Storage for Networking Professionals

Steve Hrehoriak
Cisco Systems, SE

Agenda

- Overview of Storage Approaches
  - Traditional
  - Emerging
- Storage Basics and Terms
- Fibre Channel SAN Architectures
- Cisco's Family of Storage Products

The World of Storage

- Cisco is in the Storage Business
- Leverage Experience in Switching
- Data is Data Regardless of the Transport
- Four Fundamental Approaches to Storage
  - Traditional
    - Direct Attached Storage
    - Network Attached Storage
    - Storage Area Networks
  - Emerging
    - IP Storage Area Networks

Storage Basics and Terms

I/O Channel

Simply stated, It is the technology which resides between a computer and the device used to store its data.

This relationship can exist either internal to the computer casing or can extend to external storage devices.

Storage device is only accessible by attached host computer

Examples are SCSI, Fibre Channel, ESCON

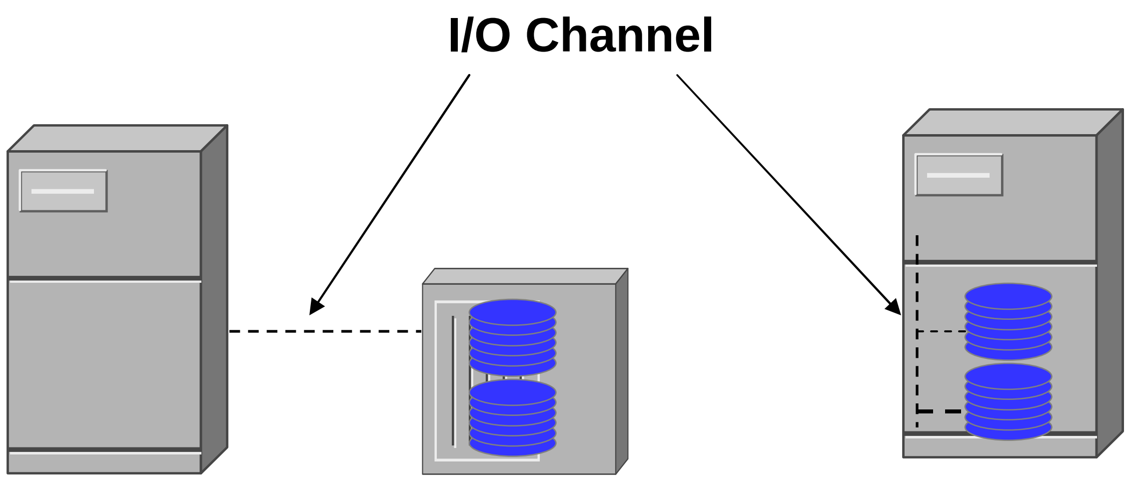

Small Computer System Interface (SCSI)

- Parallel interface I/O technology
- Maximum cable run length is 25 meters
- Speeds up to 320 MB/sec (Ultra-320 on 16 bit wide bus)
- Maximum of 16 (I/O controller + devices) SCSI devices per bus
- Several standards: SCSI-1, SCSI-2, and SCSI-3(Ultra-2&3)
- Scalability and distance limitations rule out support of large scale storage systems requiring disaster recovery

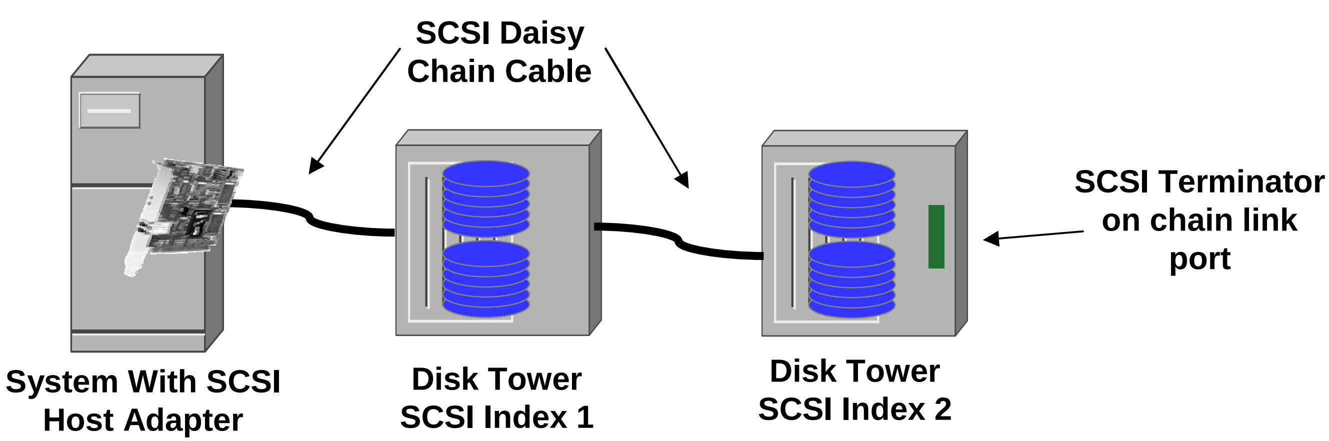

What Is Fibre Channel ?

Cisco.com

Initial effort started in 1988

Developed by ANSI

Combines the benefits of both channel and network technologies

SCSI and IP are the only upper layer protocols commercially available on Fibre Channel

Benefits of mapping SCSI onto Fibre Channel include:

- Faster speed
- Ability to connect more devices together
- Greater distances allowed between devices

Runs on copper (coax) or glass (fiber optic) cable

Fibre Channel Offers...

- Multiple protocol support (today - mainly IP and SCSI)
- Networking capability and functionality
- Heterogeneous interconnect
- Speed: 1 and 2 Gbps. 10 Gbps in near future
- Boasts: bandwidth, availability, reliability, integrity, and scalability

Fibre Channel - Port Types

Cisco.com

‘N’ port: Node ports used for connecting peripheral storage devices to switch fabric or for point to point configurations

‘F’ port: Fabric ports reside on switches and allow connection of storage peripherals (‘N’ port devices)

‘L’ port: Loop ports are used in arbitrated loop configurations to build storage peripheral networks without FC switches. These ports often also have ‘N’ port capabilities and are called ‘NL’ ports.

‘E’ port: Expansion ports are essentially trunk ports used to connect two Fibre Channel switches

‘G’ port: A generic port capable of operating as either an ‘E’ or ‘F’ port. If also capable of acting in an ‘L’ port capacity - known as a ‘GL’ port.

Fibre Channel Topologies

Point To Point

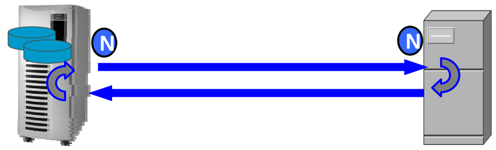

Arbitrated Loop

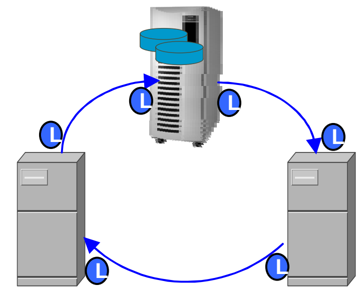

Switched Fabric

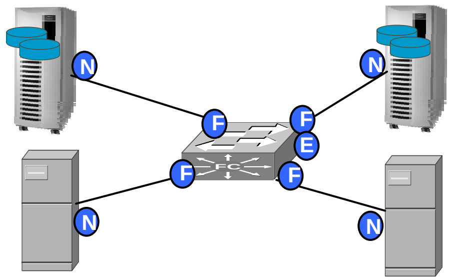

Switched Fabric

Cisco.com

- Max nodes = 16 million (24 bits)
- Max bandwidth = 200 MB/sec
- Nodes (N ports) connect to fabric (F ports)
- End to end connection managed by N ports
- Routing and addressing handled by fabric
- E port provides trunk connectivity to another Fibre Channel switch.

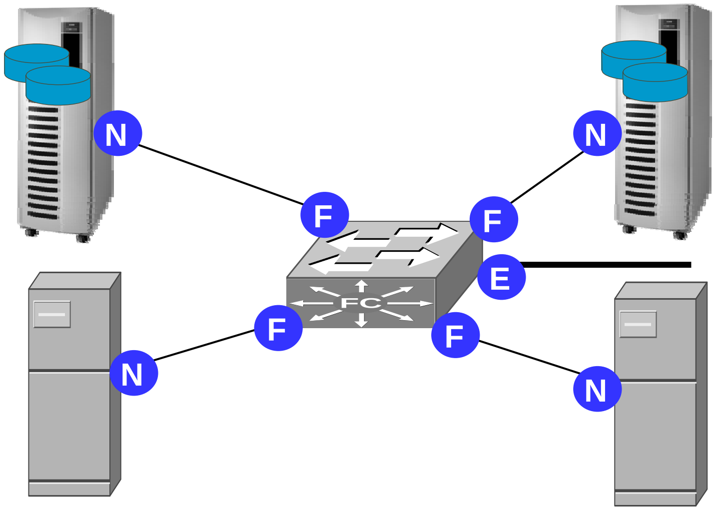

I/O Channels

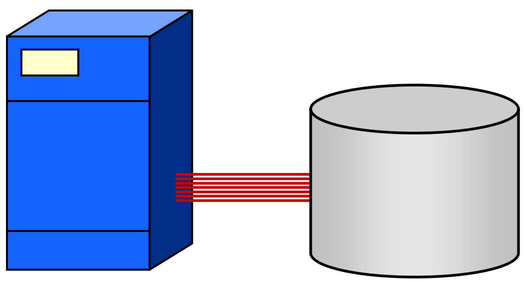

I/O Channel (SCSI)

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

Data Networks

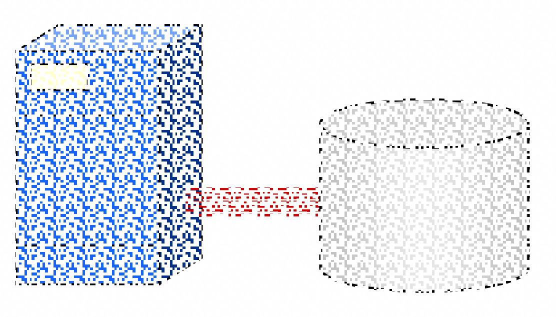

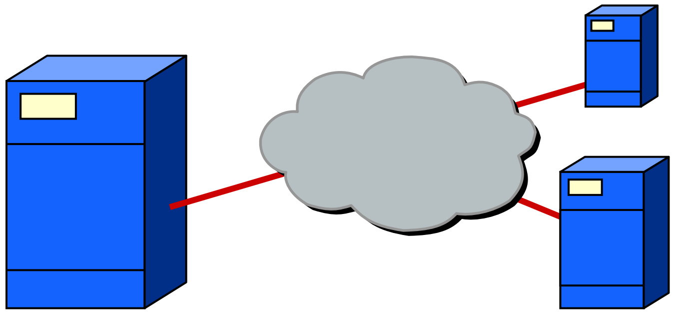

I/O Channel (SCSI)

- 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

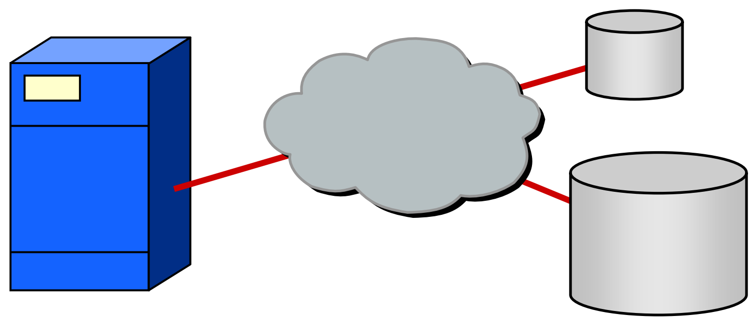

I/O Channel	Fibre Channel	Network
X Few devices	✓ Many devices	✓ Many devices
X Static	✓ Dynamic	✓ Dynamic
✓ Low latency	✓ Low latency	X High latency
X Short distances	✓ Long distances	✓ Long distances
✓ Hardware-based delivery management	✓ Hardware-based delivery management	X Software-based delivery management

Fibre Channel Summary

Cisco.com

- Nodes are "transparent" devices
- Ports are intelligent interface points
- Standard port types:
  - N_Port
  - NL_Port
  - F_Port
  - FL_Port
  - E_Port
  - B_Port
- Cisco port types:
  - TL_Port
  - TE_Port
  - SD_Port
  - Fx_Port
- Automatic port configuration: U_Port/G_Port

Just a Bunch Of Disks (JBOD)

- Drives are independantly attached to I/O channel
- Scalable, but requires servers to manage multiple volumes
- No protection in the event of drive failure
- Drives share common power supplies and physical chassis

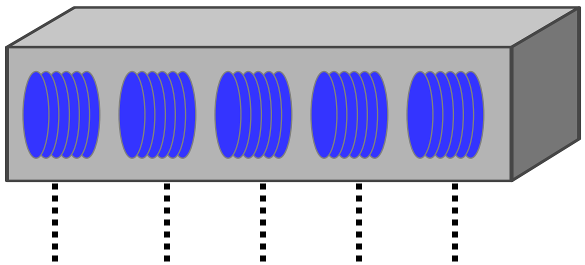

Redundant Array of Inexpensive Disks (RAID)

- Fault -tolerant grouping of disks that server views as a single disk volume.
- Combination of striping, mirroring, and parity checking
- Self Contained, manageable unit of storage
- Management : One file system across entire virtual disk
- RAID delivers Capacity, performance, reliability, and availability benefits.

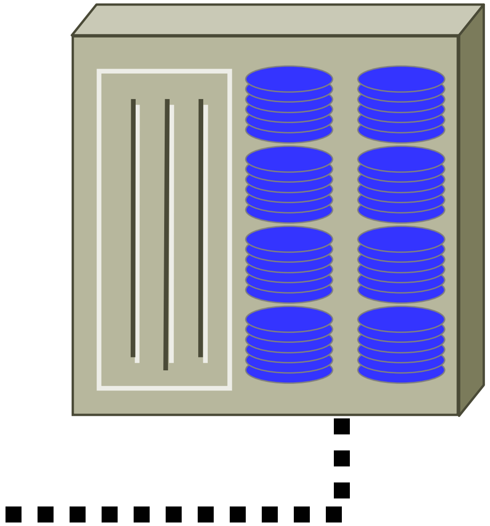

Raid Levels

Cisco.com

RAID levels

RAID levels 0 through 6 were defined in the original University of California Berkeley RAID project. RAID 2, 4, and 6 are rarely seen in commercial products. RAID 0 is merely disk striping, which has some performance advantages but stores no parity information and thus does not offer true RAID data protection.

RAID 1 offers complete duplication of data, and this 100 percent data redundancy provides the best protection--but it is much too expensive for most applications. RAID 3 and RAID 5 each use one extra disk to store parity information needed to recreate data in the event of a single disk failure. RAID 3 uses a dedicated parity disk and is typically faster for throughput-oriented applications, such as file transfer and other sequential applications. RAID 5 distributes the parity information across all disks in the array and is typically faster for transaction processing and other random access applications. These results are relevant mostly in arrays that have little or no controller cache memory. In products with significant cache memory (64MB or more) on-board the controller, performance will be higher in all cases due to the distinctly higher abilities of the controller; these products will perform in a vastly superior manner regardless of RAID mode.

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**

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.

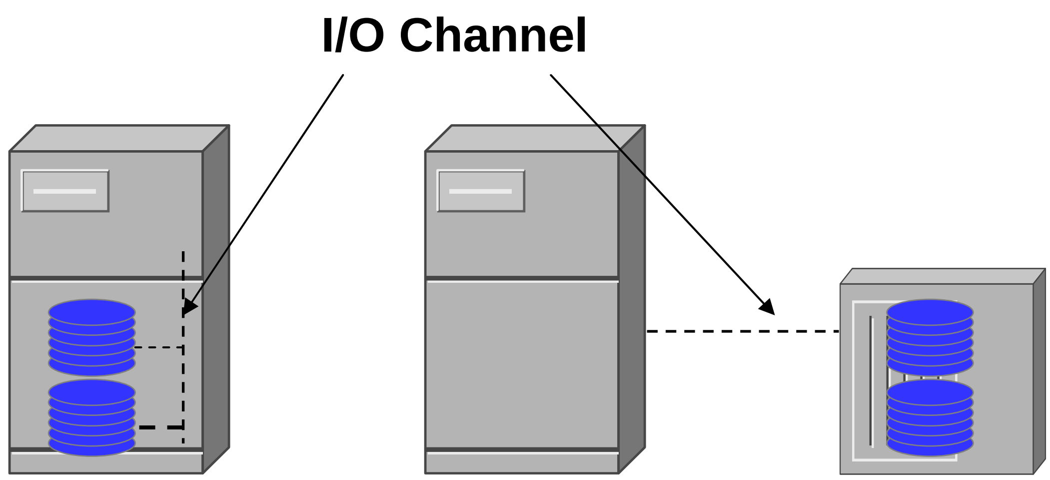

DAS Options

The "typical" storage environment

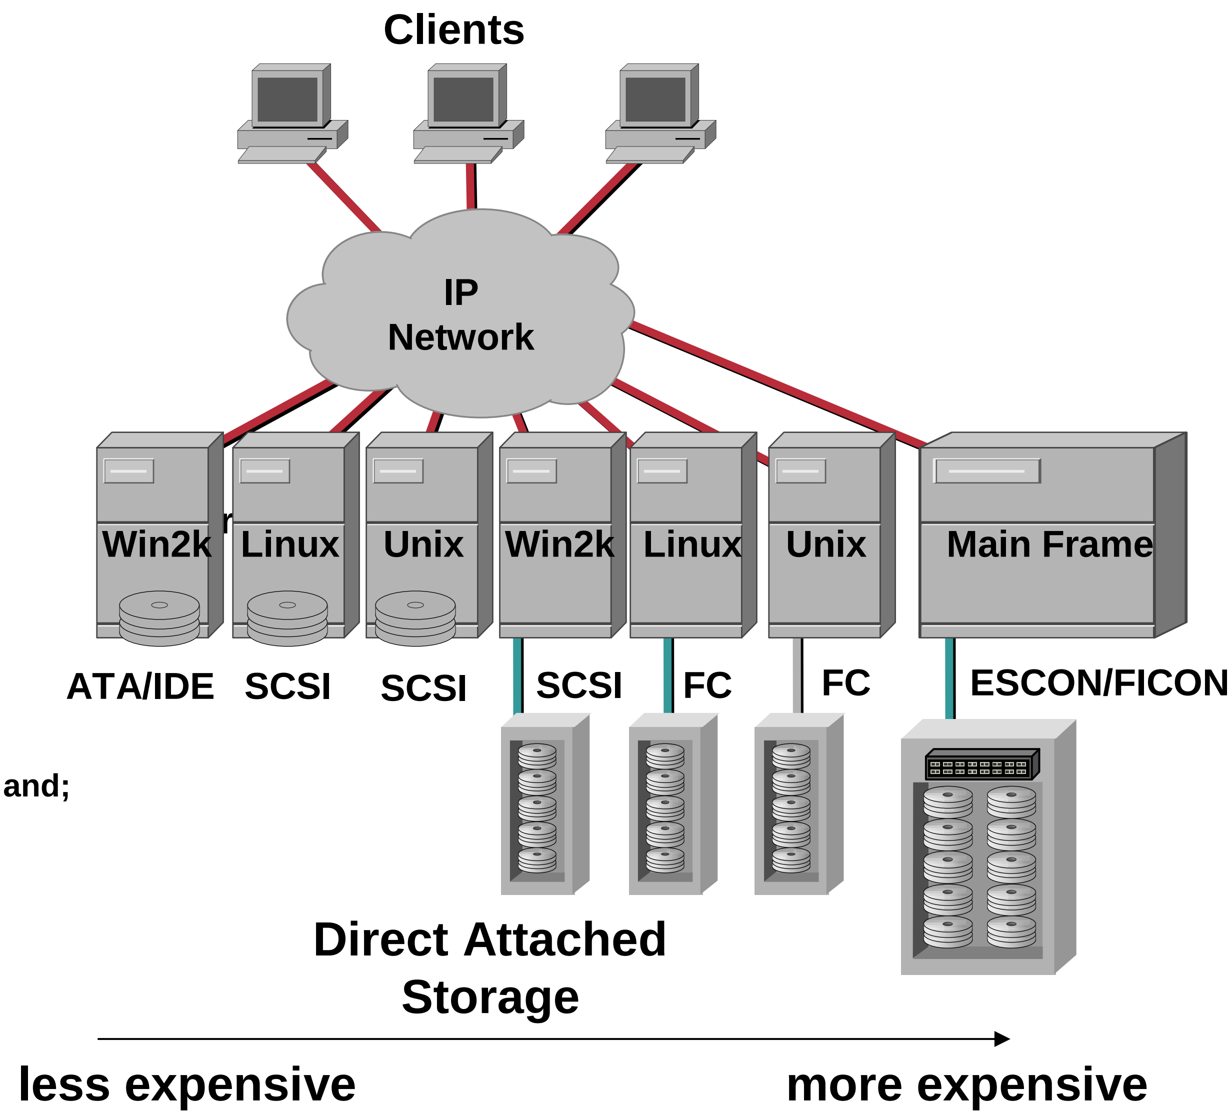

Server CPU handles I/O requests and;
•User DB inquiries
•User file/print serving
•Data integrity checks
•Comm with other devices

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)

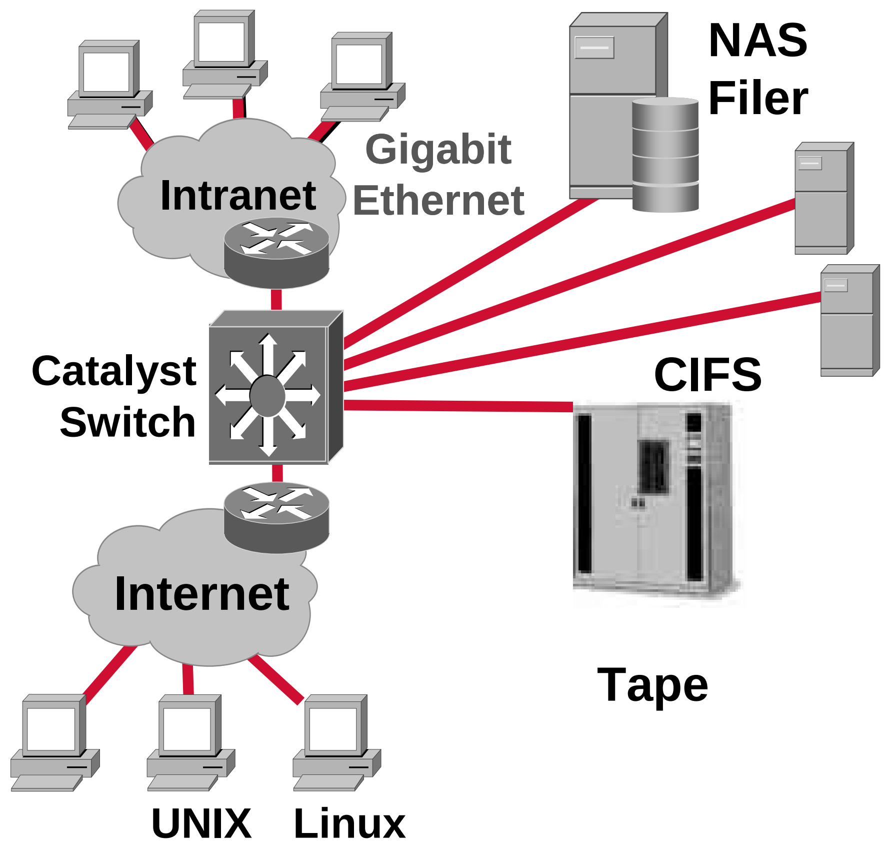

NAS Protocols

Client systems require a protocol to communicate with NAS devices:

- Windows systems use the Common Internet File System (CIFS) protocol
- Unix systems (and others) use the Network File System (NFS)

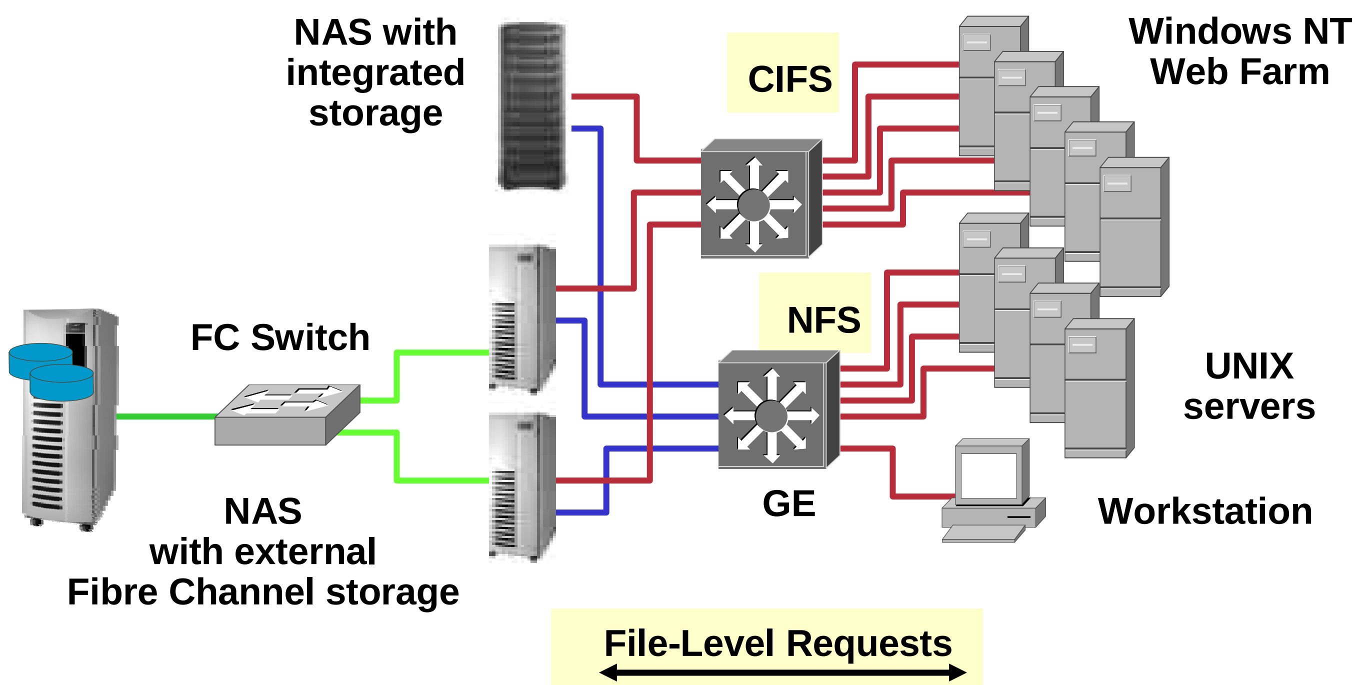

SAN: What Is It?

Cisco.com

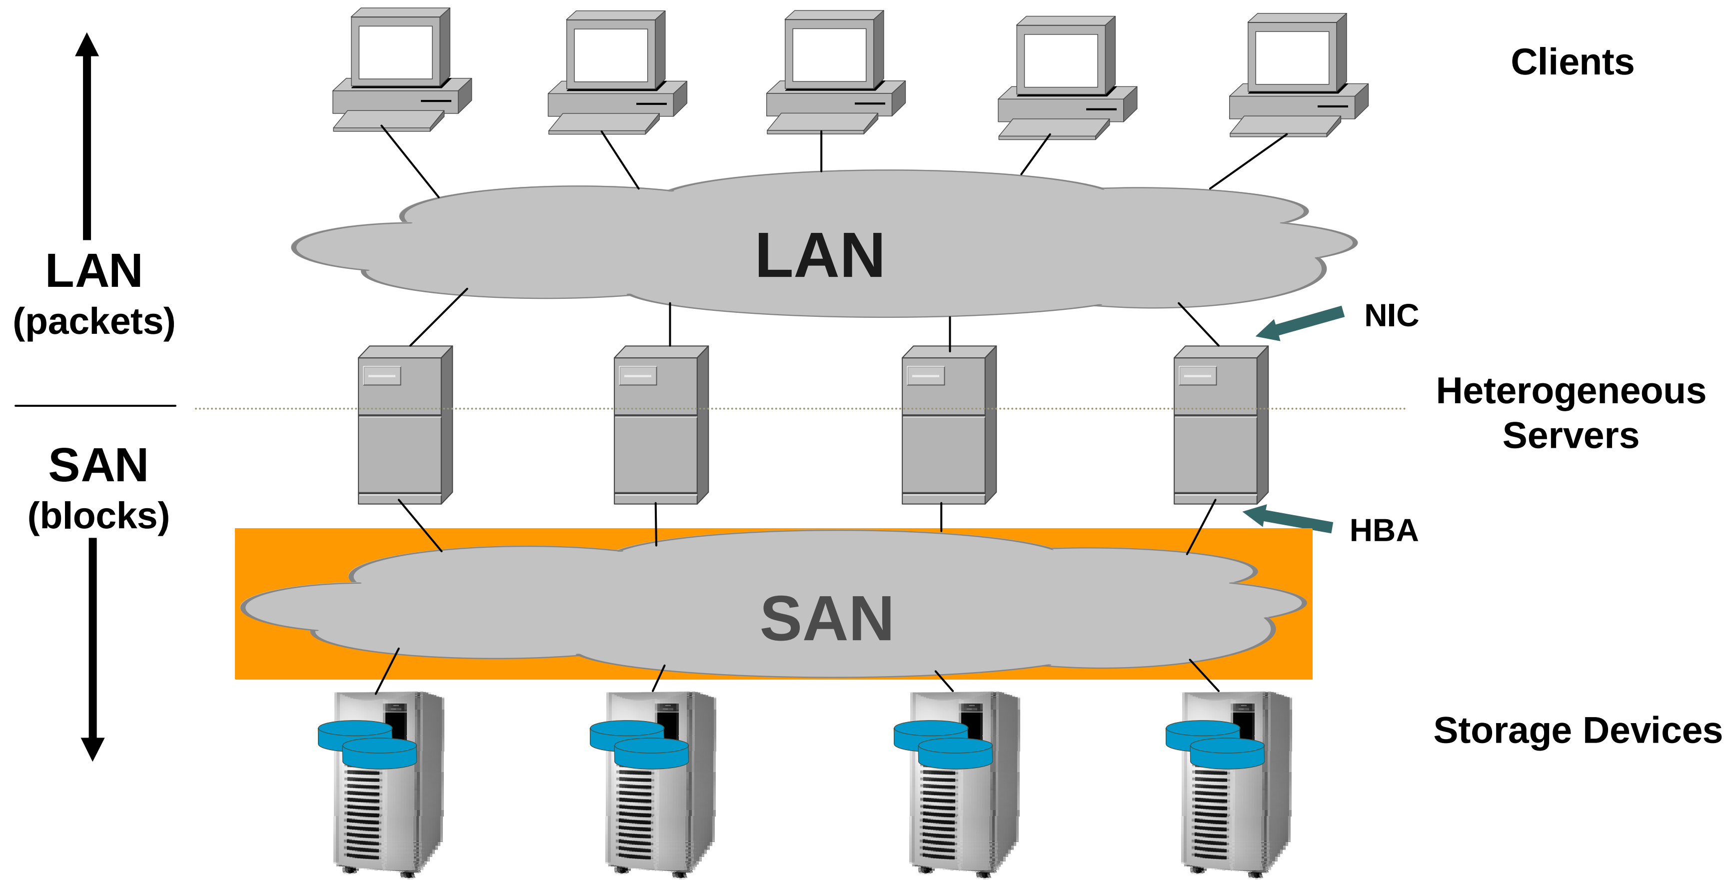

"A reliable transport for running the SCSI protocol"

SAN Components

Servers with Host Bus Adapters (HBAs)

Storage systems

- RAID
- JBOD
- Tape
- Optical

Hubs(managed and unmanaged)

Switches (loop and fabric)

Bridges and channel extenders

SAN management software

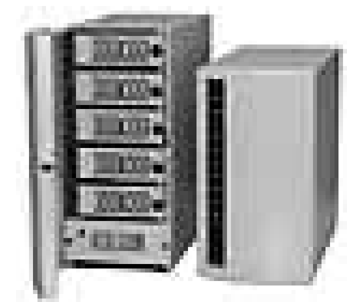

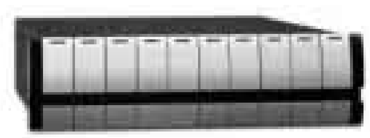

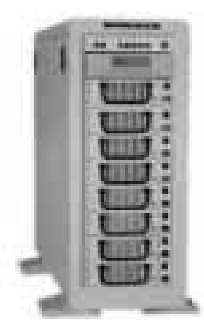

SANs: Scalability and Performance

Storage Expansion

- No impact on servers

Server Expansion

- No impact on storage

Load Balancing

- Active parallel paths

Bandwidth on Demand

- Robust topology

SAN and NAS - When To Use What??

NAS and SAN solutions solve different application storage requirements

Use SAN for DBMS (OLTP) storage and most application scenarios

Use NAS for file serving and file sharing applications

…Very few storage vendors or products can meet all business needs

IP SAN Storage Overview

SAN Definition

“A SAN is a reliable transport for running the SCSI protocol”.

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

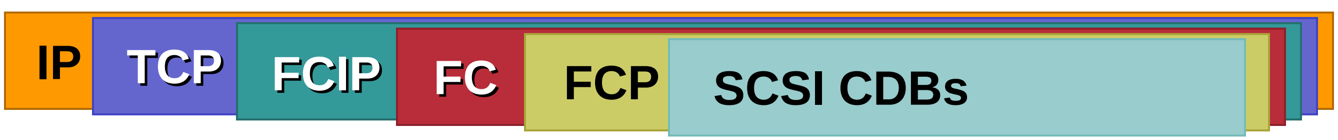

What is an iSCSI Router?

Purpose is to link iSCSI-enabled servers on IP networks with Fibre Channel based storage subsystems

- **The iSCSI router *terminates* iSCSI Command Descriptor Blocks (CDBs) and *re-initiates Fibre* Channel FCP CDBs**

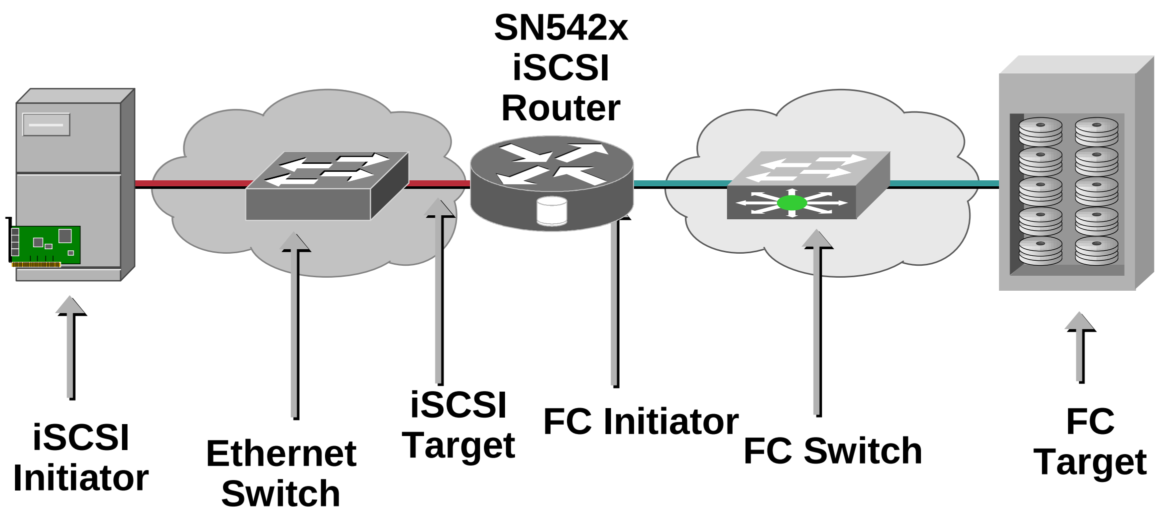

iSCSI Communication

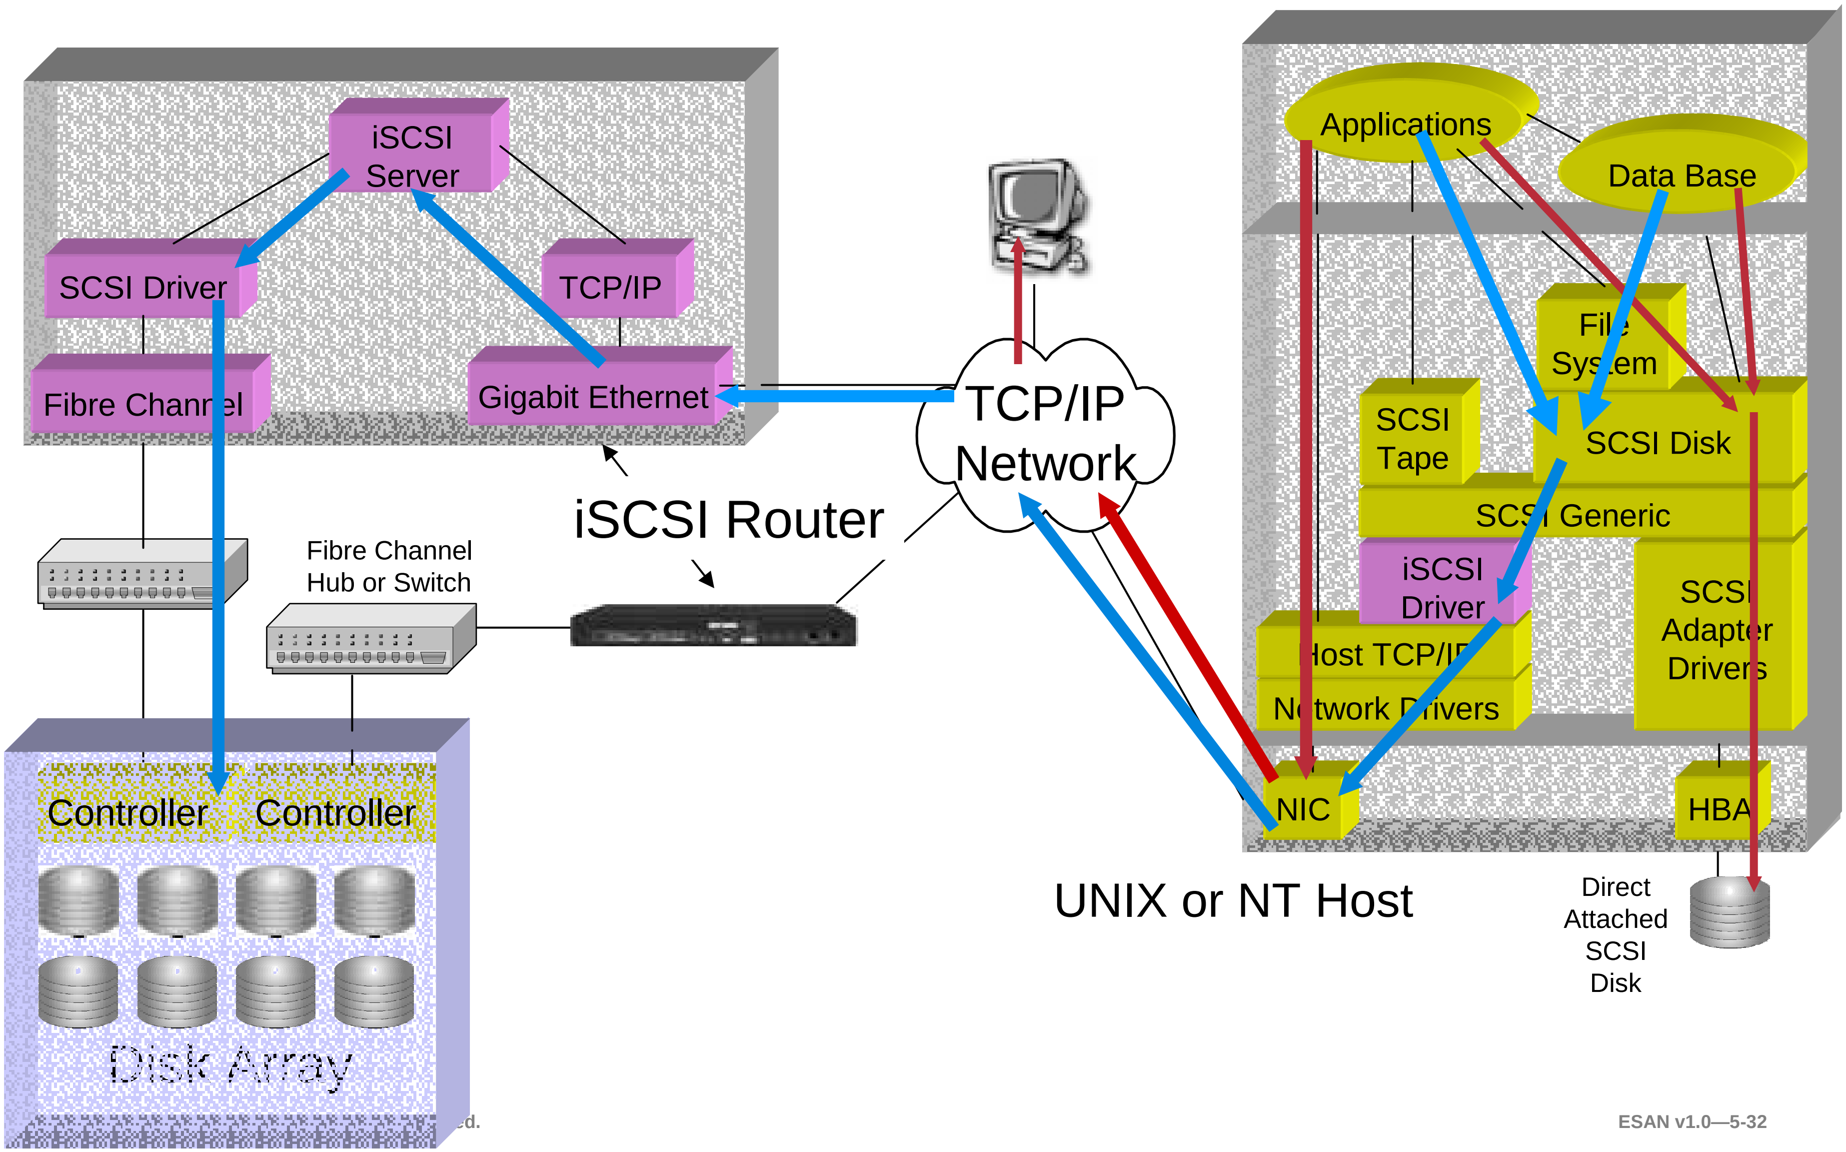

iSCSI Host Implementations

- iSCSI device driver
- Legacy NIC
- Legacy software based TCP/IP stack

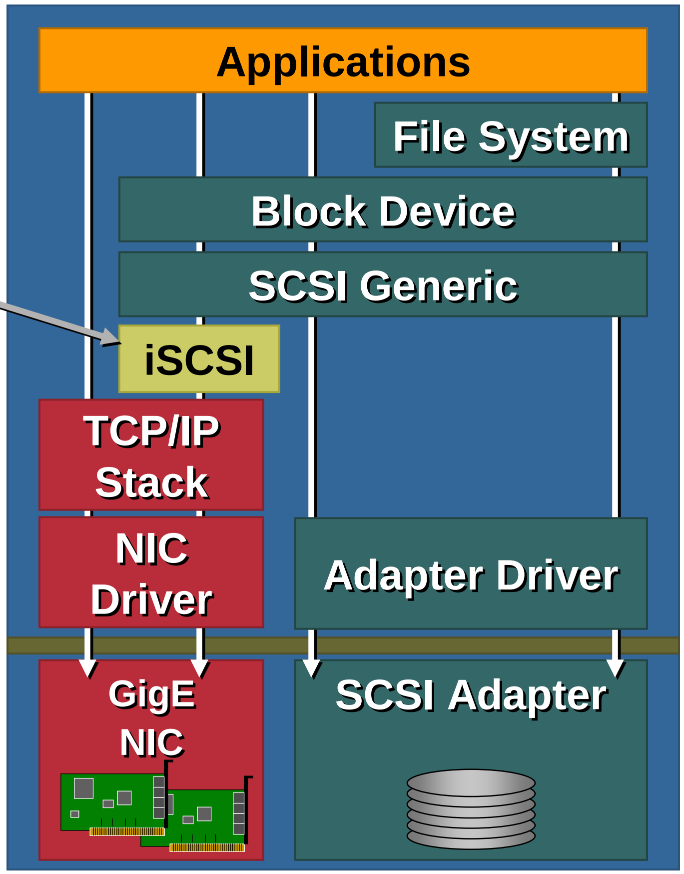

iSCSI Host Implementation (1)

iSCSI Host Implementations

Cisco.com

iSCSI device driver

TCP Offload Engine (TOE)

- Processing implemented in NIC
- CPUs used to be faster than networks, but not anymore
- Fewer interrupts
- Eliminates memory copies

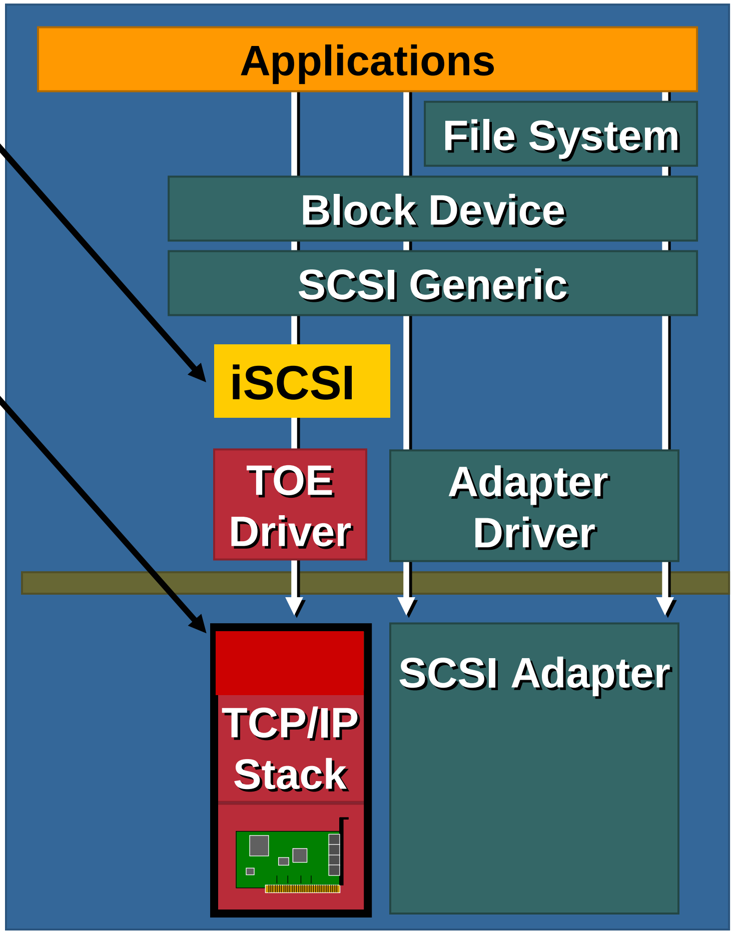

iSCSI Host Implementation (2)

iSCSI Host Implementations

iSCSI and TOE both implemented in NIC

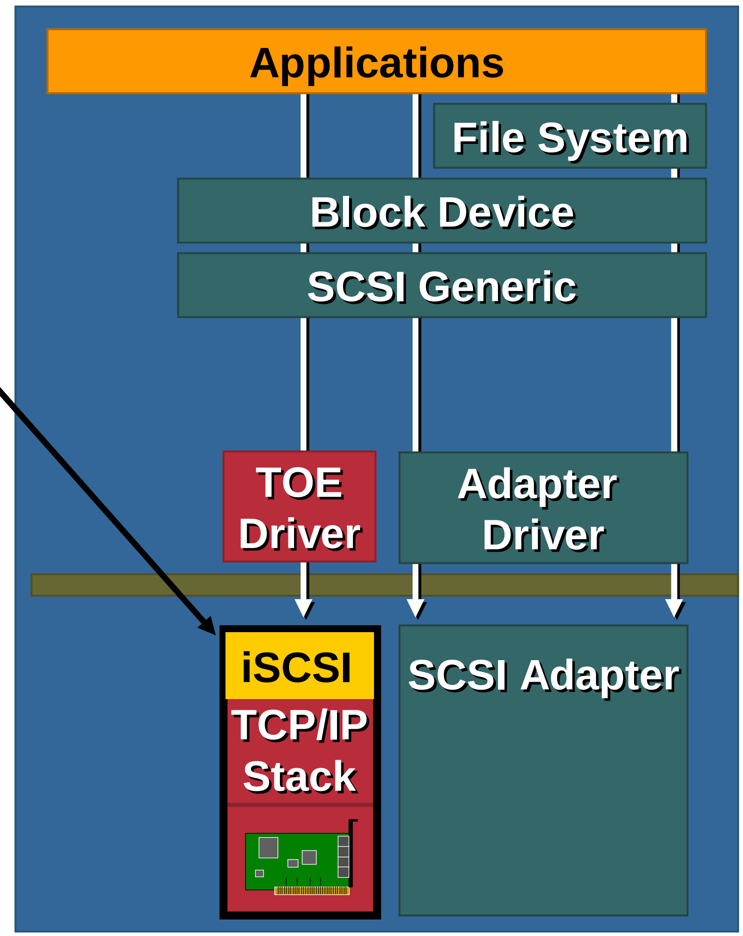

iSCSI Host Implementation (3)

IP Storage Network

Cisco.com

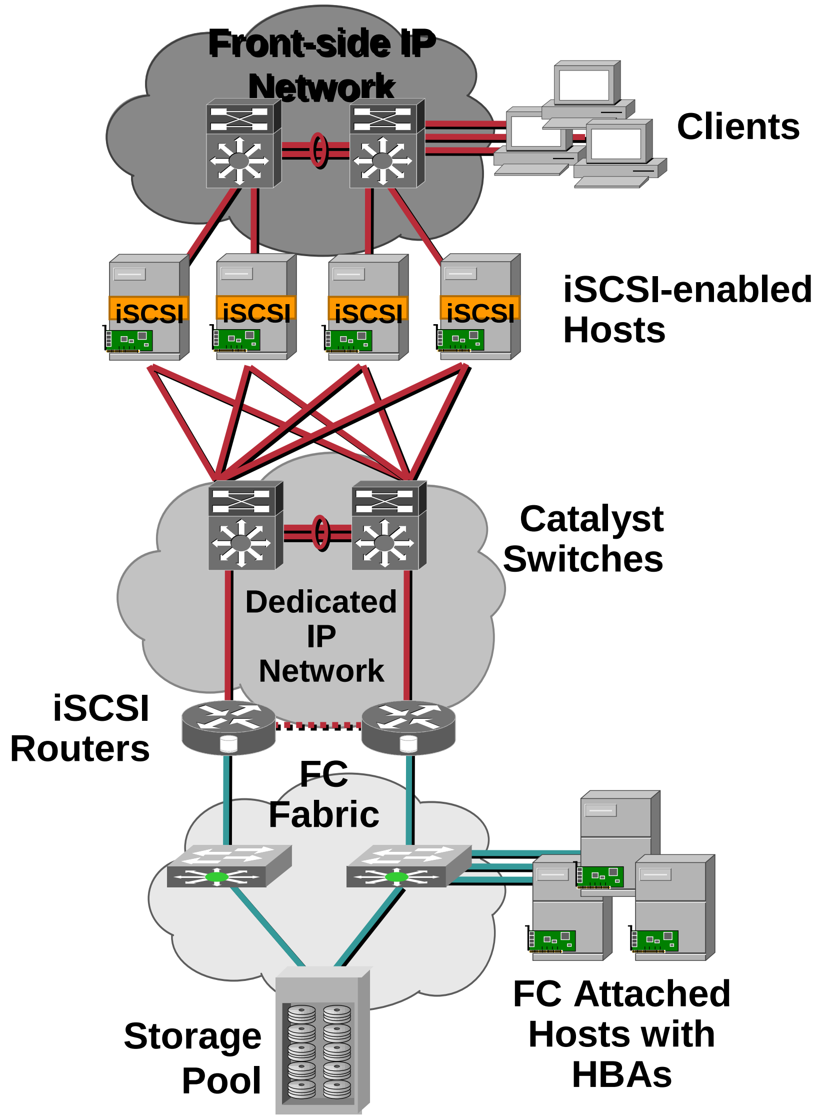

What is FCIP (Fibre Channel over IP)?

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

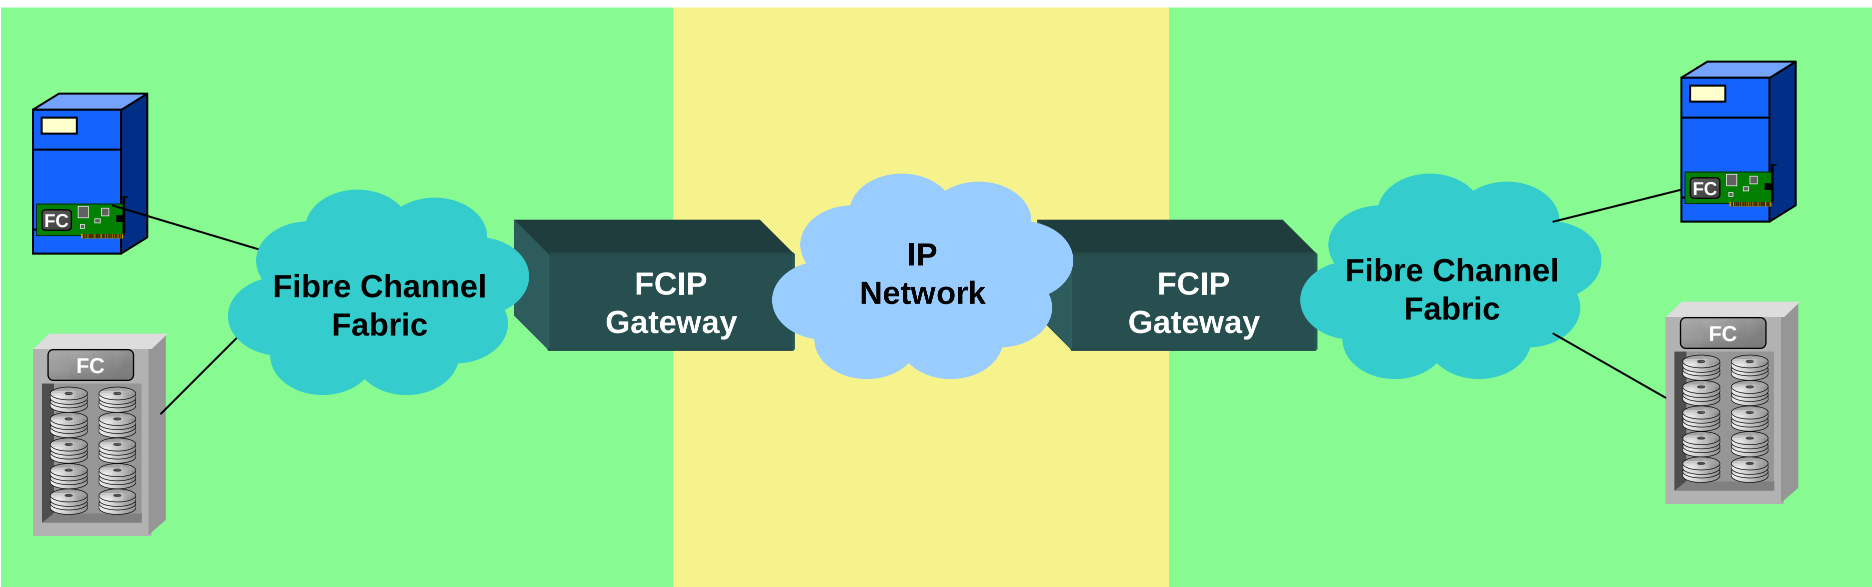

What is FCIP (Fibre Channel over IP)?

Cisco.com

- 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

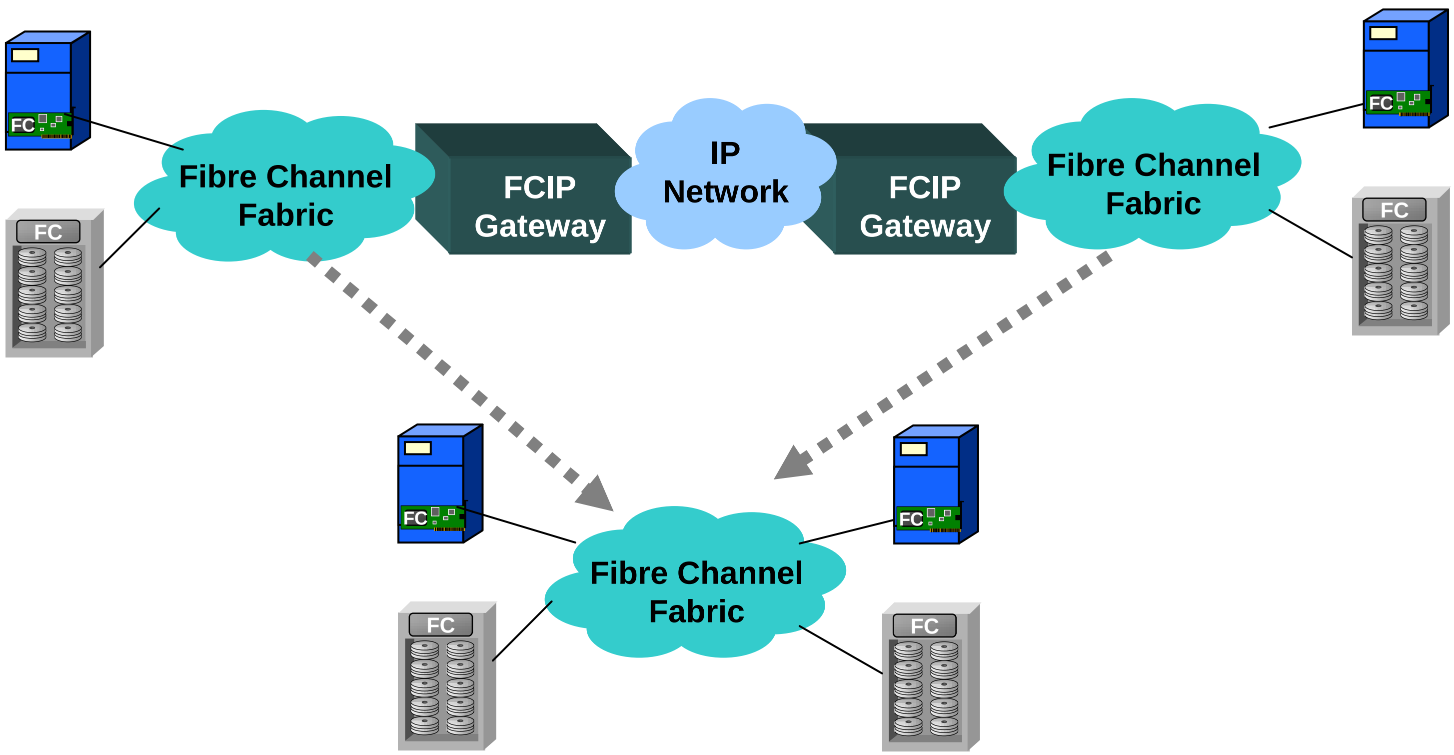

An FCIP Application Topology

- 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

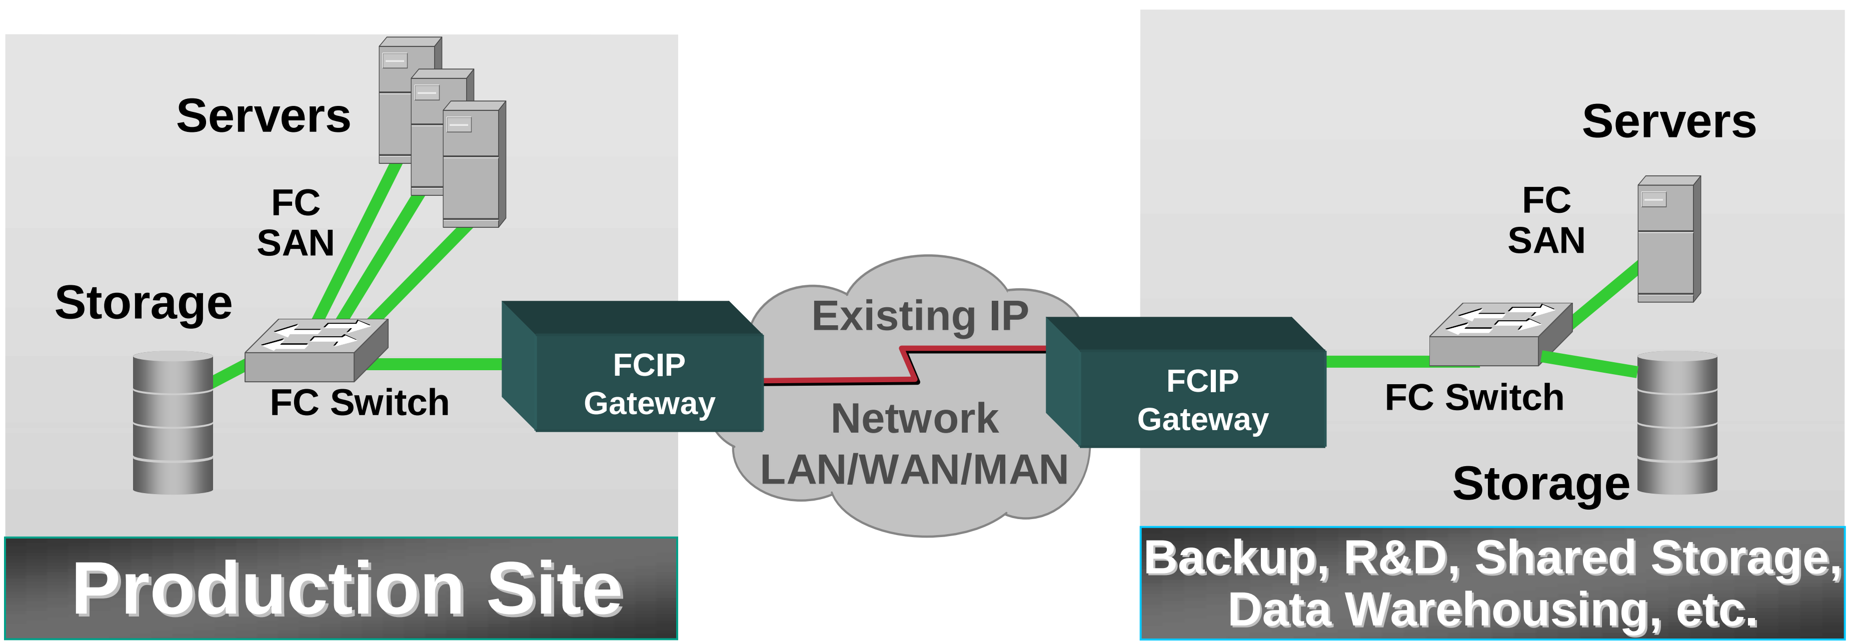

IP Storage Advantages

Cisco.com

Numerous services exist within IP to secure IP storage traffic, many of which are not available in Fibre Channel

- IPSec
- VLANs
- Access Control Lists (ACLs)
- Authentication, Authorization and Accounting
- Firewalls

IP Storage Advantages

Many services exist within IP to ensure performance in an IP SAN as well as protect IP storage traffic from potential bottlenecks

- **QoS**
- **EtherChannel**
- **Gigabit Ethernet**
- **Multi-Protocol Label Switching (MPLS)**

IP Storage Disadvantages

Cisco.com

- IP Storage is an emerging technology
- Extended distance may affect application performance
- Security needs to be considered

Fibre Channel SAN Architectures

Comparing Fabric Designs

Cisco.com

Cascade:

- Very limited inter-switch bandwidth
- No resilience to failures
- Suitable for 2- or 3- switch fabrics where performance and availability are less of a concern than cost

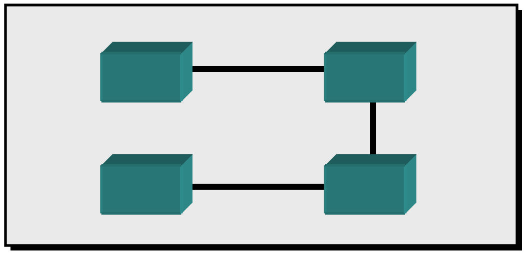

Cascade ring:

- Better performance and availability than cascade design
- Suitable for 3- to 5-switch fabrics with limited scalability, performance, and availability requirements

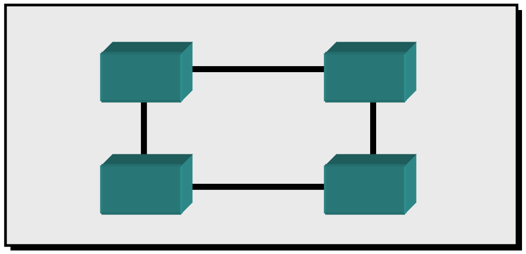

Comparing Fabric Designs

Cisco.com

Mesh:

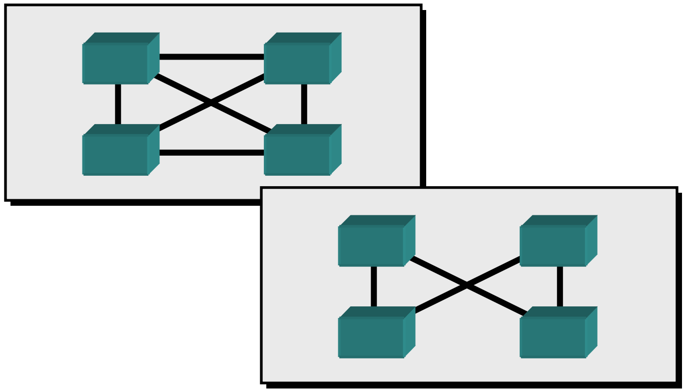

- Highly available
- Performance varies according to full/partial configuration
- Suitable for 4- to 8-switch fabrics with limited scalability requirements

Core-edge:

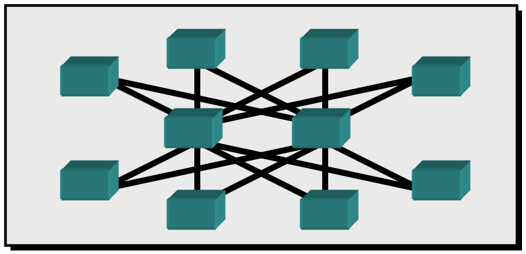

- Highly scalable solution for large fabrics
- Strong performance

Collapsed Fabric Design

Cisco.com

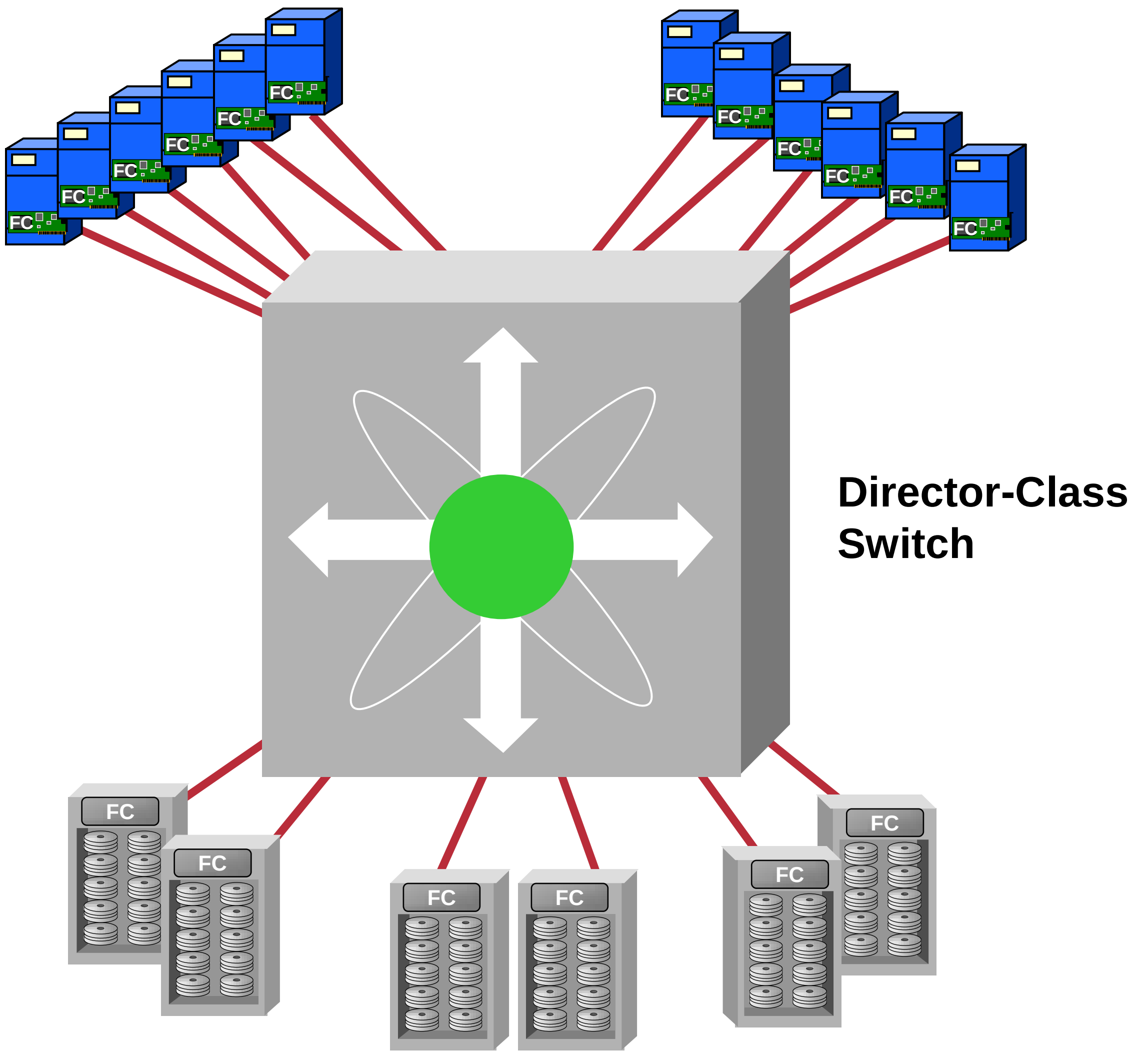

Collapsed Fabric Design

Collapsed Architecture:

- Lack of ISLs means:
  - All purchased ports are available for nodes
  - Increased reliability
  - Simplified management
- Scales easily (hot-swap blade architecture)
- Fixed latency between ports = highest performance
- Single management interface
- Cost-effective for large SANs when ISL ports and management costs are added up
- Not all “director-class switches” are the same

Building Scalable Fabrics

- To get 176 device ports with 4Gb/s of ISL bandwidth, you need…

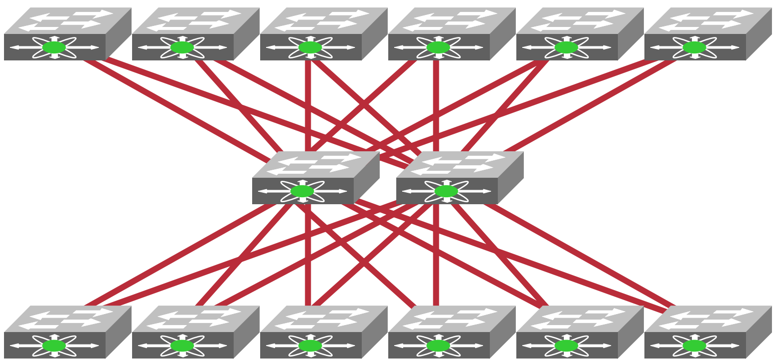

or

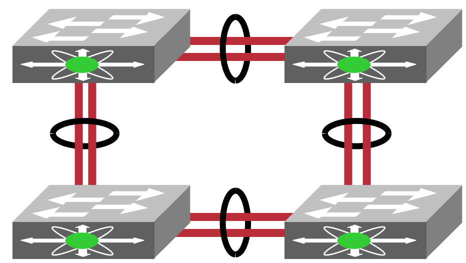

Fourteen 16-port switches:

- 24 ISLs (not load-balanced)
- 78% of ports available
- $1000 per port
- $1545 per node port

Four 48-port MDS 9216 switches:

- 8 ISLs (load-balanced)
- 92% of ports available
- $1000 per port
- $1181 per node port

Cisco's Family of Intelligent Storage Network Devices

Cisco Storage Networking Products

Cisco.com

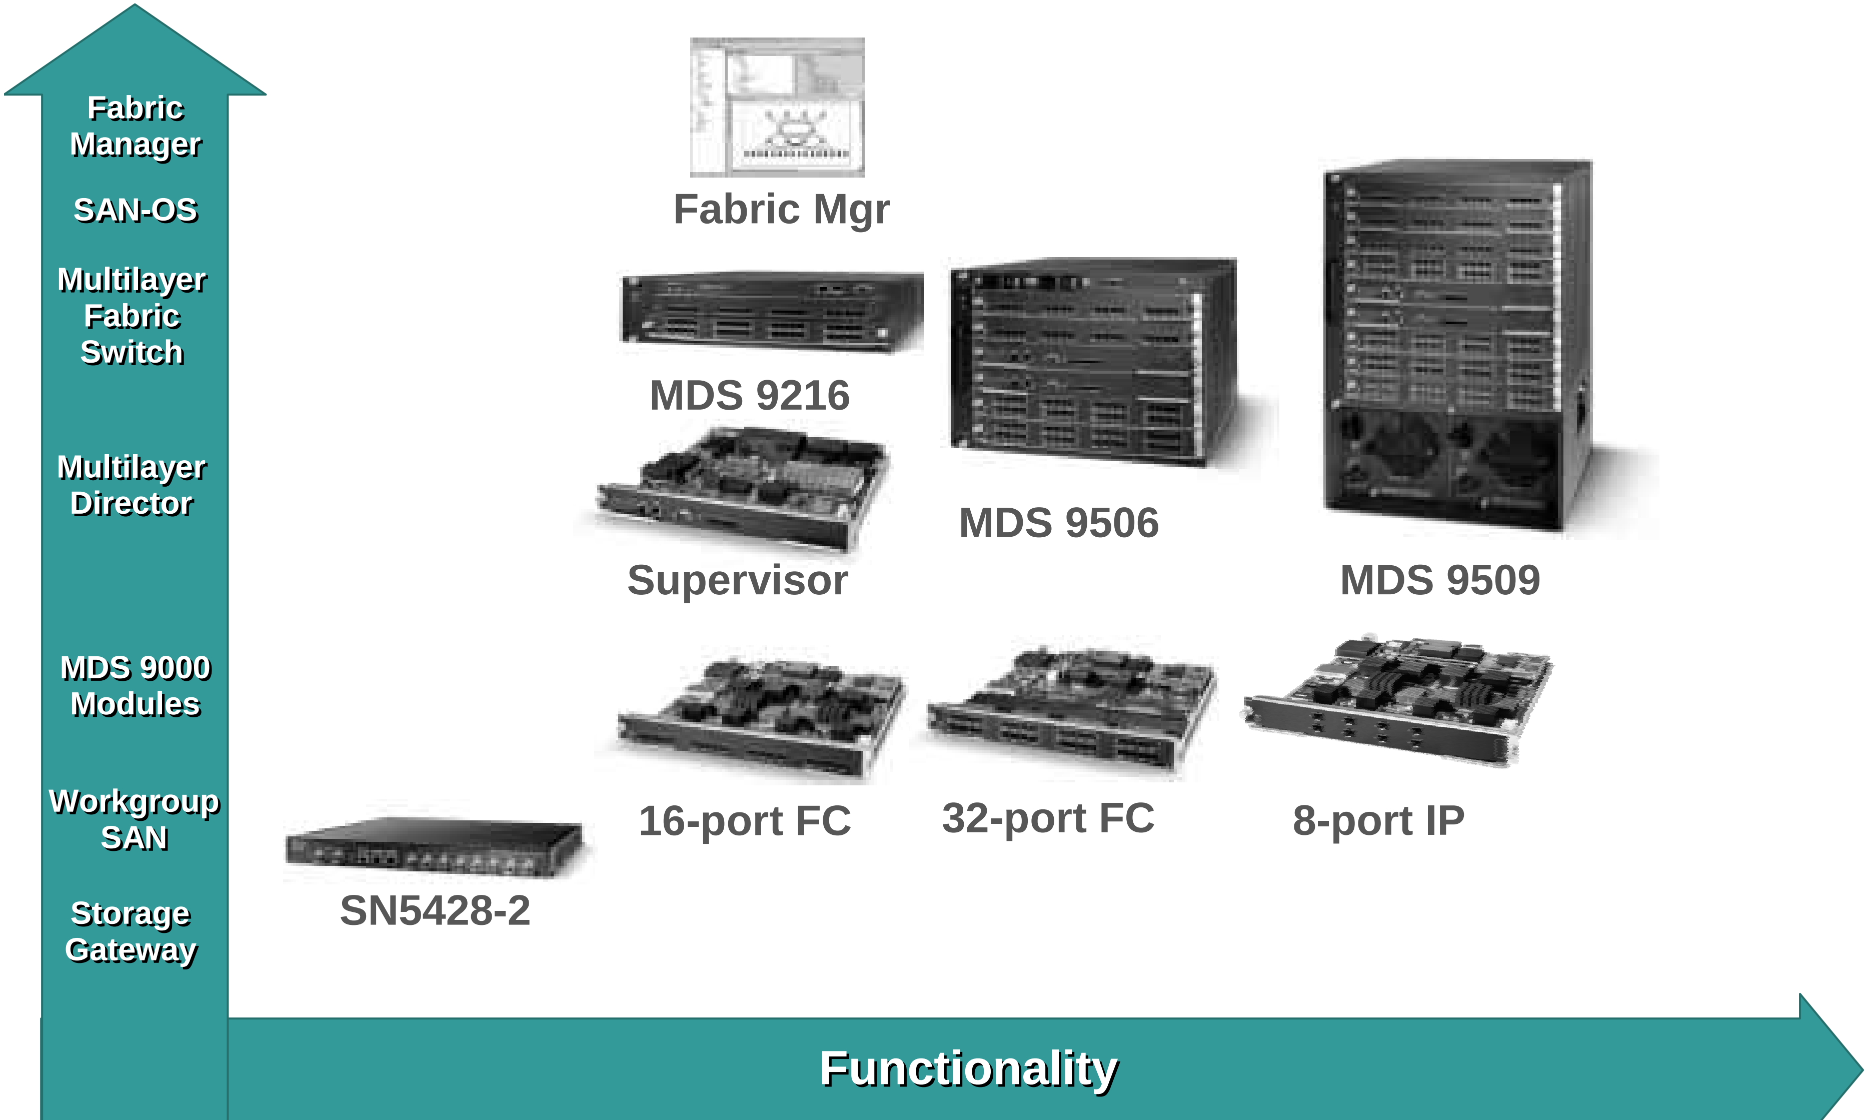

Cisco SN5428-2 Storage Router

SN5428 Sample Application

Block Access via iSCSI for Front-End Servers

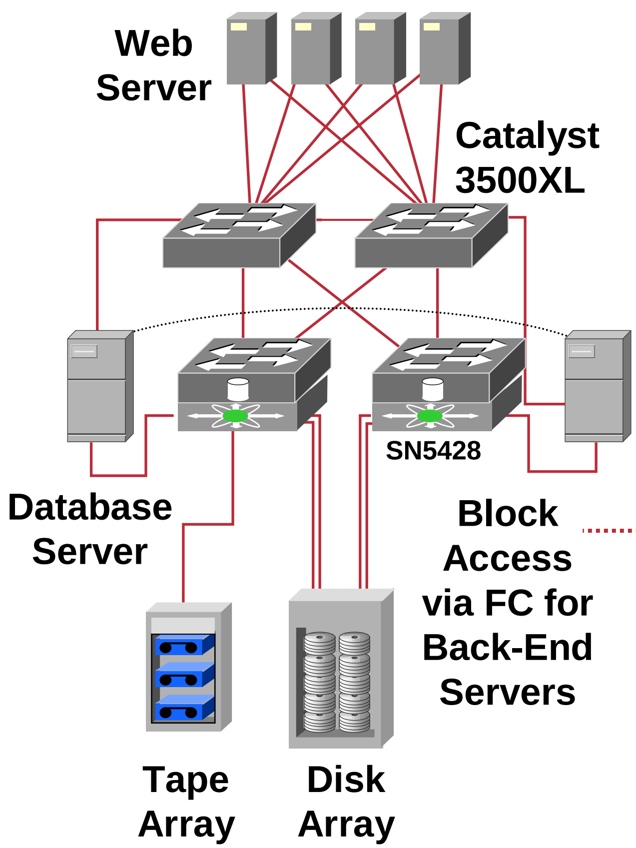

SN5428 Storage Router

- Increased port density and performance
- Integrated 8-port Fibre Channel Switch
- ISCSI + FCIP
- Combined with a Catalyst switch yields cost-effective, small storage network
- SAN+LAN connectivity for application hosts and storage

Introducing the Cisco MDS 9216 Multilayer Fabric Switch

Full-featured fabric switch with expandability

- Base configuration is 16 ports fixed
- Expansion slot allows growth to 48 ports
- 1 / 2 Gbps auto-sensing SFP/LC interfaces
- Compatible with all MDS 9000 Family switching modules
  - 16 and 32-port FC Switching Module
  - 8-port IP Storage Services Module
- Hardware-based services
  - Security services—VSANs, VLANs, ACLs
  - Traffic management—QoS, FCC
  - Enhanced services—PortChannel, load balancing
  - Diagnostics—SPAN, FC Traceroute, FC Ping, Cisco Fabric Analyzer

Introducing the Cisco MDS 9500 Multilayer Director

Redefining director-class storage switching

- Non-blocking fabric—1.44 Tbps
- 1 / 2 Gbps auto-sensing ports—10Gbps ready
- Platform for storage management software
- Hardware-based services
  - Security services—VSANs, VLANs, ACLs
  - Traffic management—QoS, FCC
  - Diagnostics—SPAN, FC Traceroute, Fabric Analyzer
  - Enhanced services—PortChannel, load balancing

Multitransport switch—FC, iSCSI, FCIP

Highly-Scalable MDS 9500 Series Supervisor Module

Integrated crossbar has many benefits

- Investment protection—ability to support new line cards including new transports
- Multiprotocol support in one system
- Highly-scalable system—1.44Tbps

High port density means fewer devices to purchase and manage

- Increase in usable ports due to minimal switch interconnects
- Common equipment amortized over more ports (power supplies, supervisors, chassis)

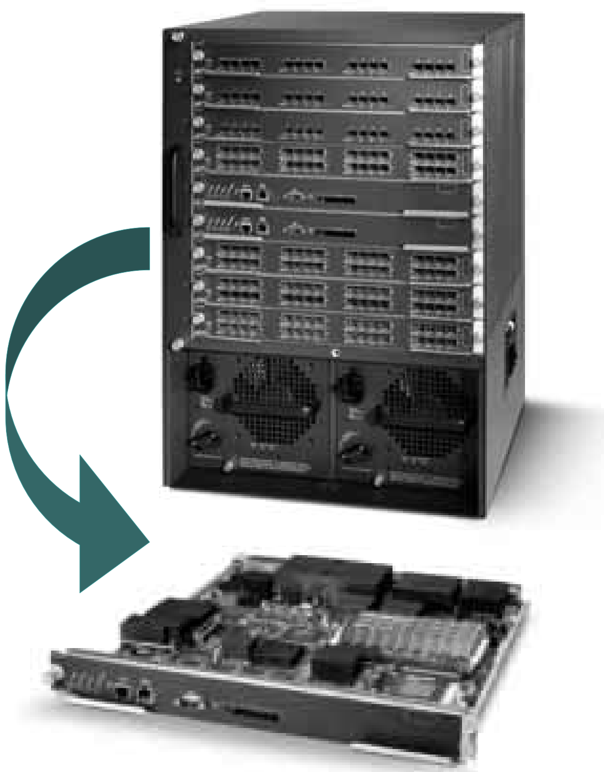

720 Gbps Multiprotocol Crossbar per Supervisor

Cisco MDS 9000 Family Switching Module Summary

Cisco.com

16-Port Fibre Channel

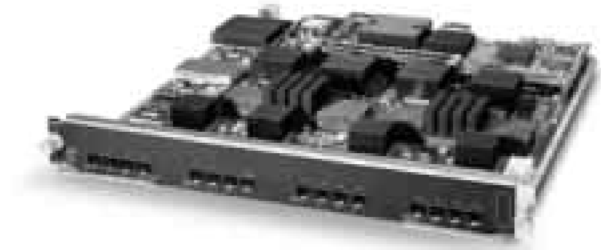

- 16-port 1 / 2-Gbps auto-sensing Fibre Channel (SFP/LC)

32-Port Fibre Channel

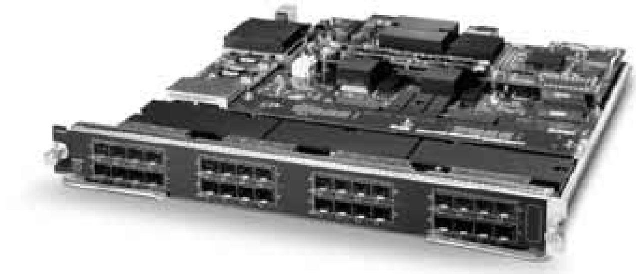

- 32-port 1 / 2-Gbps auto-sensing Fibre Channel (SFP/LC)

8-Port IP Storage Services

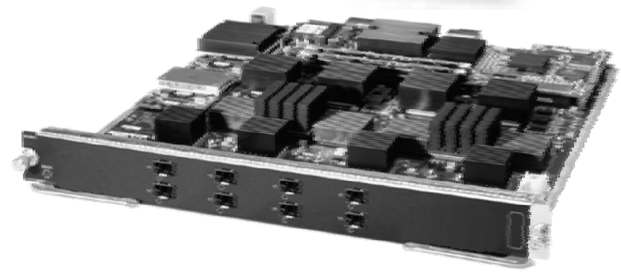

- 8-port 1-Gbps Ethernet with iSCSI and FCIP Gateway functionality (SFP/LC)

16 and 32-Port Fibre Channel Switching Module Features

Cisco.com

Interfaces

- 16 or 32-port 1 / 2 Gbps auto-sensing Fibre Channel (E, F, FL, SD, TE, and TL ports) with SFP/LC optical interfaces

Security Features

- Hardware ACL-based Port Security, Virtual SANs (VSANs), Port Zoning and LUN Zoning

Performance Features

- Up to 80 Gbps fabric bandwidth available per line card
- Up to 255 Buffer Credits per Port
- PortChannel, Multi-Path Load Balancing
- Forward Congestion Control (FCC)
- Quality of Service

MDS 9000 Family
IP Storage Services Module Features

Cisco.com

Interfaces

- 8-port 1 Gbps Ethernet with SFP/LC optical interfaces

iSCSI Feature Highlights

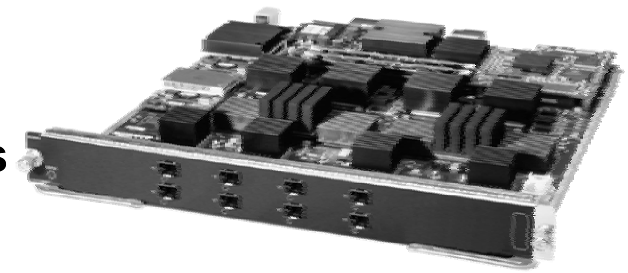

- iSCSI Initiator-Fibre Channel Target
- Transparent view of all allowed hosts/targets
- iSCSI to Fibre Channel zone mapping

FCIP

- Up to 3 FCIP tunnels per port on all ports (24 tunnels per line card)

Fibre Channel Features

- All standard Fibre Channel line card features (interfaces N/A)
- Leverages Fibre Channel interfaces on other switch modules

Multiprotocol Flexibility

- iSCSI and FCIP on each port concurrently – software configurable
- Investment protection – seamless migration to new technologies

Cisco Fabric Manager

Simplifies Management of Multiple Switches and Fabrics

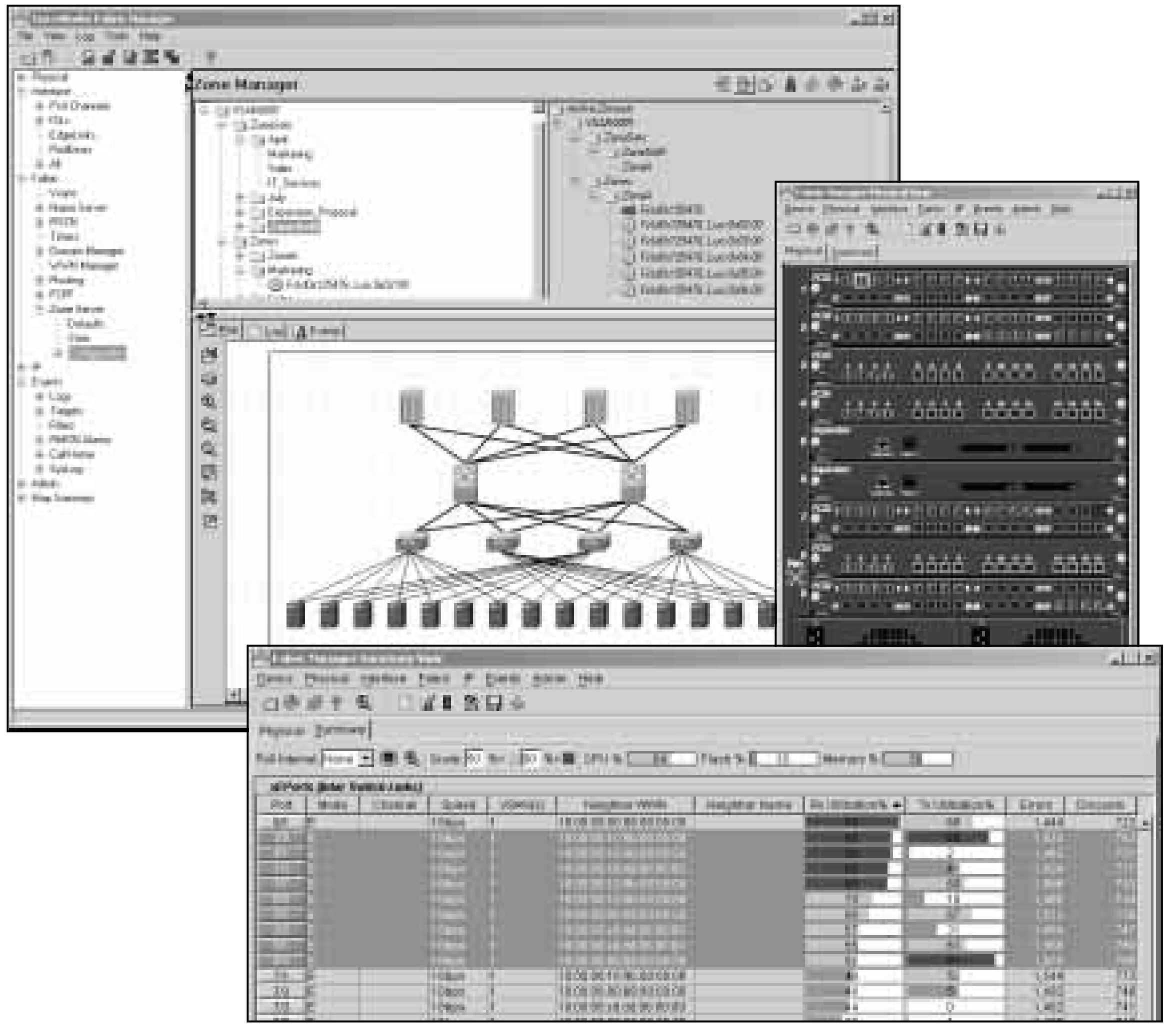

Switch-embedded Java-based Application

Discovery and Topology Mapping

Multiple Views

- **Fabric View**
- **Summary View**
- **Physical View**

Configuration

Monitoring and Alerts

Network Diagnostics

Security

- **SNMPv3**
- **SSH**
- **RBAC**

Integration with CiscoWorks

Resource Manager Essentials (RME)

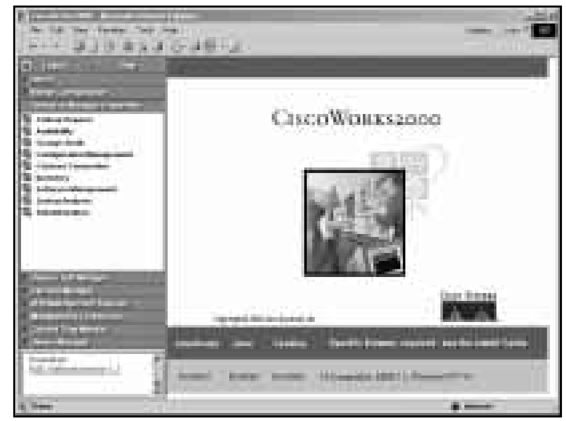

- Software image and configuration management
  - Scheduled downloads with rollback
  - Configuration file editing and difference checking
- Archive switch software images and configuration files
- Provide change auditing and syslog analysis
- Integrates with CCO to simplify support

Cisco Systems®

Empowering the Internet Generation