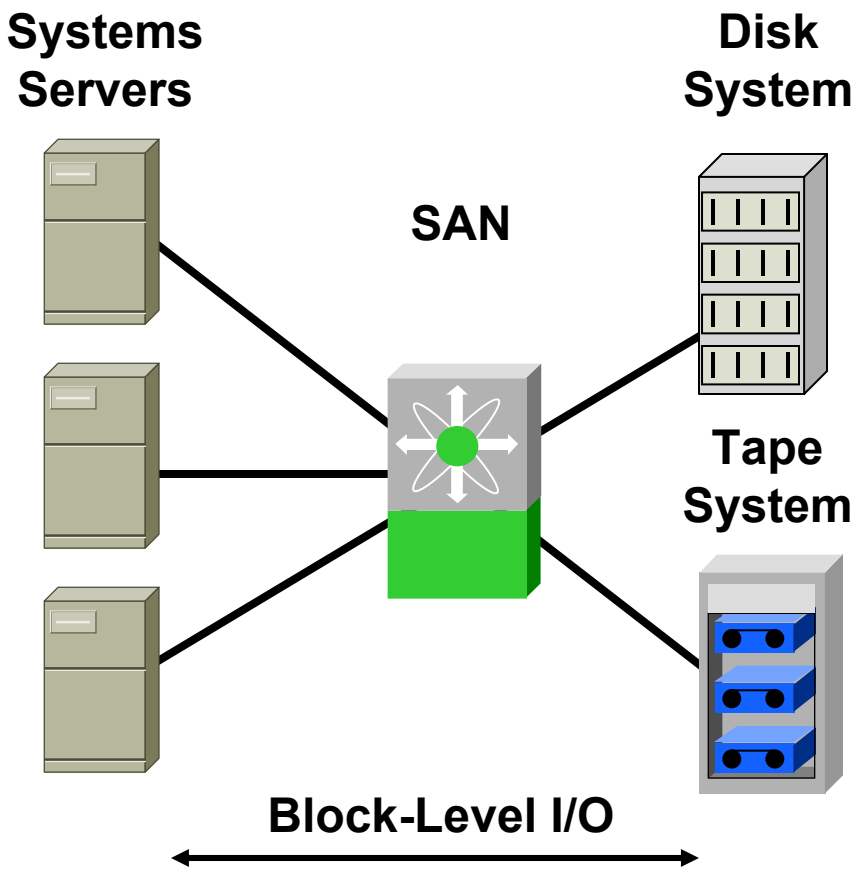

DHA1 change the diagrams

Duncan Anderson, 10/23/2002

DAS: I/O Channels

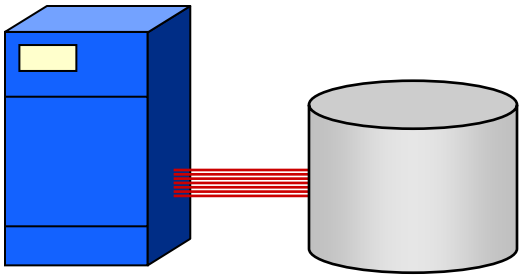

I/O Channel (SCSI)

- X Few devices
- X Static
- ✓ Low latency
- X Short distances
- ✓ Hardware-based delivery management

NAS: Data Networks

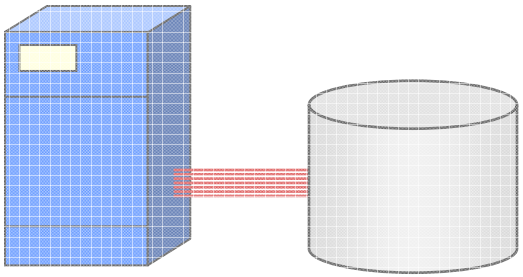

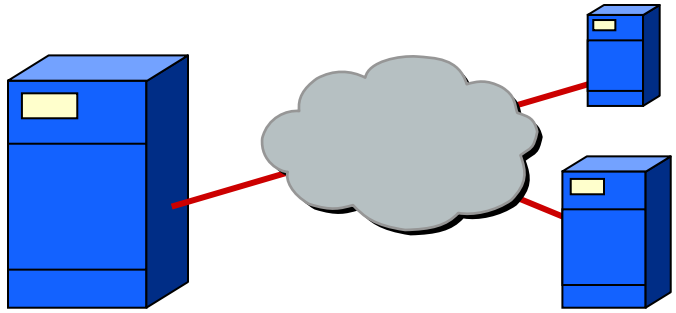

I/O Channel (SCSI)

- X Few devices
- X Static
- ✓ Low latency
- X Short distances
- ✓ Hardware-based delivery management

Network (Ethernet)

- ✓ Many devices
- ✓ Dynamic
- X High latency
- ✓ Long distances
- X Software-based delivery management

Fibre Channel: The Best of Both Worlds

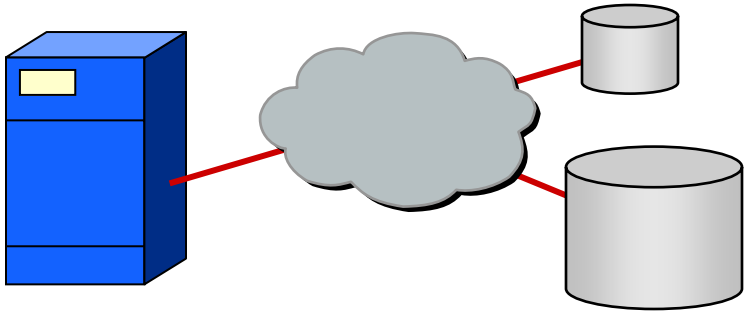

I/O Channel

- X Few devices
- X Static
- ✓ Low latency
- X Short distances
- ✓ Hardware-based delivery management

Fibre Channel

- ✓ Many devices
- ✓ Dynamic
- ✓ Low latency
- ✓ Long distances
- ✓ Hardware-based delivery management

Network

- ✓ Many devices
- ✓ Dynamic
- X High latency
- ✓ Long distances
- X Software-based delivery management

Grow storage and maintain current staffing levels

3-Year TCO by Storage Architecture
$ per Megabyte of User Data

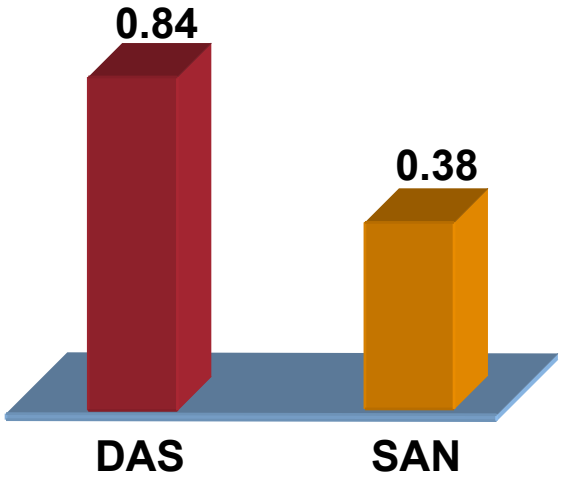

Source: Aggregate of Sources

3-Year TCO by Storage Architecture
Distribution of Cost

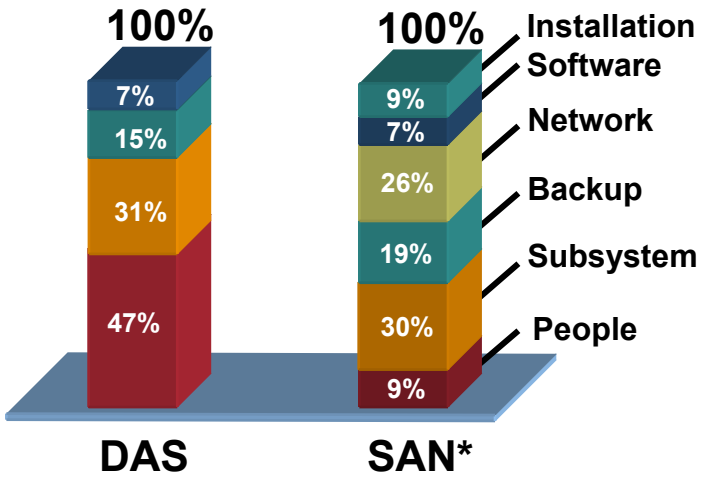

*** Large SAN (30 servers)**

Source: Aggregate of Sources

- **Networked Storage Cost Savings**
  - **Centralized Management** – 5-10x increase in storage managed per administrator
  - **Storage Utilization** – storage capacity is not stranded behind hosts that don't need it
  - **Independent scaling of computing and storage resources** – minimize capital expenditures
- **Other Operational Benefits**
  - **Availability and Performance** – allows multiple paths for redundancy and load sharing
  - **Disk to Disk Backup** – reduces congestion on LAN and reduces CPU cycles on hosts
  - **Billing / Reporting** – centralization enables enterprise-wide billing and reporting

SCSI

Introducing SCSI

- SCSI = Small Computer System Interface.
- SCSI is a *standard* that defines an interface between an initiator (usually a computer) and a target (usually a storage device such as a hard disk).
- *Interface* refers to connectors, cables, electrical signals, optical signals and the command protocol that allow initiators and targets to communicate.

SCSI Example

In this case, a file is being written to the hard drive by an application on the workstation.

The SCSI command protocol is used to communicate between SCSI devices.

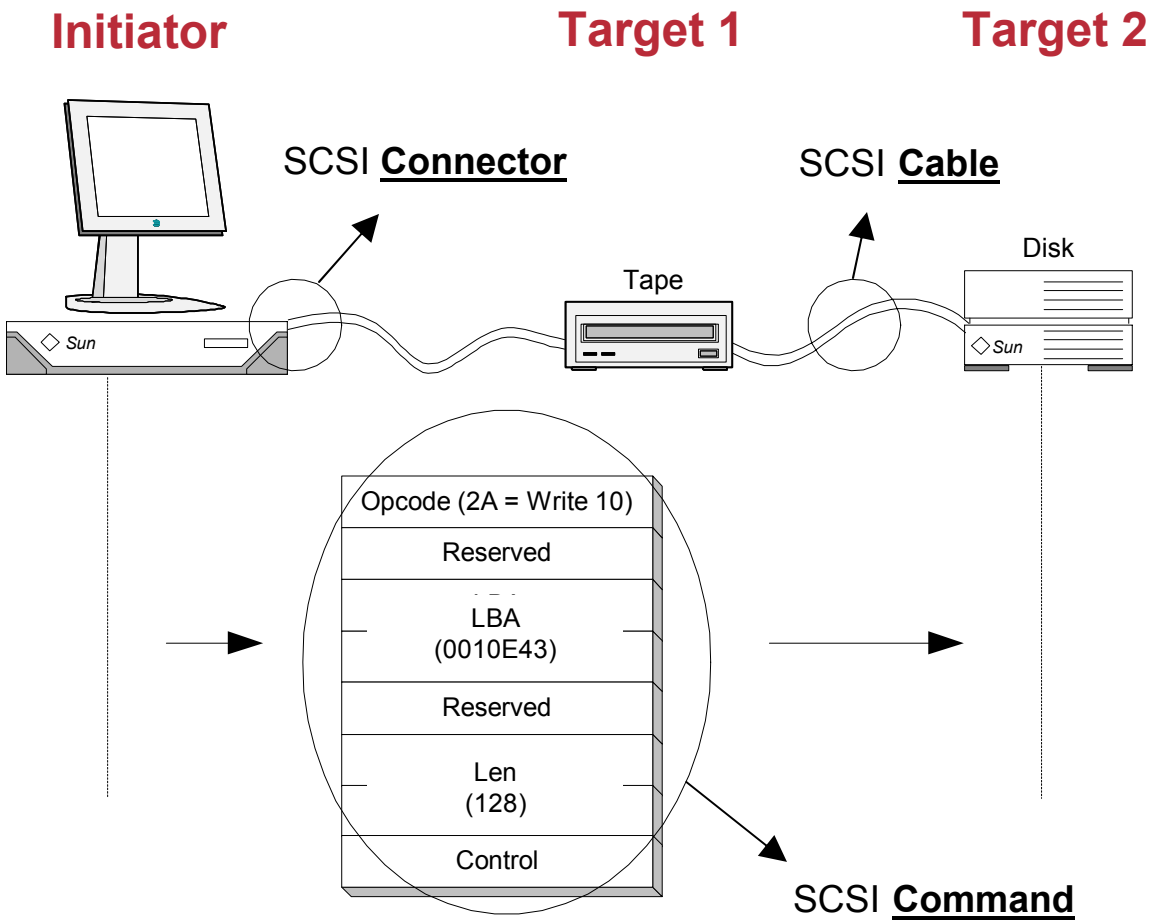

Standards

- SCSI has evolved since it was introduced as SASI in 1979 by Shugart Associates – it was approved as a standard by ANSI in 1986 and is now referred to as SCSI-1.
- SCSI-2 was approved by X3 in 1990 and by ANSI in 1994.
- SCSI-3 refers to a *collection* of standards, each of which defines a very specific part of SCSI: physical interface, transport interface, command interface, architecture model, programming interface, etc.

Sample SCSI Standard Components

SCSI Primary Commands (SPC-2)

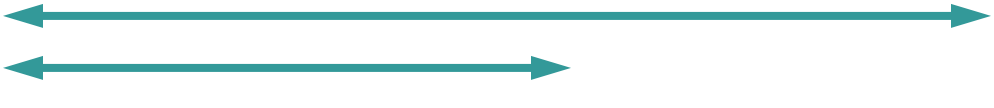

Initiator Target 1 Target 2

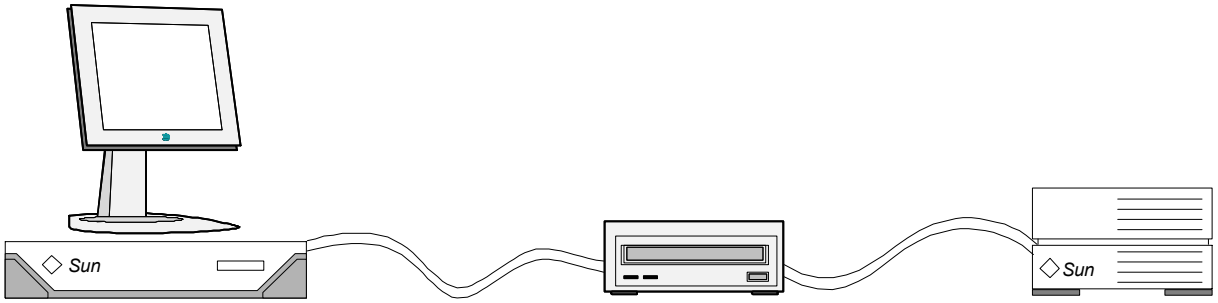

SCSI Block Commands (SBC)

SCSI Stream Commands (SSC)

End to End Interaction

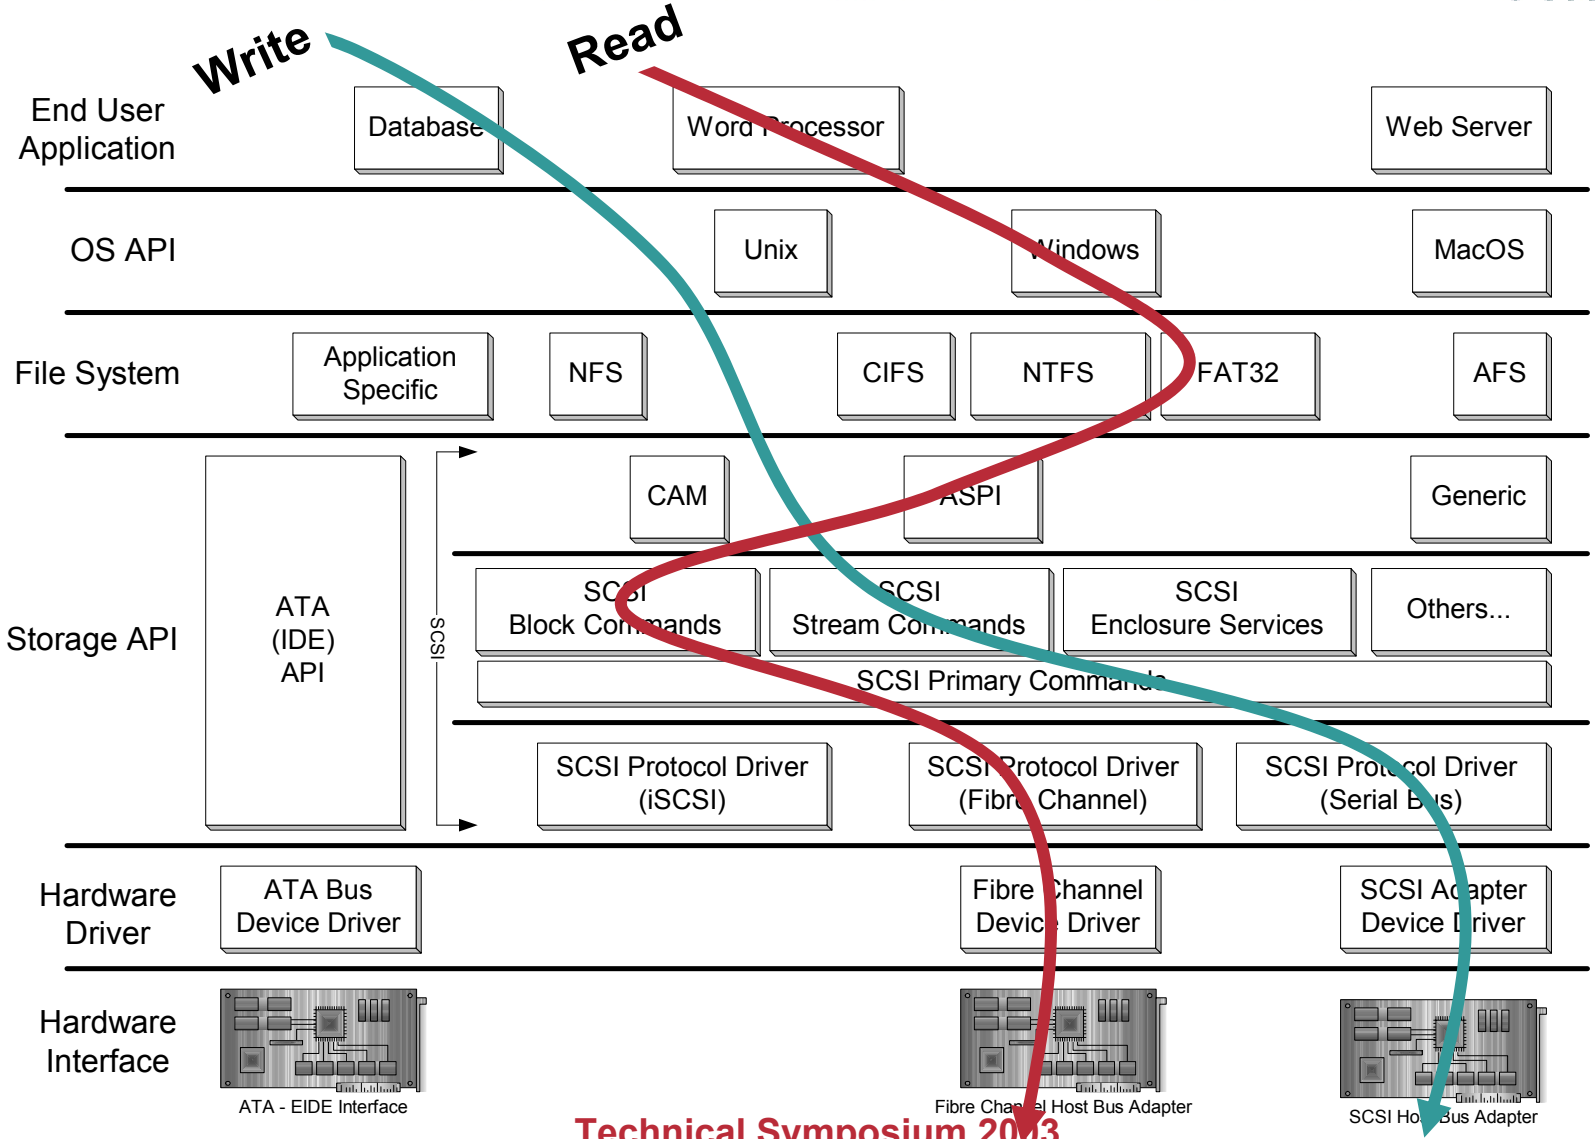

SCSI-3 Standards:

- **SCSI-3 Parallel Interface (SPI)**

  Defines the mechanical, timing, phases, and electrical parameters of the parallel cable we all know and love. Some of the electrical and cable parameters are tightened/improved over SCSI-2.

- **SCSI-3 Interlock Protocol (SIP)**

  Defines the messages and how the phases are invoked. No real change from SCSI-2, except for some new messages.

- **SCSI-3 Architectural Model (SAM)**

  In a nutshell, defines a common set of functions and services and definitions for how a physical transport properly gets commands, data, and status exchanged between two devices, complete with error handling and queueing.

- **SCSI-3 Primary Commands (SPC)**

  All of the commands executed by any and all SCSI devices, like REQUEST SENSE and INQUIRY, etc.

- **SCSI-3 Block Commands (SBC)**

  Disk commands.

- **SCSI-3 Stream Commands (SBC)**

  Tape commands.

- **SCSI-3 Controller Commands (SCC)**

  RAID box commands.

- **SCSI-3 Multimedia Commands (MMC)**

  For CDROMS etc.

- **SCSI-3 Fibre Channel Protocol (FCP)**

  SCSI commands over gigabit Fibre Channel.

- **SCSI-3 Serial Bus Protocol (SBP)**

  SCSI commands over IEEE 1394 High Speed Serial Bus (Apple's "Firewire").

- **SCSI-3 Serial Storage Protocol (SSP)**

  SCSI commands over SSA.

Sample SCSI Standard Components (1)

Cisco.com

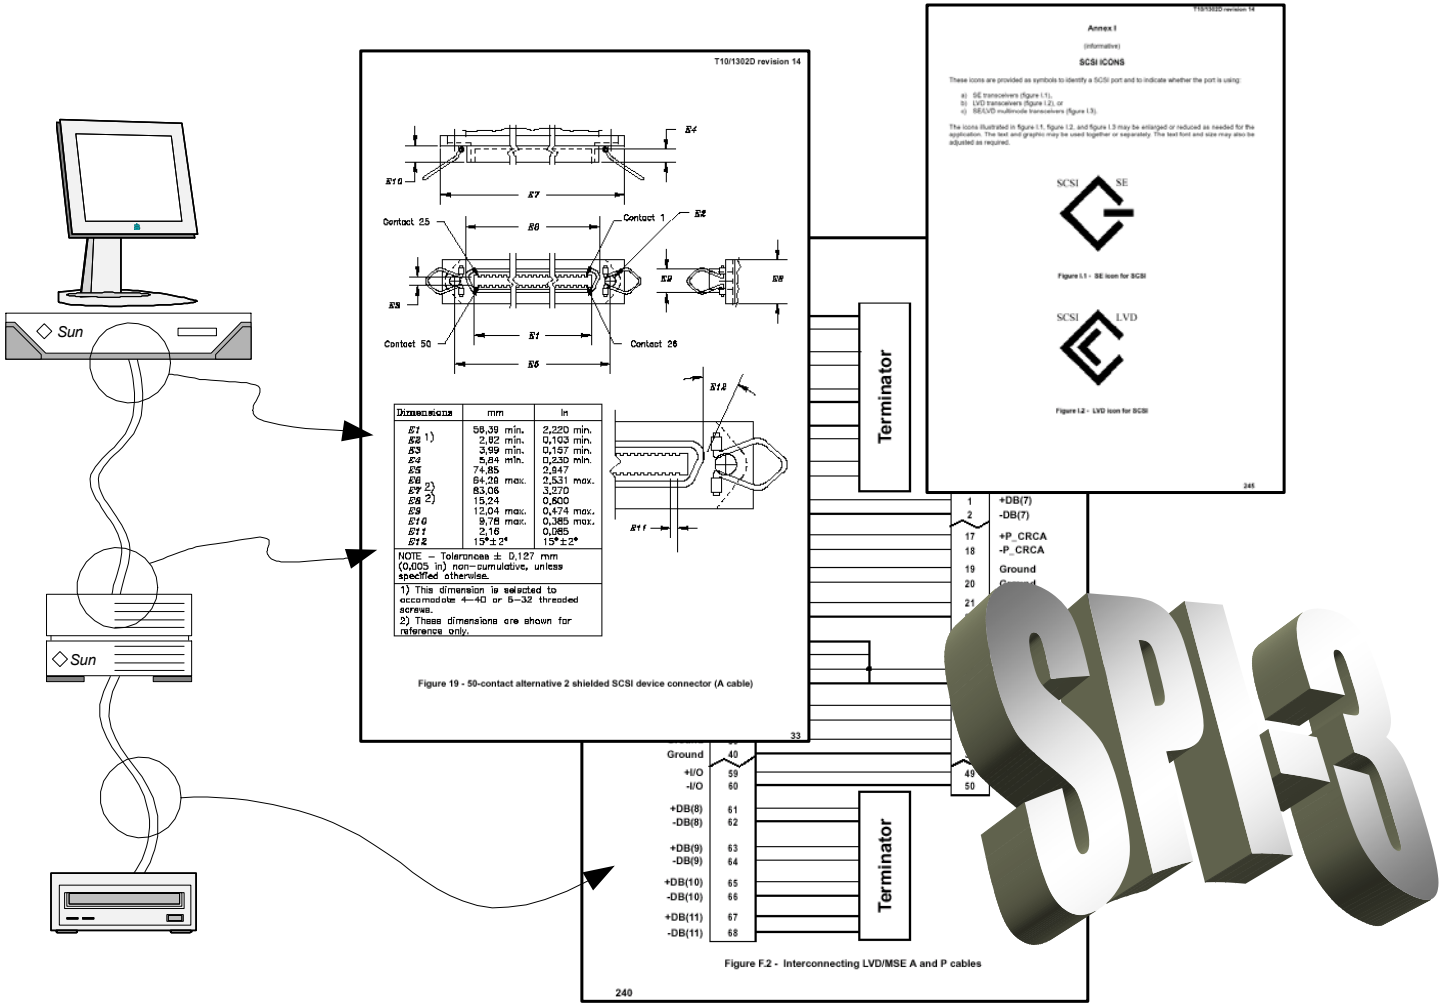

SCSI: Connectors (Now SCSI-3 SPI)

SCSI I DB-25 Male	SCSI I DB-25 Female	SCSI I CENTRONICS 50 Male	SCSI I CENTRONICS 50 Female
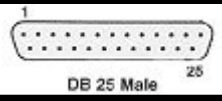	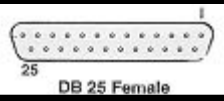		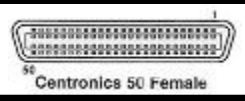
SCSI II HP-50 Male	**SCSI II** HP50 Female	**SCSI II** IDC 50 Male	**SCSI II** IDC 50 Female
	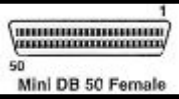	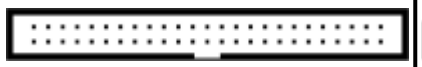	
SCSI III HP-68 Male	**SCSI III** HP68 Female	**SCSI III** SCA 80 Male	**SCSI III** SCA 80 Female
	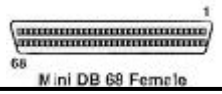	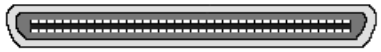	

What the difference between Fast, Ultra, and Wide?

- SCSI-1 (Async) - 5Mhz, 8-bit, 4MB/s, 7 devices, 50-pin.
- SCSI-1 (Sync) - 5Mhz, 8-bit, 5MB/s, 7 devices, 50-pin.
- SCSI-2 (Wide) - 5Mhz, 16-bit, 10MB/s, 15 devices, 68-pin.
- SCSI-2 (Fast) - 10Mhz, 8-bit, 10MB/s, 7 devices, 50-pin.
- SCSI-2 (Fast/Wide) - 10Mhz, 16-bit, 20MB/s, 15 devices, 68-pin.
- SCSI-3 (Ultra) - 20Mhz, 8-bit, 20MB/s, 7 devices, 50-pin.
- SCSI-3 (Ultra/Wide) - 20Mhz, 16-bit, 40MB/s, 7 devices, 68-pin.
- SCSI-3 (Ultra2) - 40Mhz, 8-bit, 40MB/s, 7 devices, 50-pin.
- SCSI-3 (Ultra2/Wide) - 40Mhz, 16-bit, 80MB/s, 15 devices, 68-pin.
- SCSI-3 (Ultra3(Ultra160)) - 40Mhz, 16-bit, 160MB/s, 15 devices, 68-pin.

SCSI: Common Parallel Interfaces

Ultra2 SCSI:

- Up to 80MBps
- Up to 15 devices
- 68 pin interface
- Up to 12 meters

Ultra160 SCSI:

- Up to 120MBps
- Up to 15 devices
- 68 pin interface
- Up to 12 meters

Parallel Communications: Review: 5Mhz * 8bits = 4MBps

- SCSI-1 (Async) - 5Mhz, 8-bit, 4MB/s, 7 devices, 50-pin.
- SCSI-1 (Sync) - 5Mhz, 8-bit, 5MB/s, 7 devices, 50-pin.
- SCSI-2 (Wide) - 5Mhz, 16-bit, 10MB/s, 15 devices, 68-pin.
- SCSI-2 (Fast) - 10Mhz, 8-bit, 10MB/s, 7 devices, 50-pin.
- SCSI-2 (Fast/Wide) - 10Mhz, 16-bit, 20MB/s, 15 devices, 68-pin.
- SCSI-3 (Ultra) - 20Mhz, 8-bit, 20MB/s, 7 devices, 50-pin.
- SCSI-3 (Ultra/Wide) - 20Mhz, 16-bit, 40MB/s, 7 devices, 68-pin.
- SCSI-3 (Ultra2) - 40Mhz, 8-bit, 40MB/s, 7 devices, 50-pin.
- SCSI-3 (Ultra2/Wide) - 40Mhz, 16-bit, 80MB/s, 15 devices, 68-pin.
- SCSI-3 (Ultra3(Ultra160)) - 40Mhz, 16-bit, 160MB/s, 15 devices, 68-pin.

Serial Communications: Fibre Channel

Fibre Channel 1Gbps:

- ~100MBps
- 100's of devices (per HBA)
- Optical Interface (LC Connector)
- ~10km

Fibre Channel 2Gbps:

- ~200MBps
- 100's of devices (per HBA)
- Optical Interface (LC Connector)
- ~10km

Fibre Channel Protocol

FCP: Fibre Channel Protocol

- FCP defines how SCSI-3 commands are transported over a Fibre Channel network
- Transport involves more than encapsulation
- Individual SCSI commands are mapped to FC sequences
- The set of SCSI commands that form a SCSI transaction are mapped to a FC exchange
- Fibre Channel preserves the frame order

SCSI as a FC Upper Layer Protocol

Fibre Channel Hierarchy

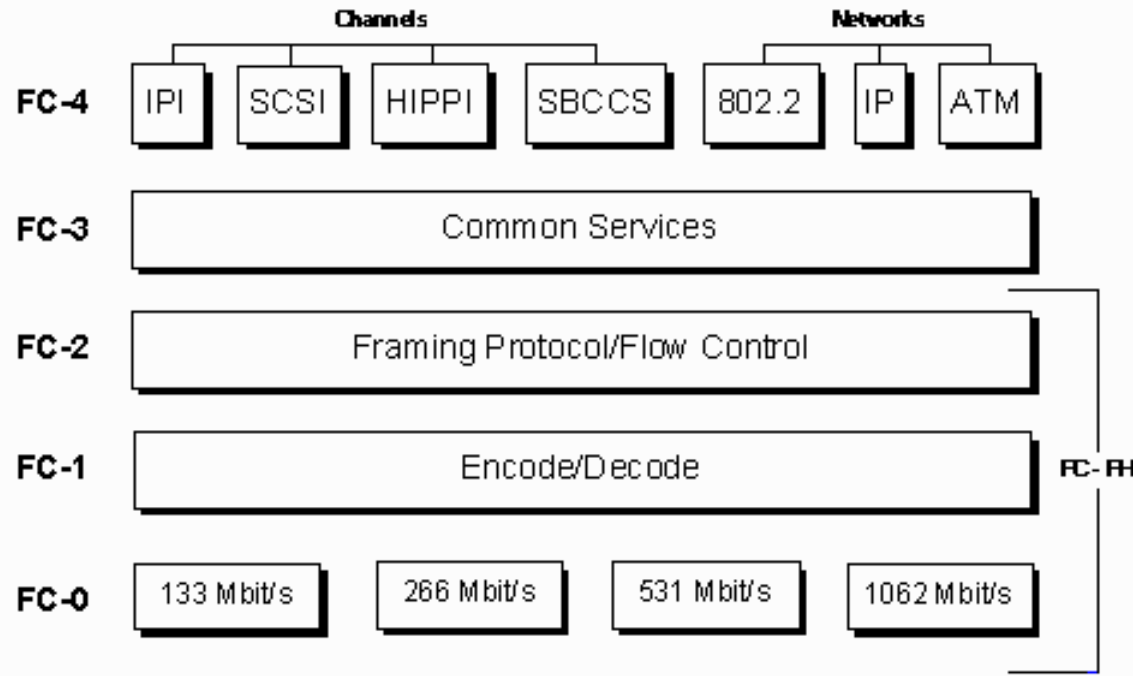

SCSI Protocol Encapsulation

- Parallel SCSI

- Fibre Channel Protocol

Fibre Channel Architecture Model

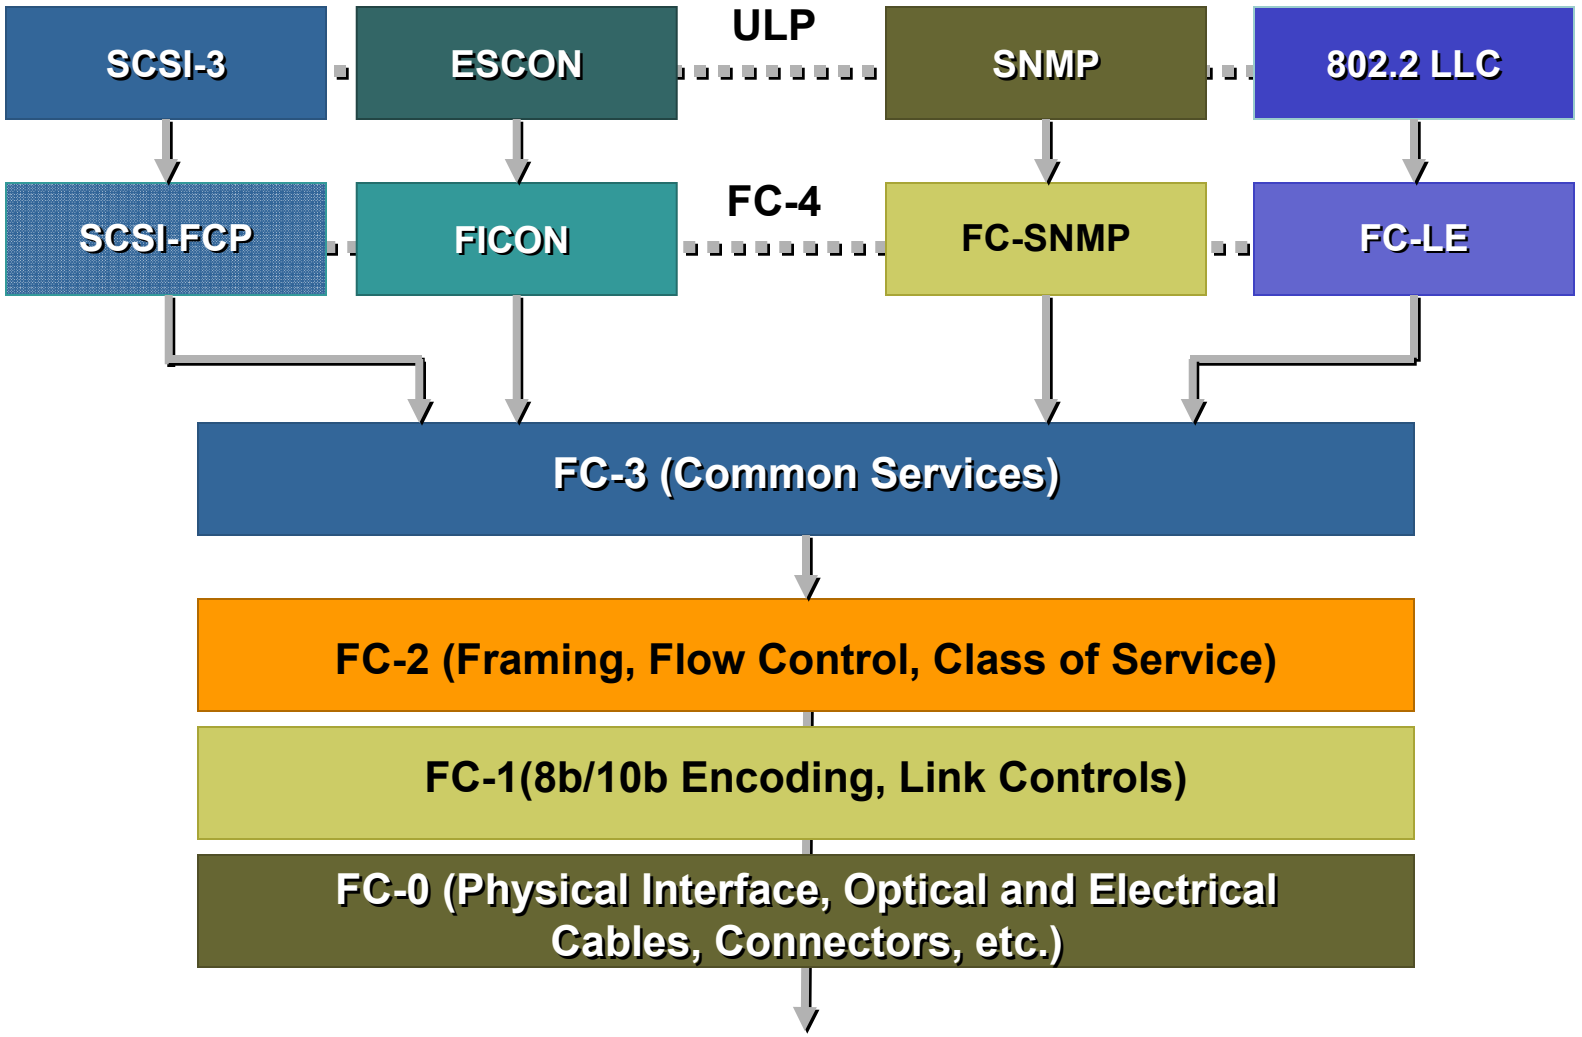

Fibre Channel Topologies

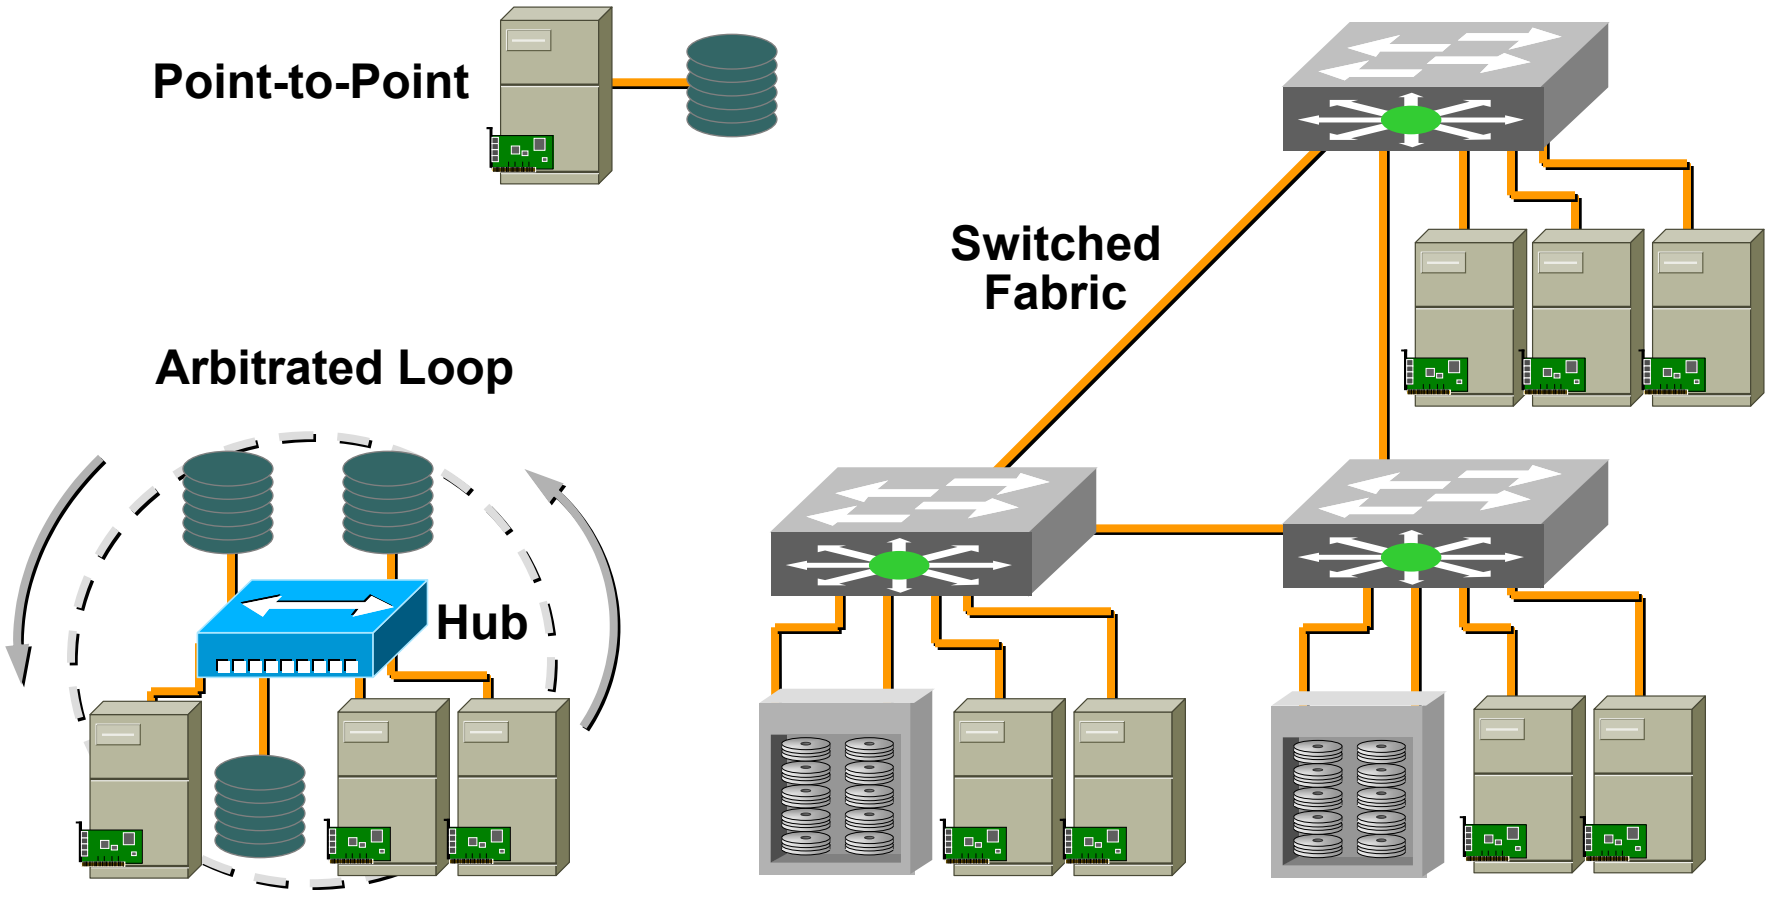

FC Port Types

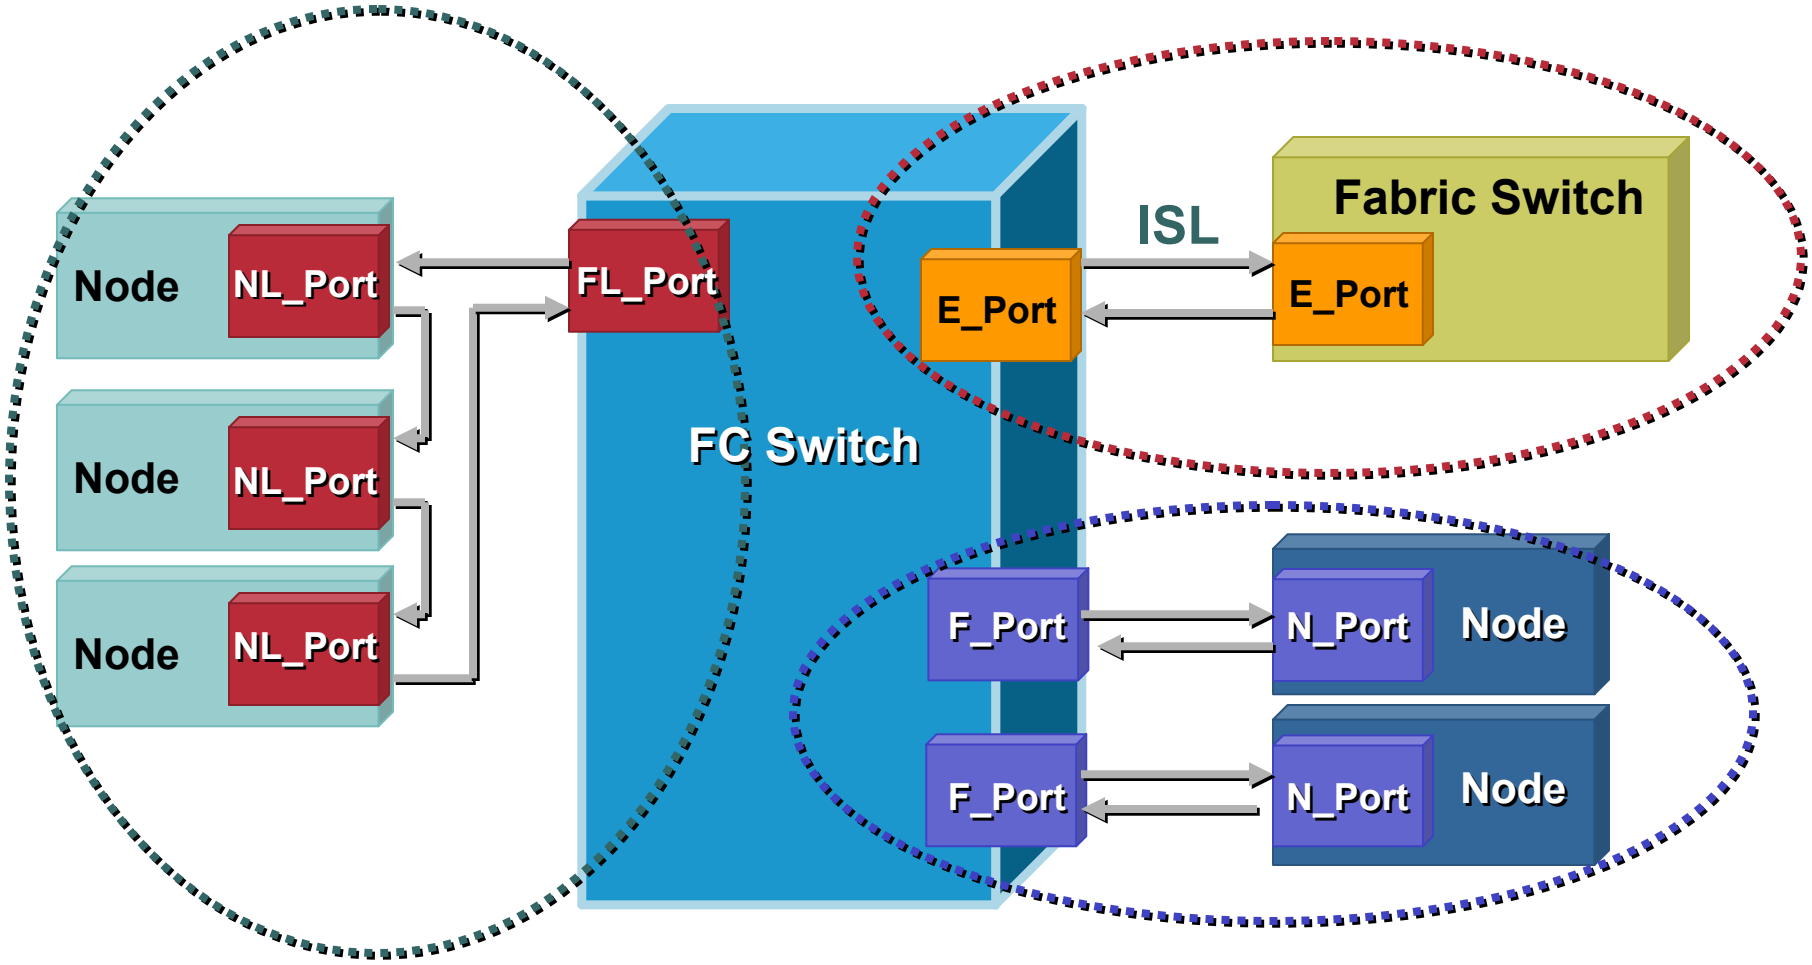

Addressing

- IP/Ethernet
  - Each node has
    - a pre-defined 48-bit MAC Address
    - a user-configurable 32-bit IP Address
  - IP Addresses *can* be assigned by DHCP
- Fibre Channel
  - Each node has
    - a pre-defined 64-bit WWN (World Wide Name)
    - a 24-bit FC_ID
  - FC_IDs *are* assigned to nodes during Fabric Login

Fibre Channel ID Format

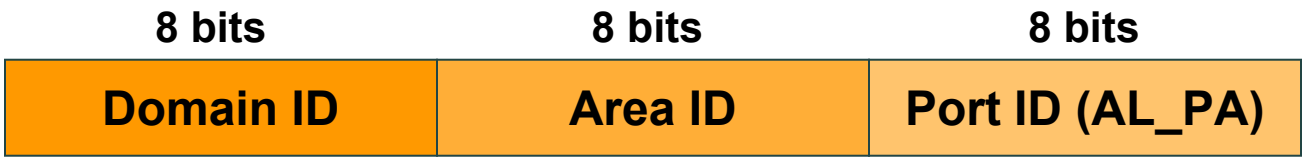

- **Domain ID**

  Identifies the switch

- **Area ID**

  Identifies different loops connected to the same switch

- **Port ID (or AL_PA)**

  Identifies the port on the switch (for N_Ports) or the specific node on the loop (for NL_Ports)

Fibre Channel Scalability Limits (In Theory)

- 239 Switches per Fabric

  256 available Domain IDs minus 17 reserved

- 256 Loops per Switch
- 256 Ports per Switch
- 128 Nodes per Loop

  Not all of the AL_PA values can be used on a loop

Address Resolution

- **IP/Ethernet**
  - **IP to MAC Addresses using ARP (Broadcast)**
  - **Names to IP Addresses using DNS (DB Query)**
- **Fibre Channel**
  - **WWNs never used to explicitly address a node (similar to IP names)**
  - **FC_IDs discovered using FC DNS (DB Query)**
  - **Broadcast is never used, except for IP-FC**

Flow Control

- **Ethernet**
  - Point-to-point Flow Control is optional
    - ON/OFF protocol (e.g. IEEE 802.3x)
  - End-to-end flow control is delegated to transport protocols (e.g. TCP)
- **Fibre Channel**
  - Point-to-point Flow Control is mandatory
    - Implemented by BB_Credit
  - End-to-end Flow Control is optional
    - Implemented by EE_Credit

Flow Control and Credits

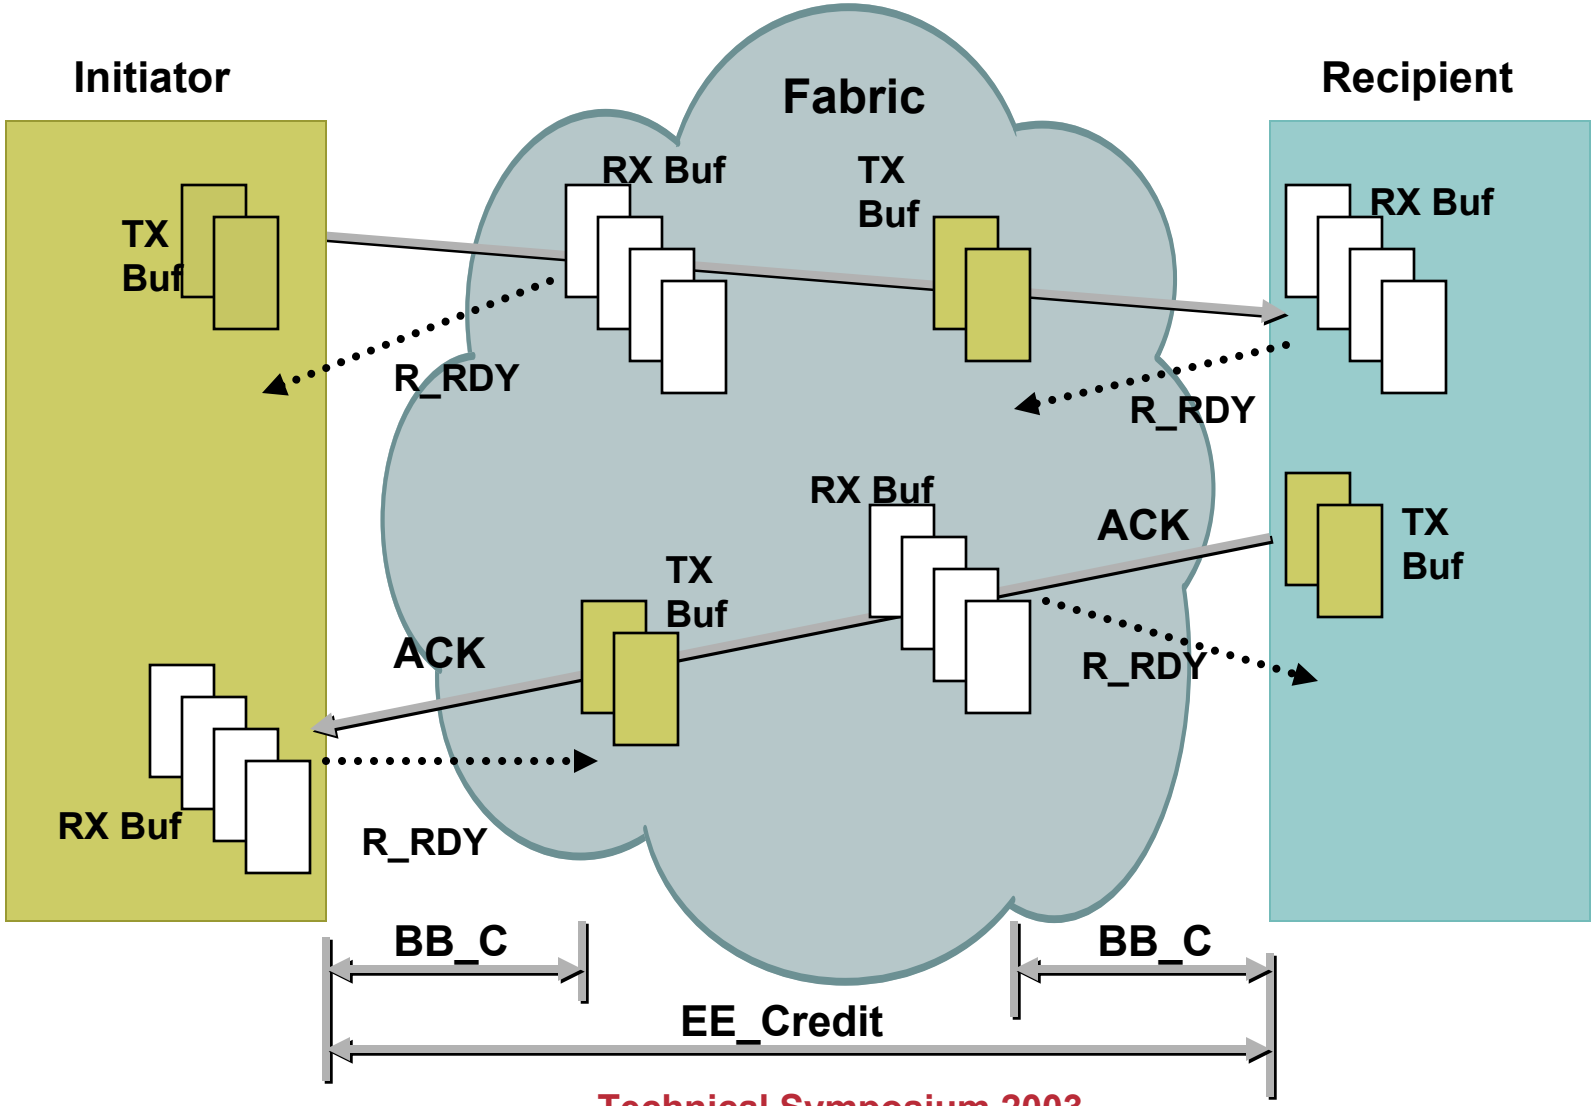

How Many Credits?

	Layer 1 Mbit/s	Data Link Layer Mbit/s	Mbyte/s	
FC 1Gb/s	1,062.50	850.00	106.25	
FC 2Gb/s	2,125.00	1,700.00	212.50	
FC 10Gb/s	10,518.75	10,200.00	1,275.00	Base-R

1Gb/s FC

Size (Bytes) Payload	Frame	Inter Frame	Frame/s	Frame Length meters	feet	Frame/Km	Minimum Credit/Km	
0	36	24	1,770,833	113	371	8.85	18.00	ACK
2048	2,084	24	50,403	3,968	13,018	0.25	1.00	Standard SCSI Payload

2Gb/s FC

Size (Bytes) Payload	Frame	Inter Frame	Frame/s	Frame Length meters	feet	Frame/Km	Minimum Credit/Km	
0	36	24	3,541,667	56	185	17.71	36.00	ACK
2048	2,084	24	100,806	1,984	6,509	0.50	2.00	Standard SCSI Payload

10Gb/s FC

Size (Bytes) Payload	Frame	Inter Frame	Frame/s	Frame Length meters	feet	Frame/Km	Minimum Credit/Km	
0	36	8	28,977,273	7	23	144.89	290.00	ACK
2048	2,084	8	609,465	328	1,077	3.05	7.00	Standard SCSI Payload

How Much Memory?

- **In Ethernet**
  - **Proportional to the round trip delay**
  - **Depends on speed**
  - **At 1Gb, 8ms delay ➔ 80Mb/port, i.e. 10MB/port)**
- **In Fibre Channel**
  - **As many buffers as credits**
  - **Independent on speed**
  - **At 1 or 2Gb, 32 Credits, 2048 byte/maxsize_frame ➔ 64KB/port**

Packet Loss

- **In IP/Ethernet**
  - **It's part of the game!**
  - **Used by TCP/IP to handle congestions**
- **In SCSI/Fibre Channel**
  - **Will throw you out of the market!**
  - **There is no congestion control!**

Congestion Example in a FC Network

- B talks with C
- A talks with D
- C is congested
- A suffers for the congestion of C

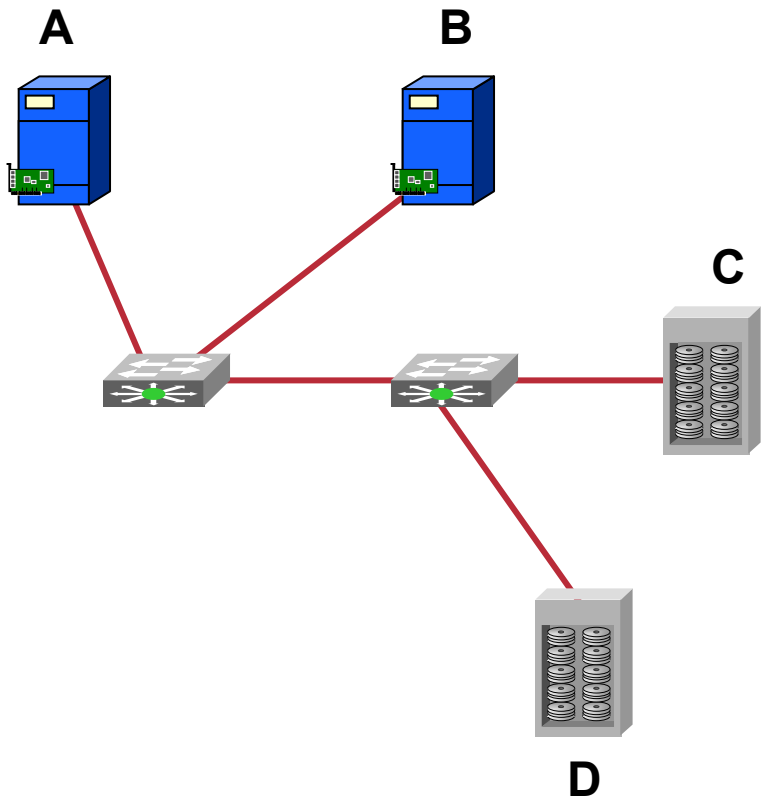

Congestion Example in a TCP/IP Network

- B talks with C
- A talks with D
- C is congested
- Packets from B to C are dropped
- B reduces the window
- A does NOT suffer for the congestion of C

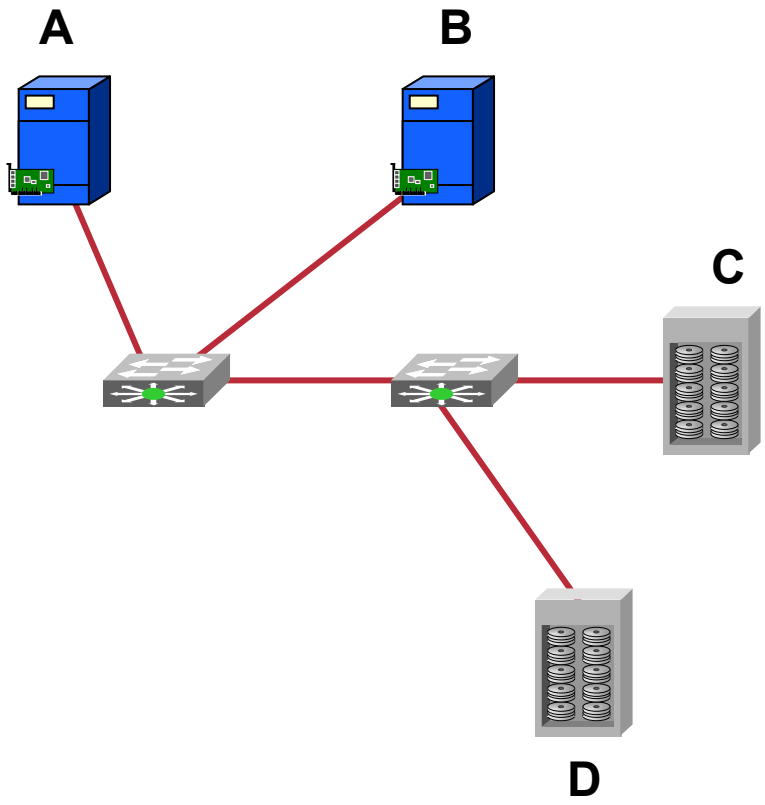

Routing

- Ethernet
  - Spanning Tree to define an active topology
  - Local forwarding decisions in each switch
  - More sophisticated routing delegated to network layer protocols (e.g. OSPF, BGP)
- Fibre Channel
  - FSPF to build routing tables into switches
    - FC_ID Allocation simplifies routing tables
  - Multiple paths used at the same time
  - No TTL (Time To Live)
    - Potential loops may be present during convergence
  - There is no such a thing as a network layer

Scaling to Large Networks

- **Ethernet/IP**

  Designed for The Internet

  Multiple levels of hierarchy allows to scale to larger networks

    VLANs, Subnets, Administrative Domains,…

- **FC/SCSI**

  Designed for I/O Channels

  Scalability requirements limited to few devices connected to a single bus or inside a Data Center

FC Framing and Sequencing

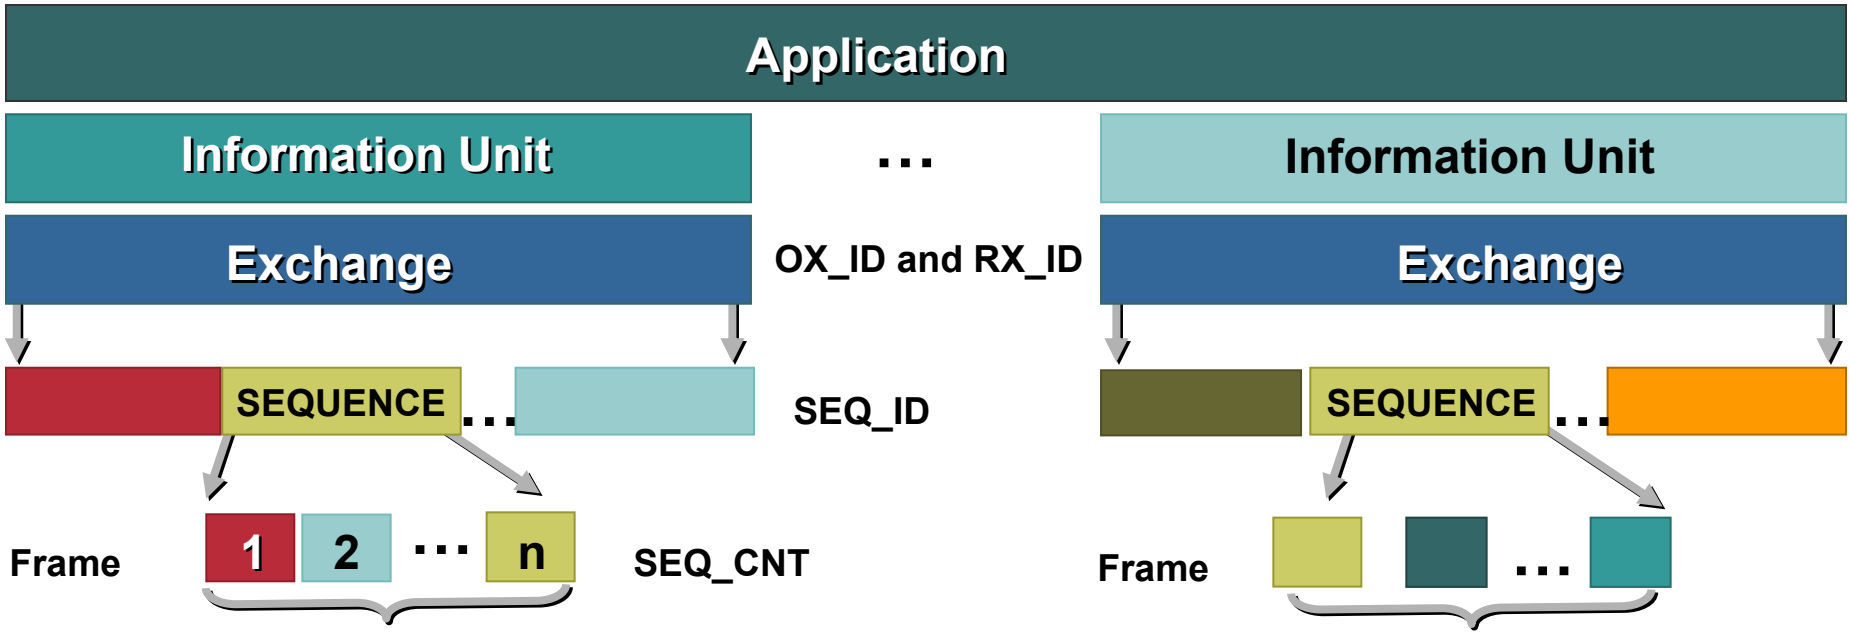

FC Frame Format

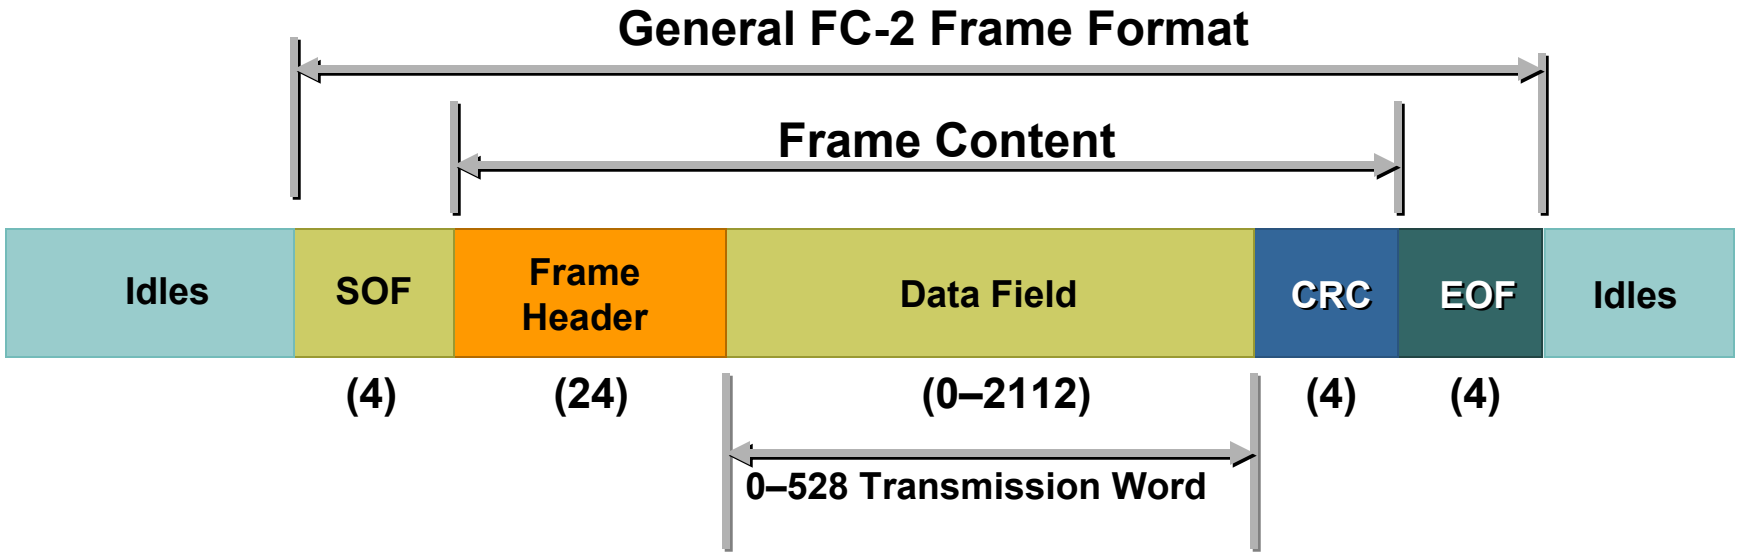

Data Field

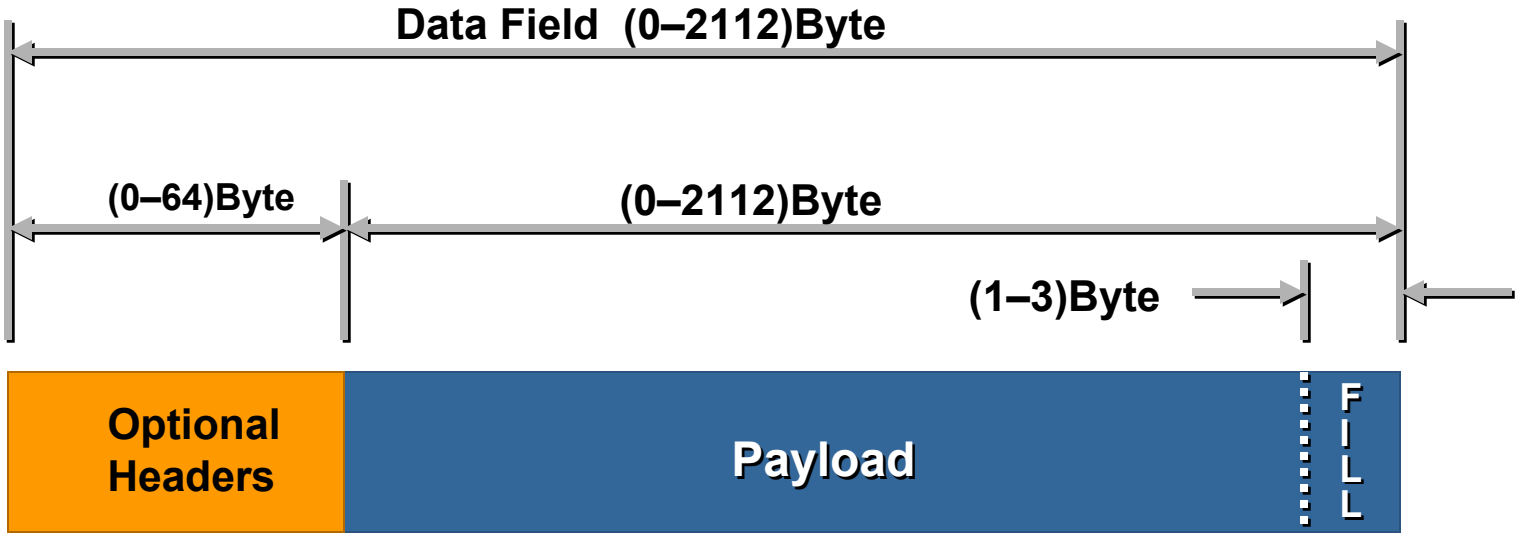

Typical MTU 0–2048

Minimum 128 Bytes

Fibre Channel Session Management

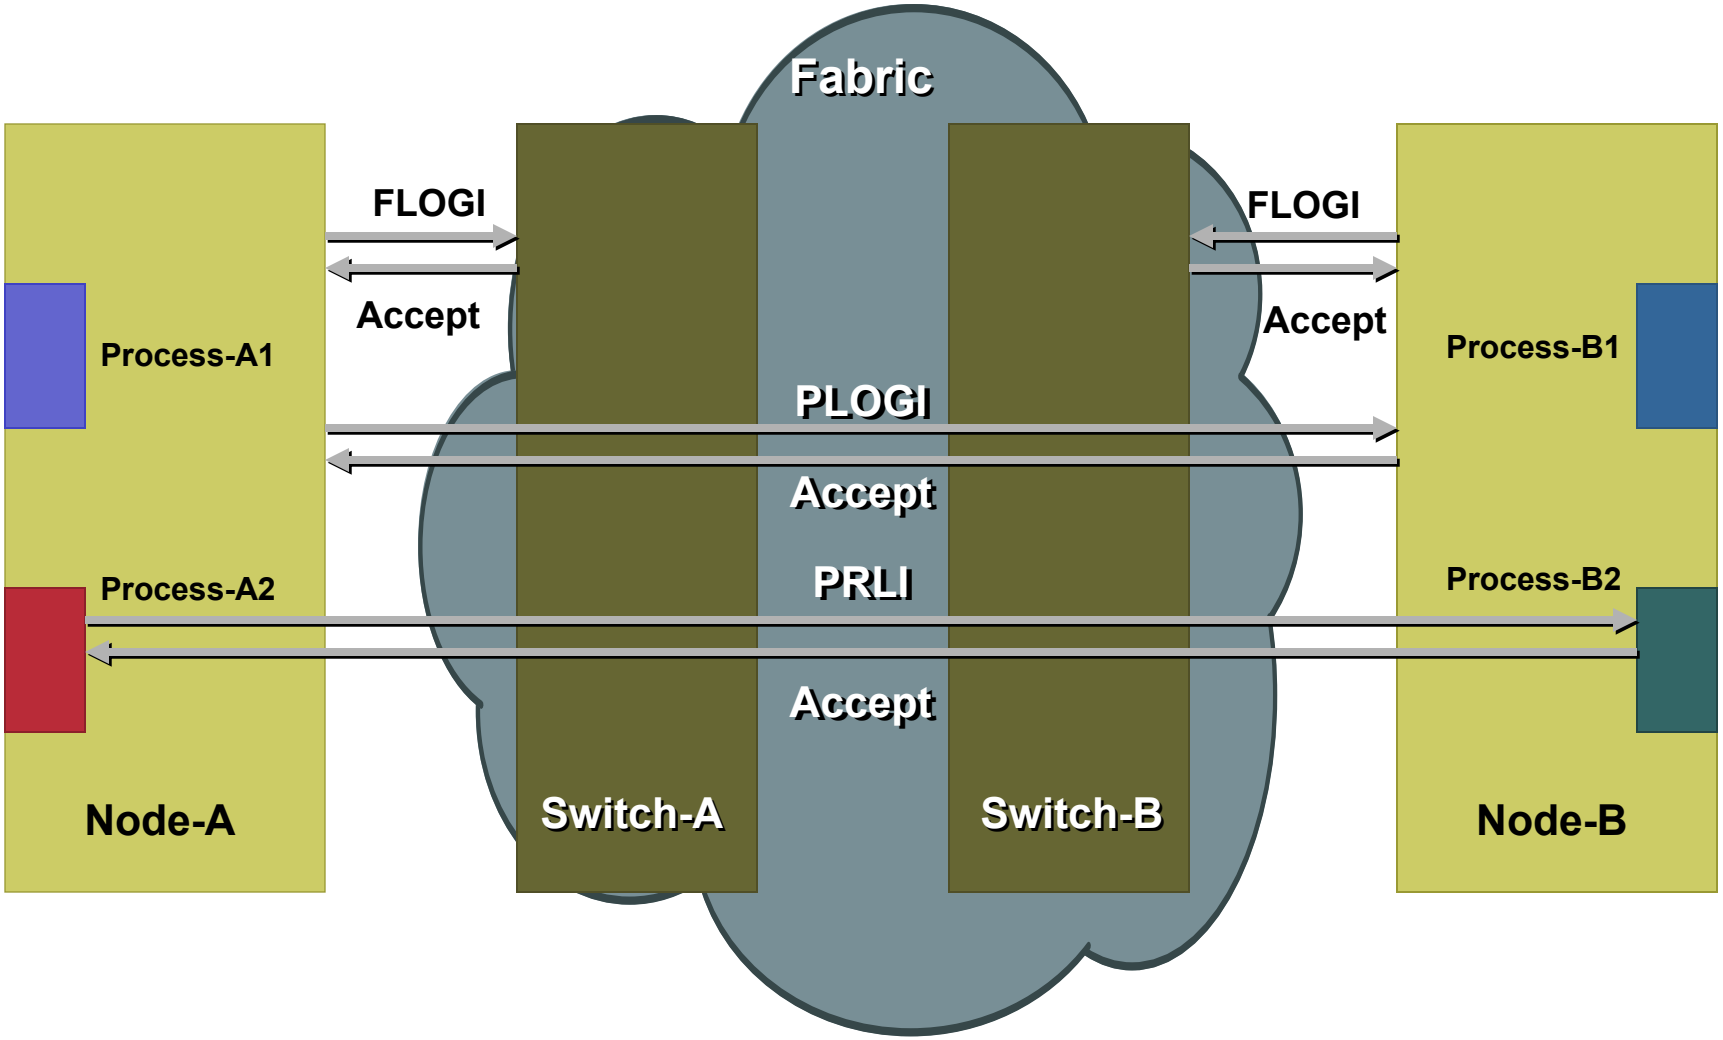

Zoning and ACLs

- **IP/Ethernet**

  ACLs to limit the capability of nodes to talk to each other

  VLANs and IP Subnets to segment broadcast domains and increase network security, scalability and availability

- **In Fibre Channel**

  Zoning to define who CAN talk to whom

  Nothing to define who CAN NOT talk to whom

  Nothing like VLANs or Subnets to increase security, scalability or availability of a Fabric (at least so far ☺)

FCIP & iSCSI

IP Protocol Encapsulation

- IP SANs carry block I/O traffic on top of IP

  Leverage Gigabit Ethernet performance for local traffic

  Use TCP: A reliable transport for delivery in MAN/WANs

- Two primary protocols:

  iSCSI—"IP-SCSI" IP-native transport of SCSI CDBs and data within TCP/IP connections

  FCIP—"Fibre-Channel-over-IP"— Tunneling of Fibre Channel frames within TCP/IP connections, including FC fabric management frames

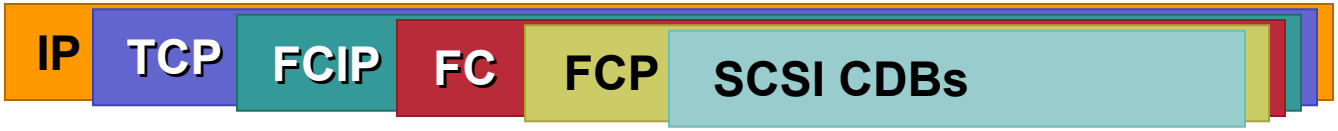

What is FCIP (Fibre Channel over IP)?

IT creates one logical fabric between remote SANS, and the switches think they are connected. IP is only used for tunneling through the WAN.

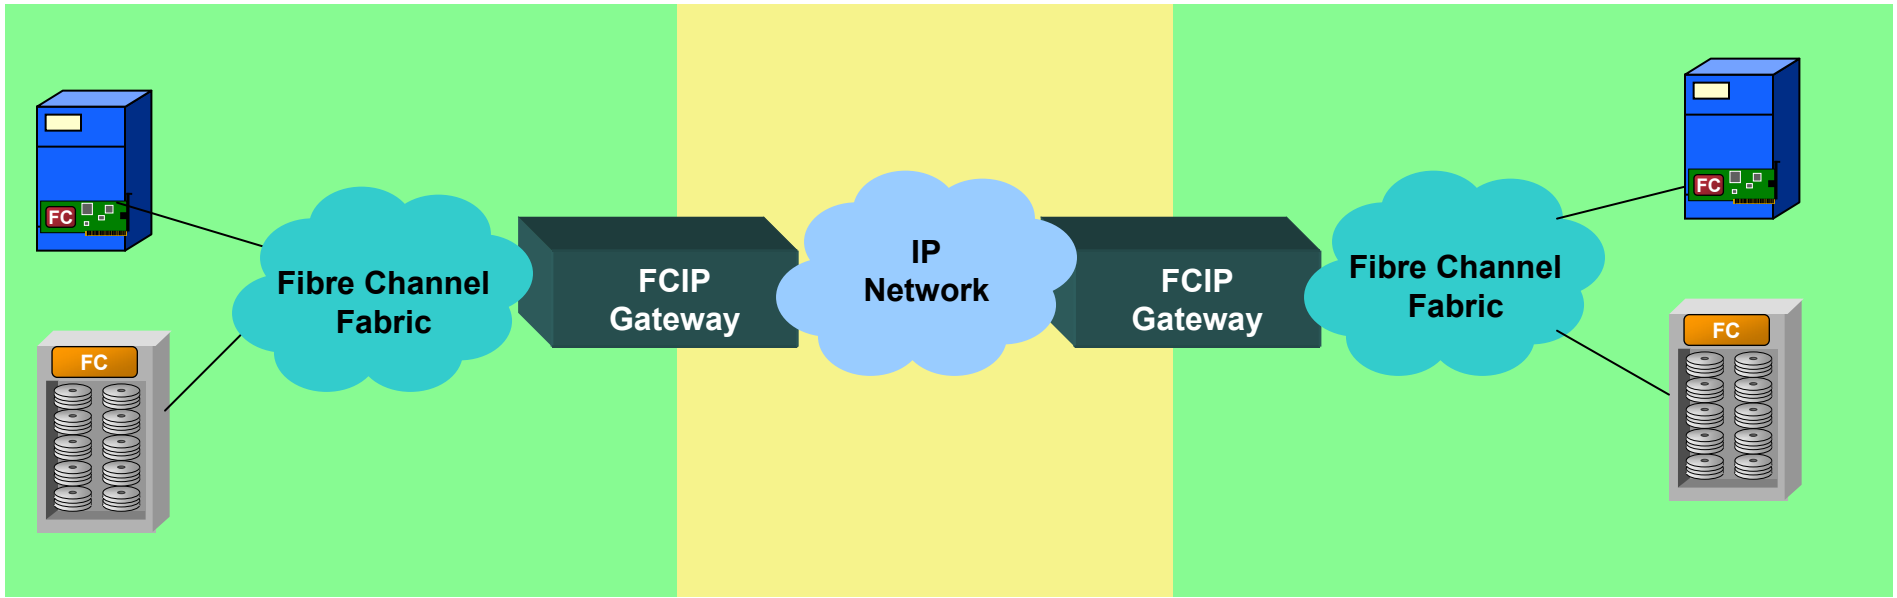

What is FCIP (Fibre Channel over IP) cont…

Remote FC resources are viewed as local

FCIP creates a Virtual FC Inter-Switch Link (ISL)

Fabric service information is extended across the FCIP ISLs

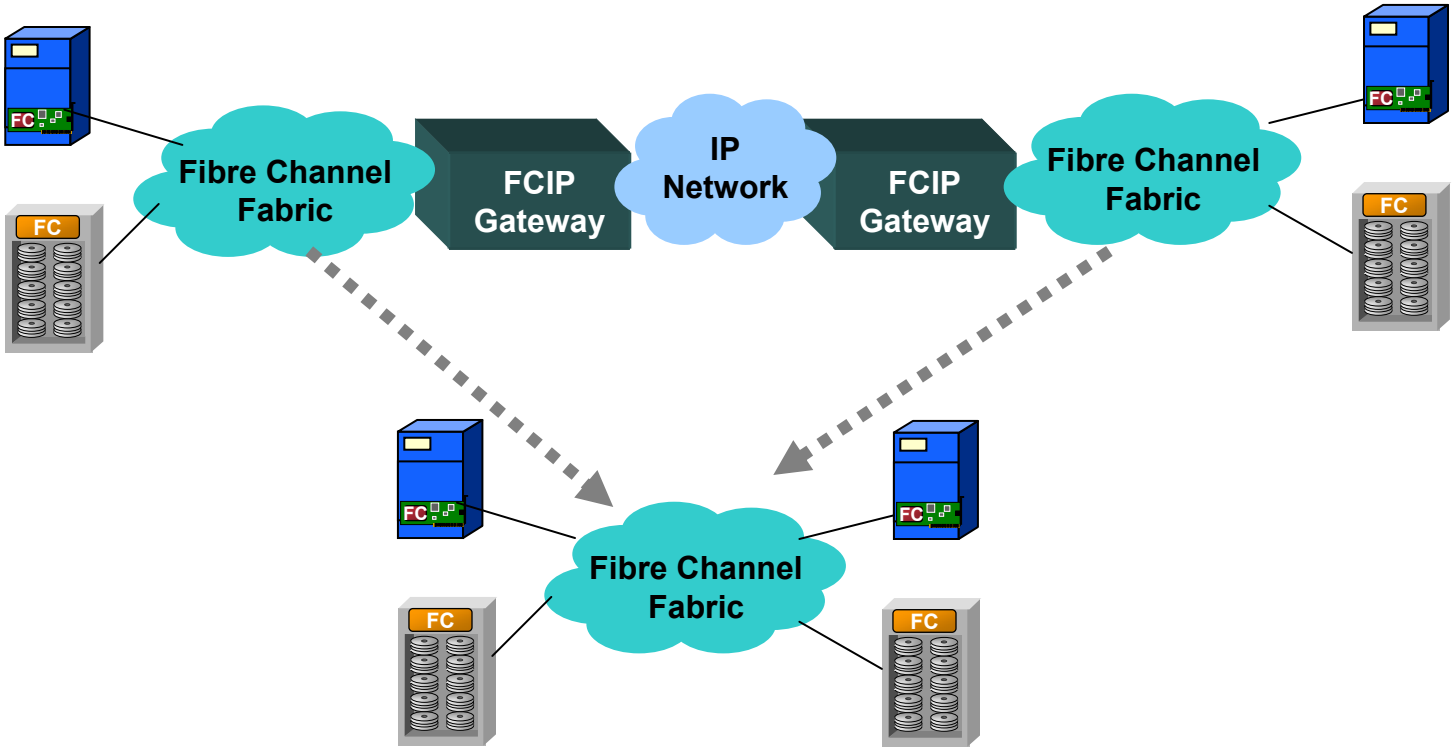

An FCIP Application Topology

FCIP Gateways perform Fibre Channel encapsulation process into IP Packets and reverse that process at the other end

FC Switches connect to the FCIP gateways through an E_Port for SAN fabric extension to remote location

A tunnel connection is set up through the existing IP network routers and switches across LAN/WAN/MAN

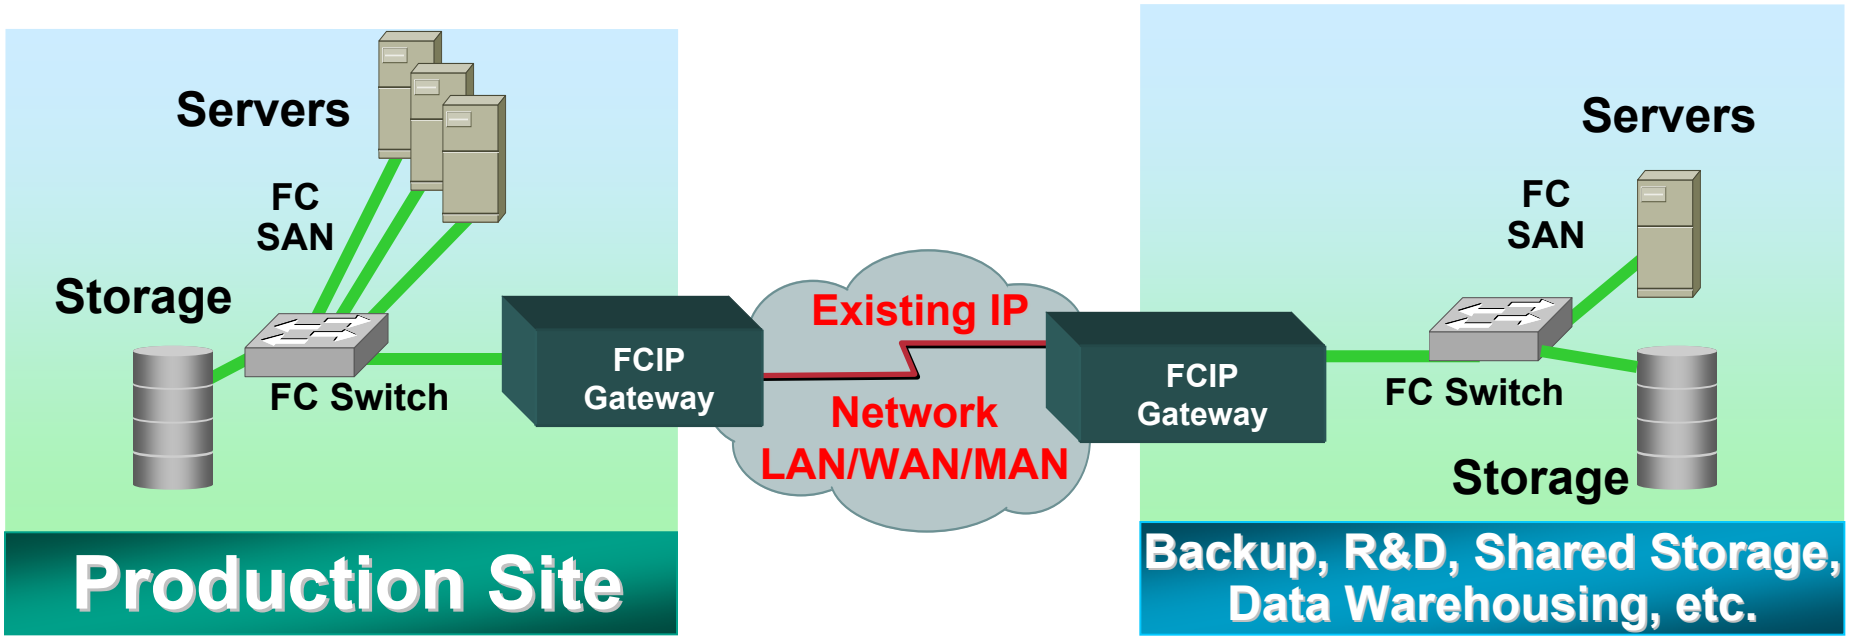

What is iSCSI?

- A way to access storage across an IP network as though it was locally attached.
- Transports SCSI protocol commands and data across an IP network
- Standards status update

  Draft RFC in IETF (last call)

  Current draft : draft-ietf-ips-iSCSI-13

  Expected to be ratified as a standard in Q3/4-02

  Major industry support (Cisco, IBM, EMC, Intel..)

iSCSI Architectural Model

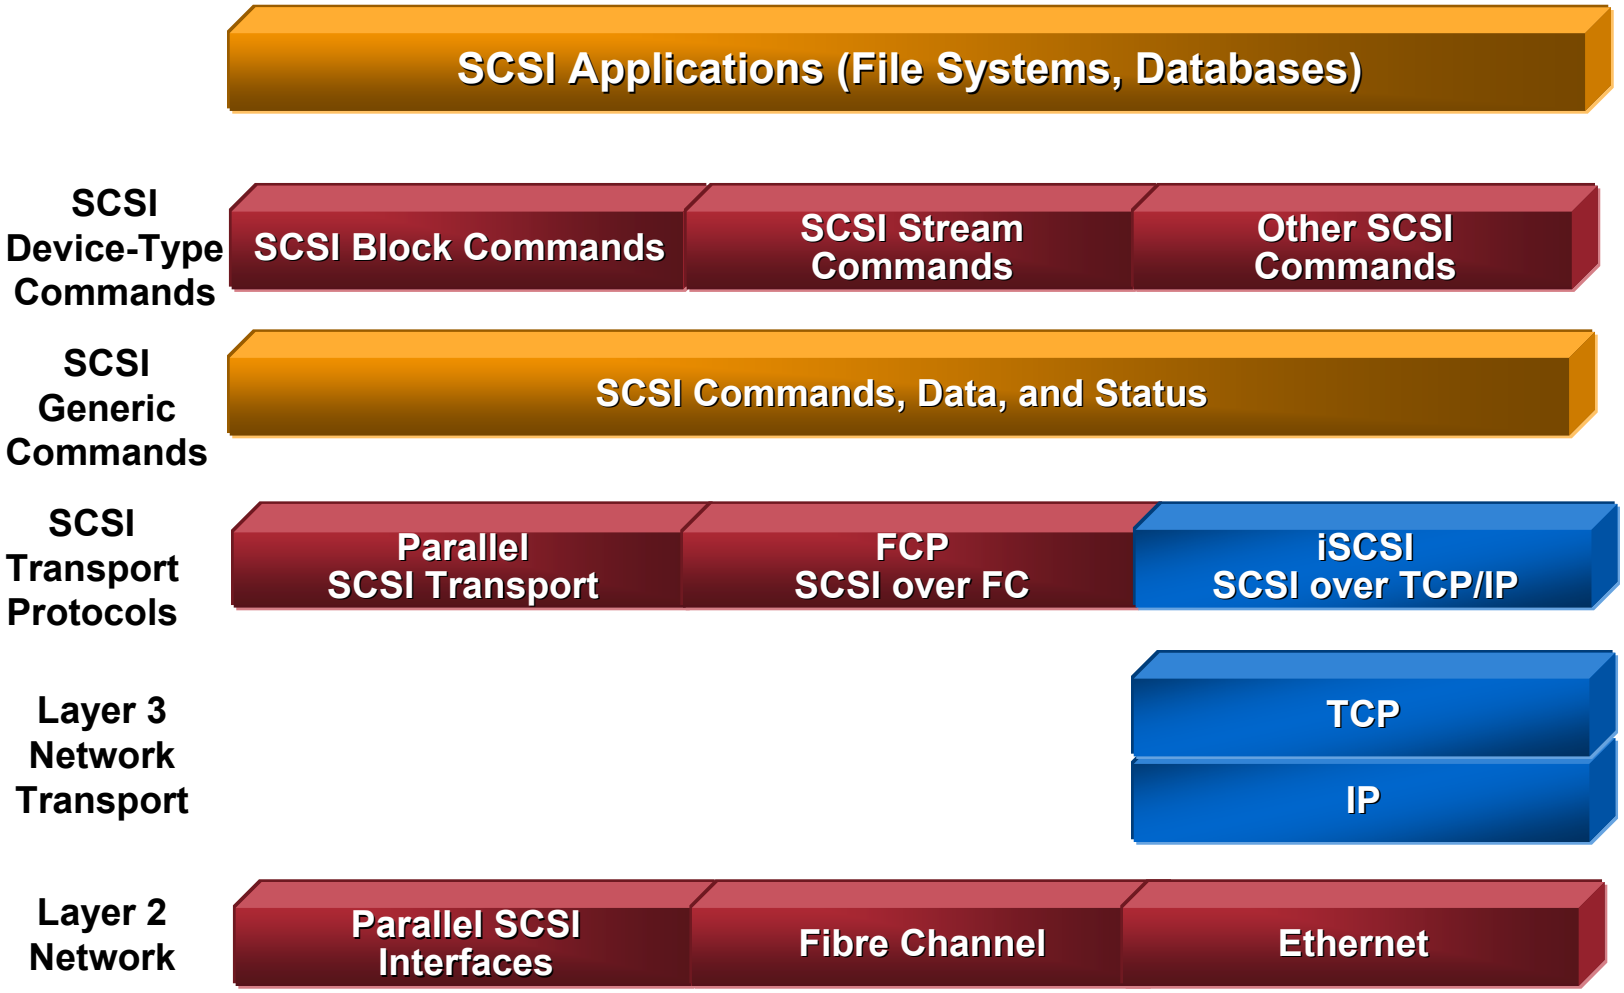

Initiator and Target Model for iSCSI

Initiator - SCSI device which is capable of originating SCSI commands and task management requests

Target - SCSI device which is capable of executing SCSI commands and task management requests

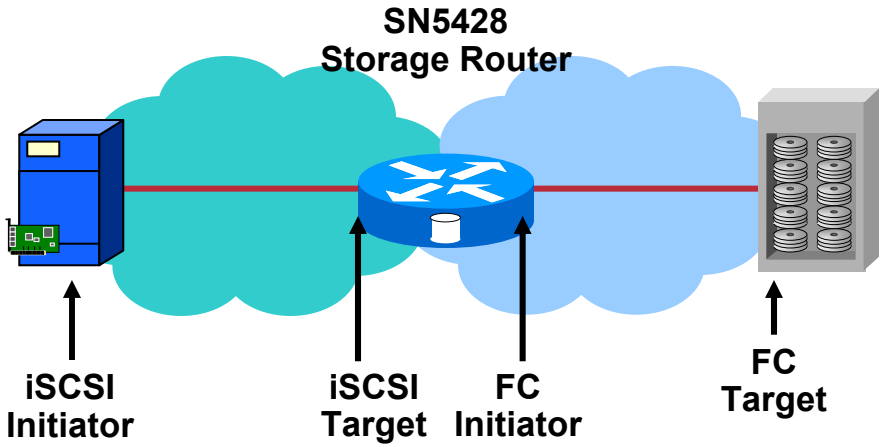

Initiators and Targets in a Multi-Protocol SAN

iSCSI Encapsulation

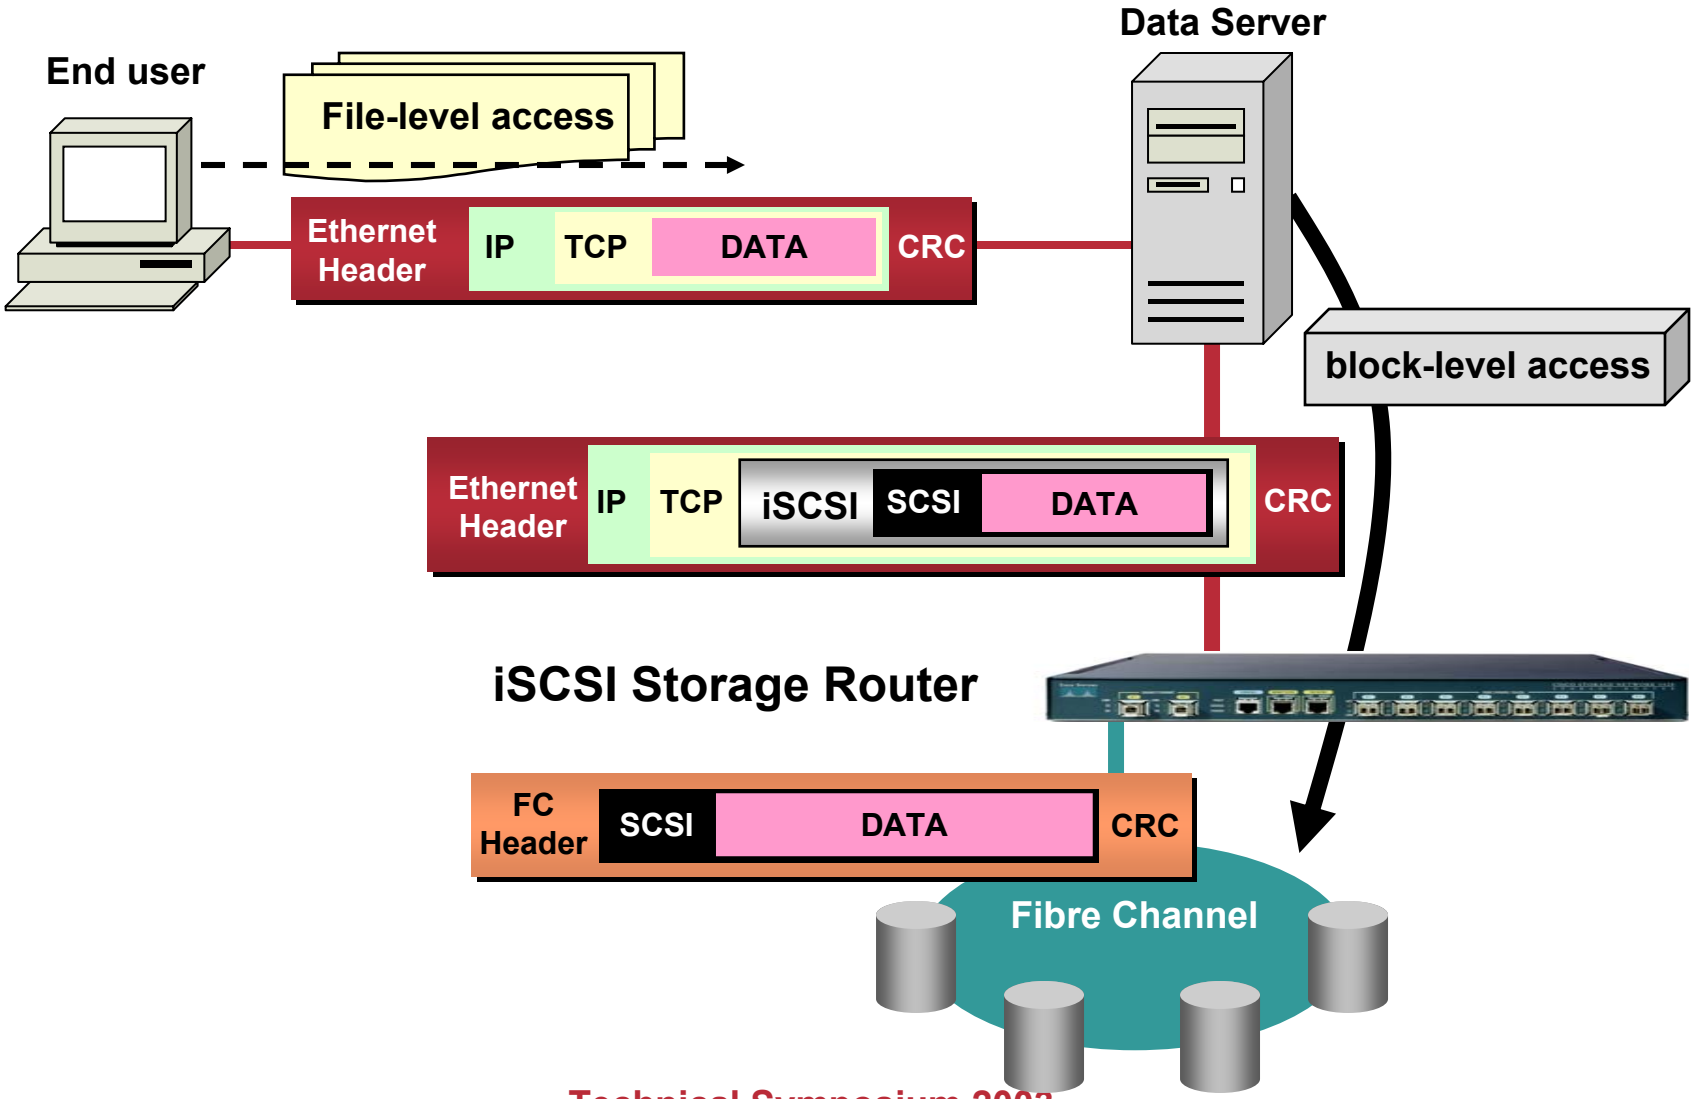

Cisco iSCSI Solution Architecture

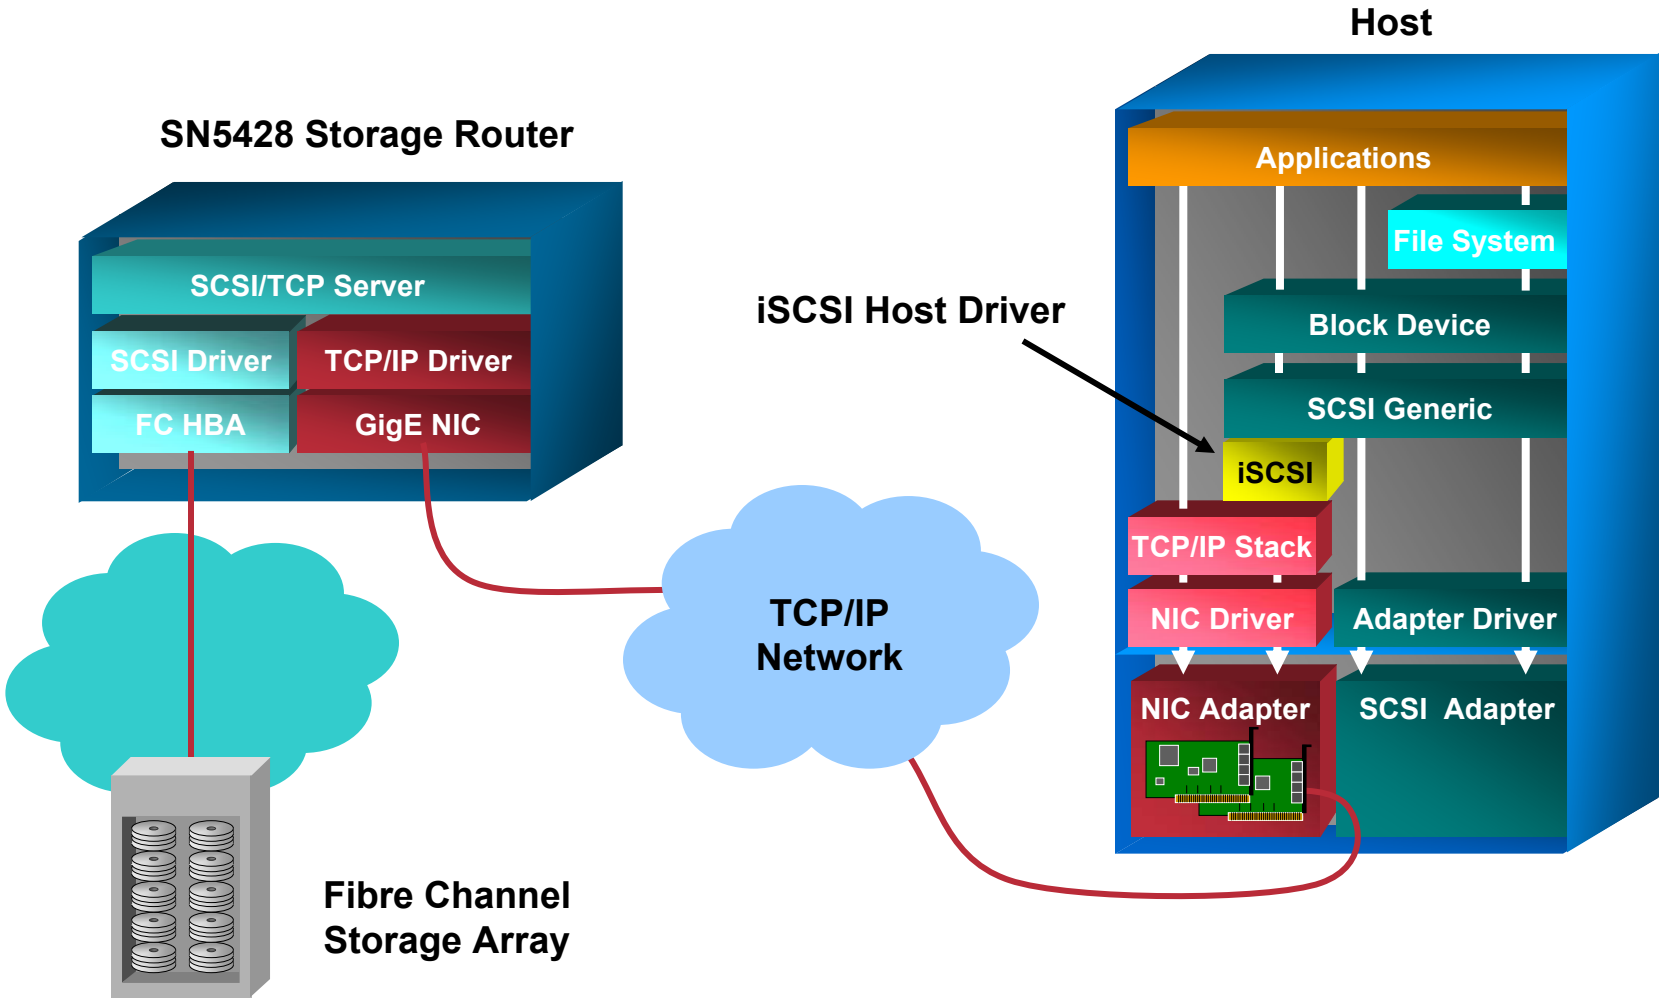

iSCSI + TCP Offload Engine (TOE)

- Hardware implementation of iSCSI and TCP/IP on a NIC
- Relieves host CPU from iSCSI and TCP processing, reducing CPU processing
- Two forms of TOEs
  - Offloads storage traffic only
  - Offloads TCP/IP for storage and data traffic
- Wire-rate iSCSI performance

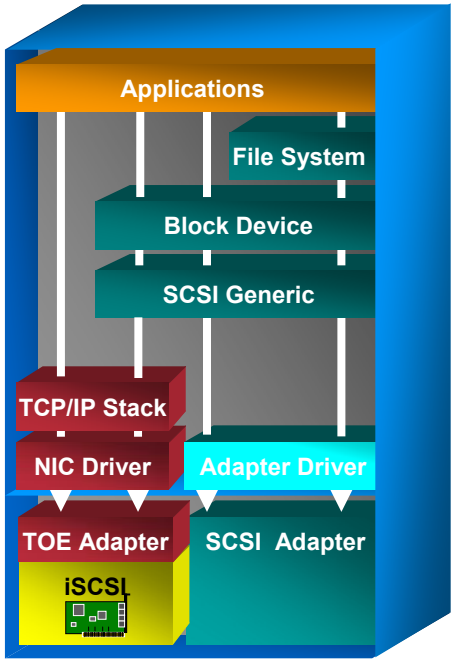

- Alacritech
- Adaptec
- Emulex
- Intel
- Qlogic

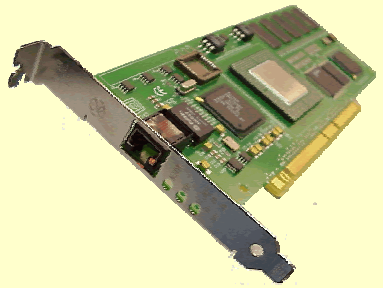

Cisco iSCSI Drivers

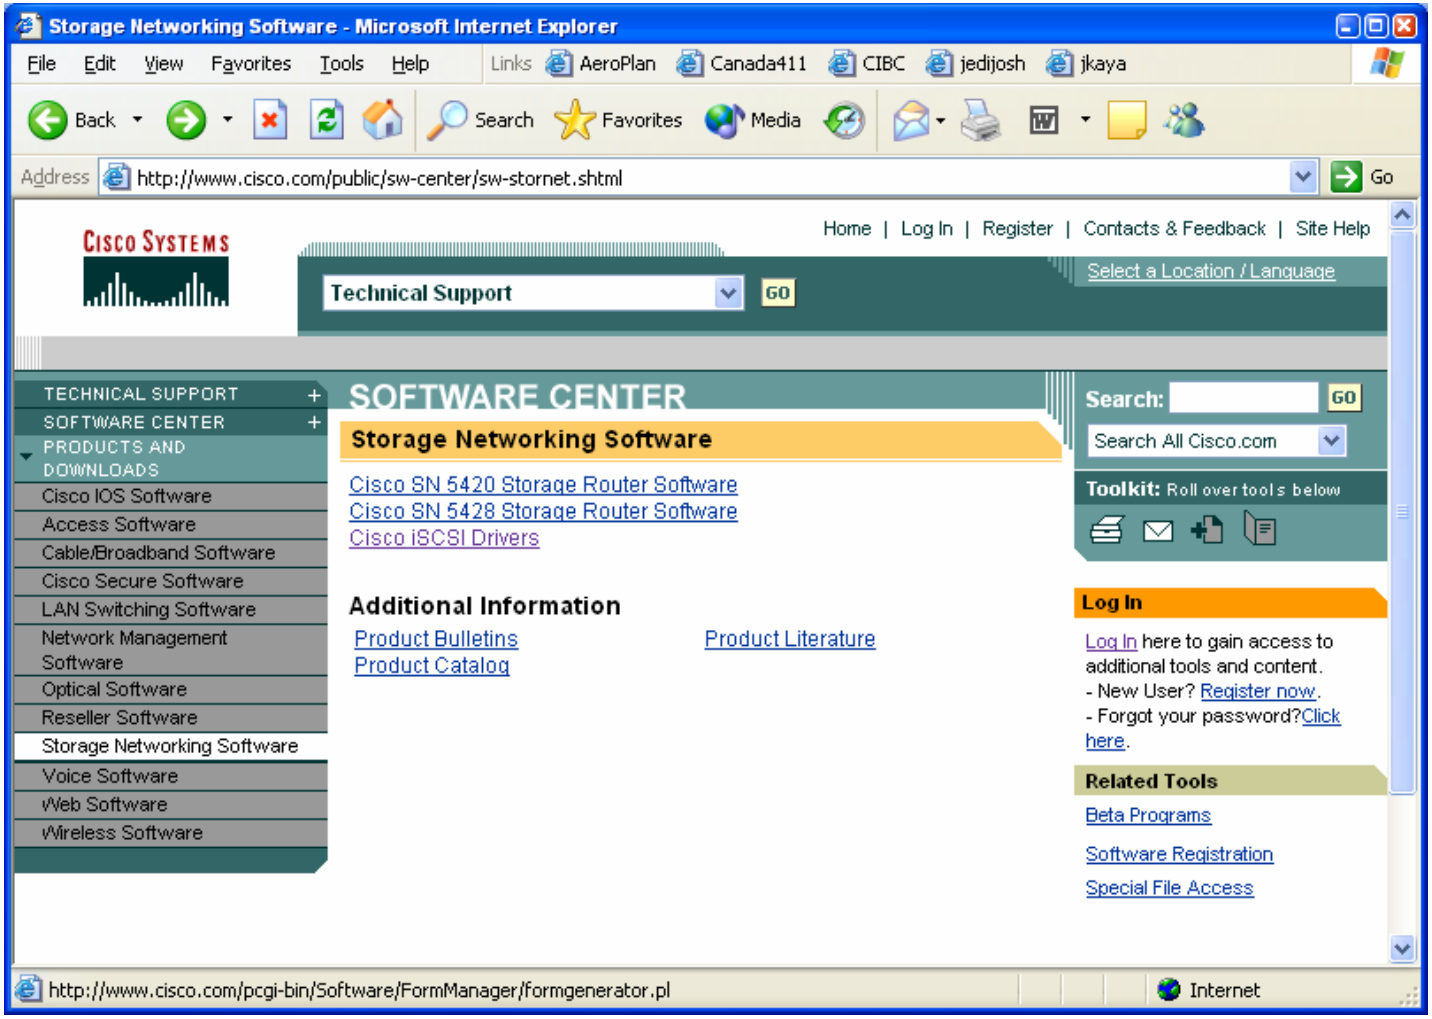

Cisco iSCSI Driver: Download

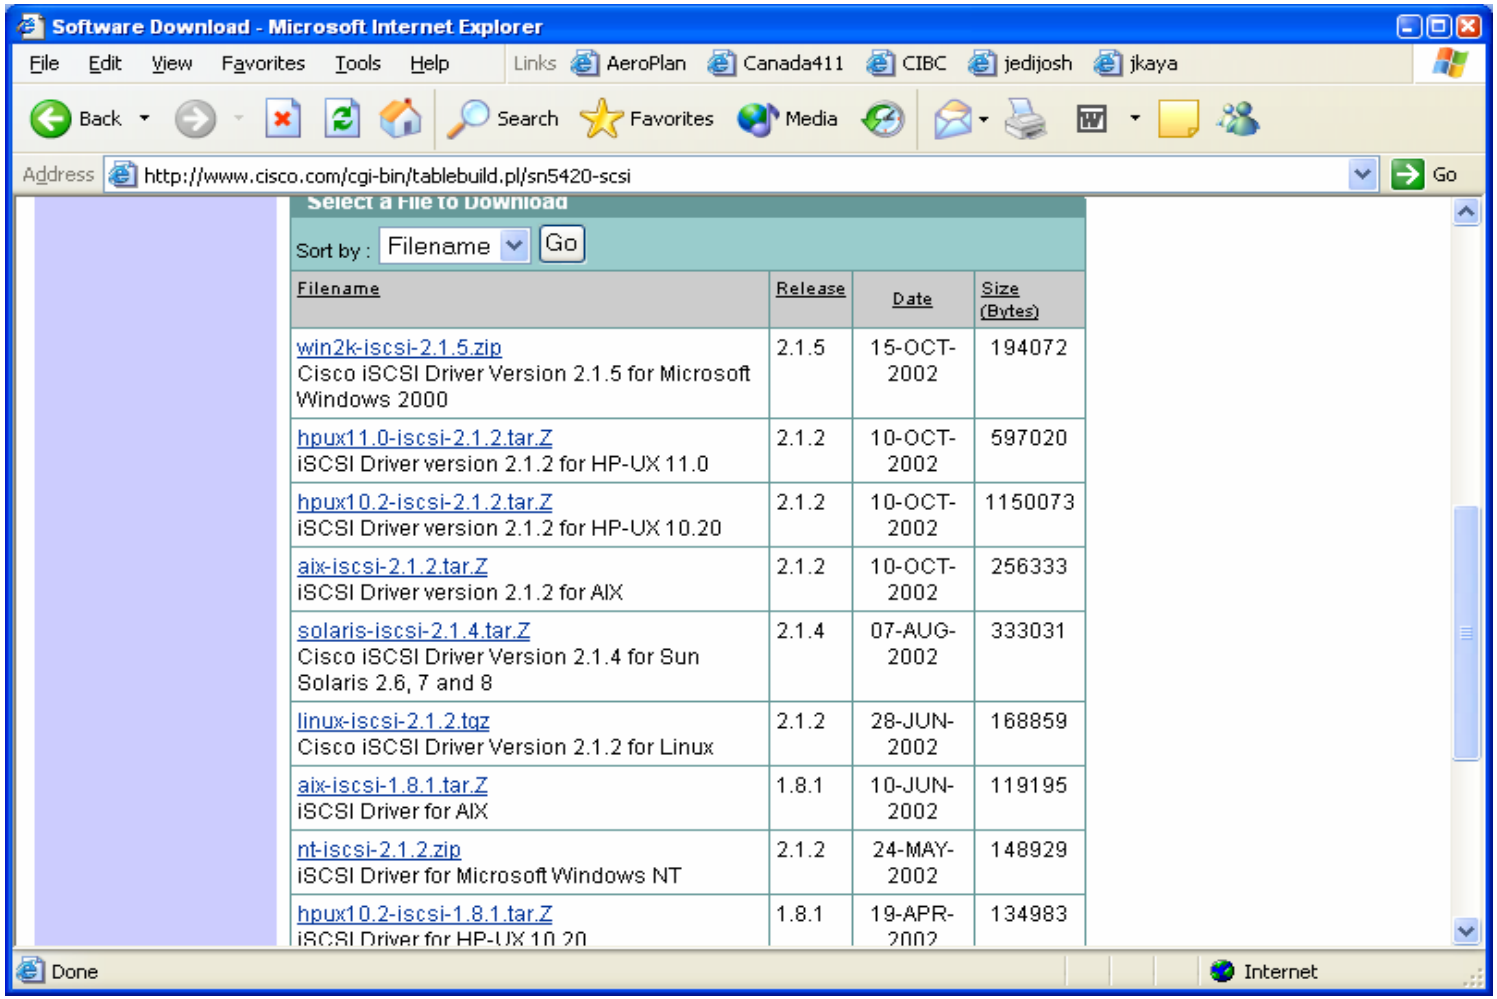

Filename	Release	Date	Size (Bytes)
win2k-iscsi-2.1.5.zip Cisco iSCSI Driver Version 2.1.5 for Microsoft Windows 2000	2.1.5	15-OCT-2002	194072
hpux11.0-iscsi-2.1.2.tar.Z iSCSI Driver version 2.1.2 for HP-UX 11.0	2.1.2	10-OCT-2002	597020
hpux10.2-iscsi-2.1.2.tar.Z iSCSI Driver version 2.1.2 for HP-UX 10.20	2.1.2	10-OCT-2002	1150073
aix-iscsi-2.1.2.tar.Z iSCSI Driver version 2.1.2 for AIX	2.1.2	10-OCT-2002	256333
solaris-iscsi-2.1.4.tar.Z Cisco iSCSI Driver Version 2.1.4 for Sun Solaris 2.6, 7 and 8	2.1.4	07-AUG-2002	333031
linux-iscsi-2.1.2.tgz Cisco iSCSI Driver Version 2.1.2 for Linux	2.1.2	28-JUN-2002	168859
aix-iscsi-1.8.1.tar.Z iSCSI Driver for AIX	1.8.1	10-JUN-2002	119195
nt-iscsi-2.1.2.zip iSCSI Driver for Microsoft Windows NT	2.1.2	24-MAY-2002	148929
hpux10.2-iscsi-1.8.1.tar.Z iSCSI Driver for HP-UX 10.20	1.8.1	19-APR-2002	134983

Configuration Applet

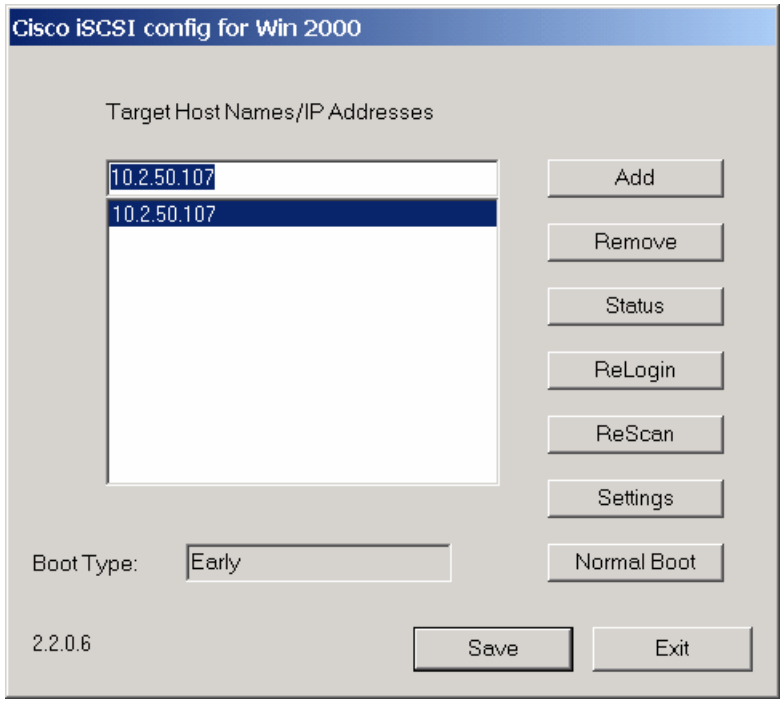

iSCSI Configuration Applet

iSCSI Connection Status

- T:0 means TargetID 0
- Correlate TargetID with Disk Manager

iSCSI Disks

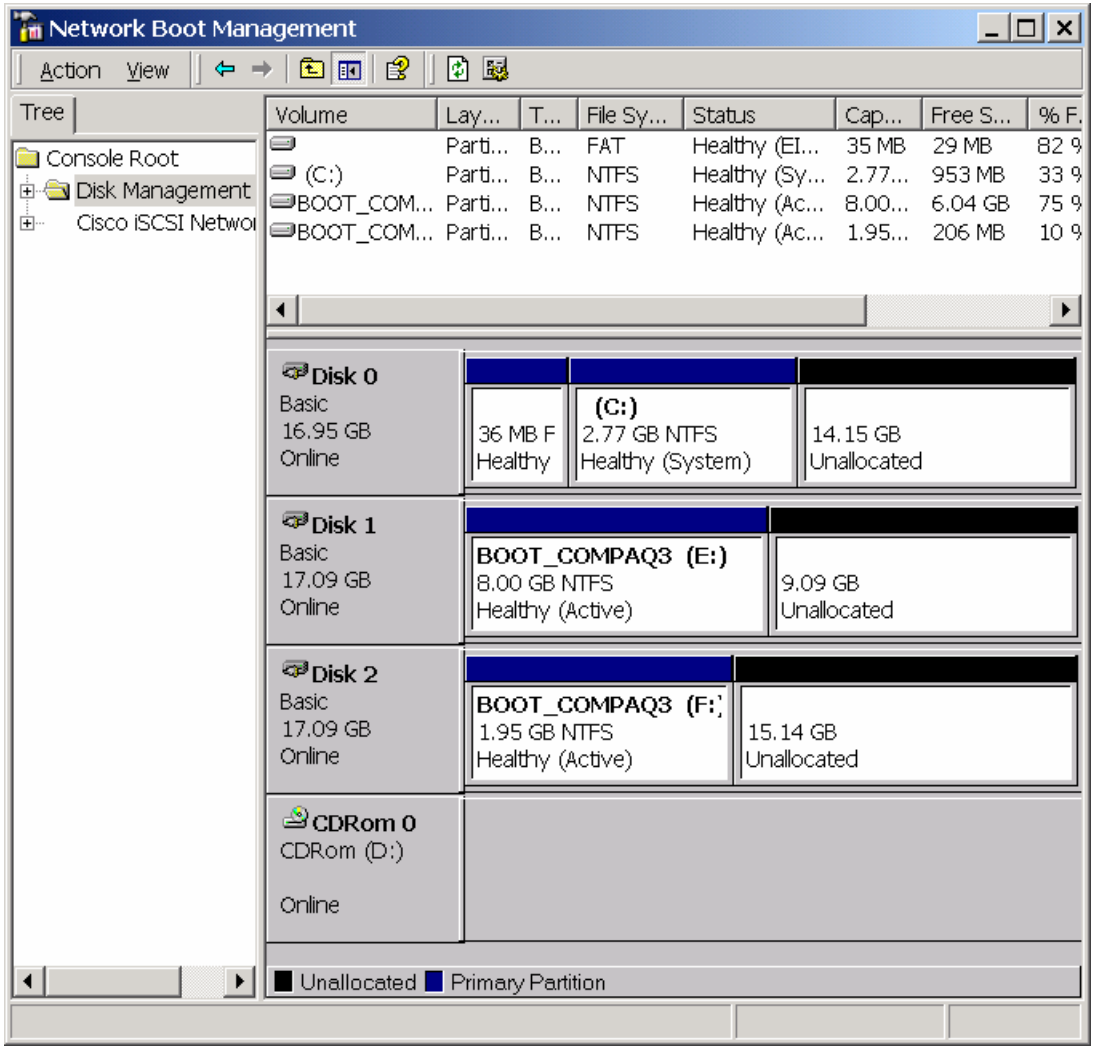

- **iSCSI disks appear in normal disk manager app**
- **Disk 0 is internal**
- **Disks 1 and 2 are iSCSI disks**

The Workgroup SAN Overview

- Workgroup SAN Environments
  - Enterprise Department
  - Small-Medium Business
- 10 - 40 Servers with Midrange Applications
  - Email – Exchange, Notes
  - Database – SQL, Oracle
  - Financials – Great Plains, Lawson, Oracle
  - Web Servers – IIS
  - File Servers
  - Print Servers
  - Customer Developed
- Most servers cost under $10,000
- The network/system/storage admin are the same person

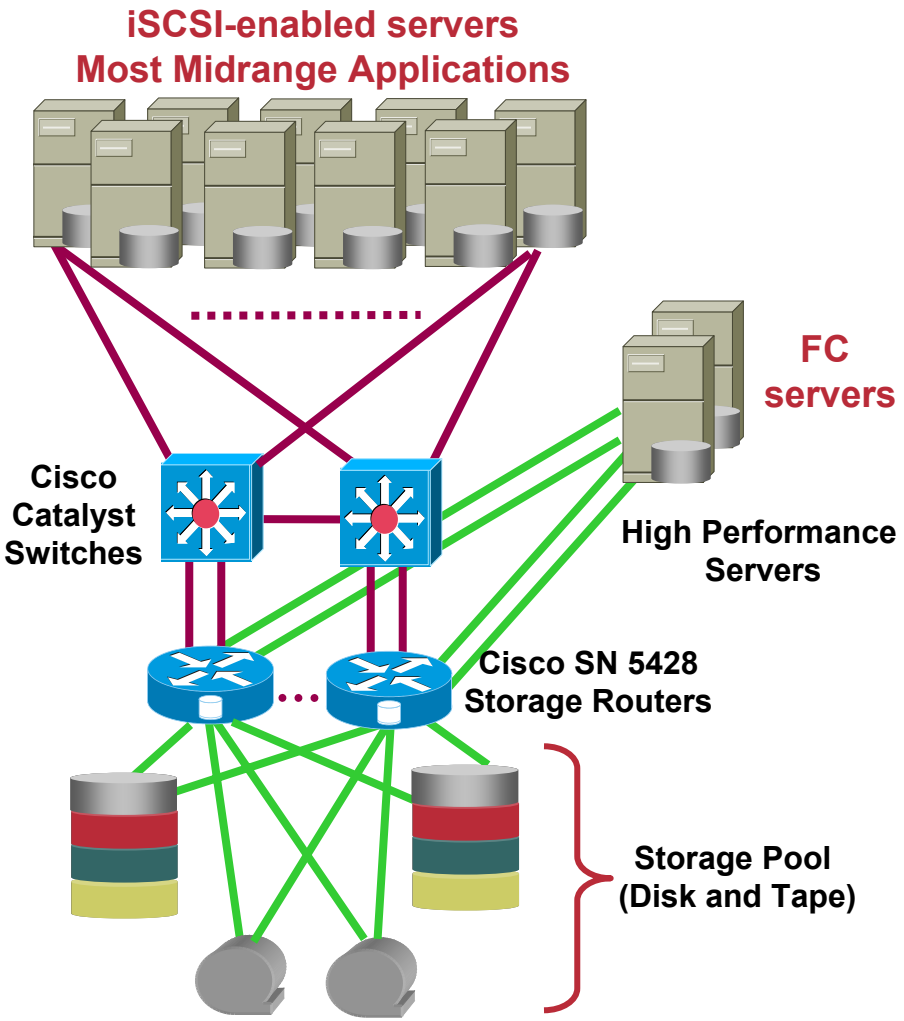

Useful Information

- Storage Network Industry Association (SNIA)
  http://wwwin.snia.org
- Internet Engineering Task Force – IP Storage
  http://www.ietf.org/html.charters/ips-charter.html
- ANSI T11 – Fibre Channel
  http://www.t11.org/index.htm
- Cisco Storage Networking
  http://www.cisco.com/go/storagenetworking
- Cisco AVVID Storage Networking Partner Program
  http://www.cisco.com/warp/public/779/largeent/partner/esap/storage.html
- Cisco Storage Router Product Information
  http://www.cisco.com/go/storagenetworking
- Cisco Metro Optical Product Information
  http://www.cisco.com/go/comet