Disk Systems

- DAS
  - ATA
  - SCSI
  - RAID
- NAS
  - NAS architecture
    - NAS appliances
  - SAN architecture
    - FC
    - iSCSI

NAS architecture

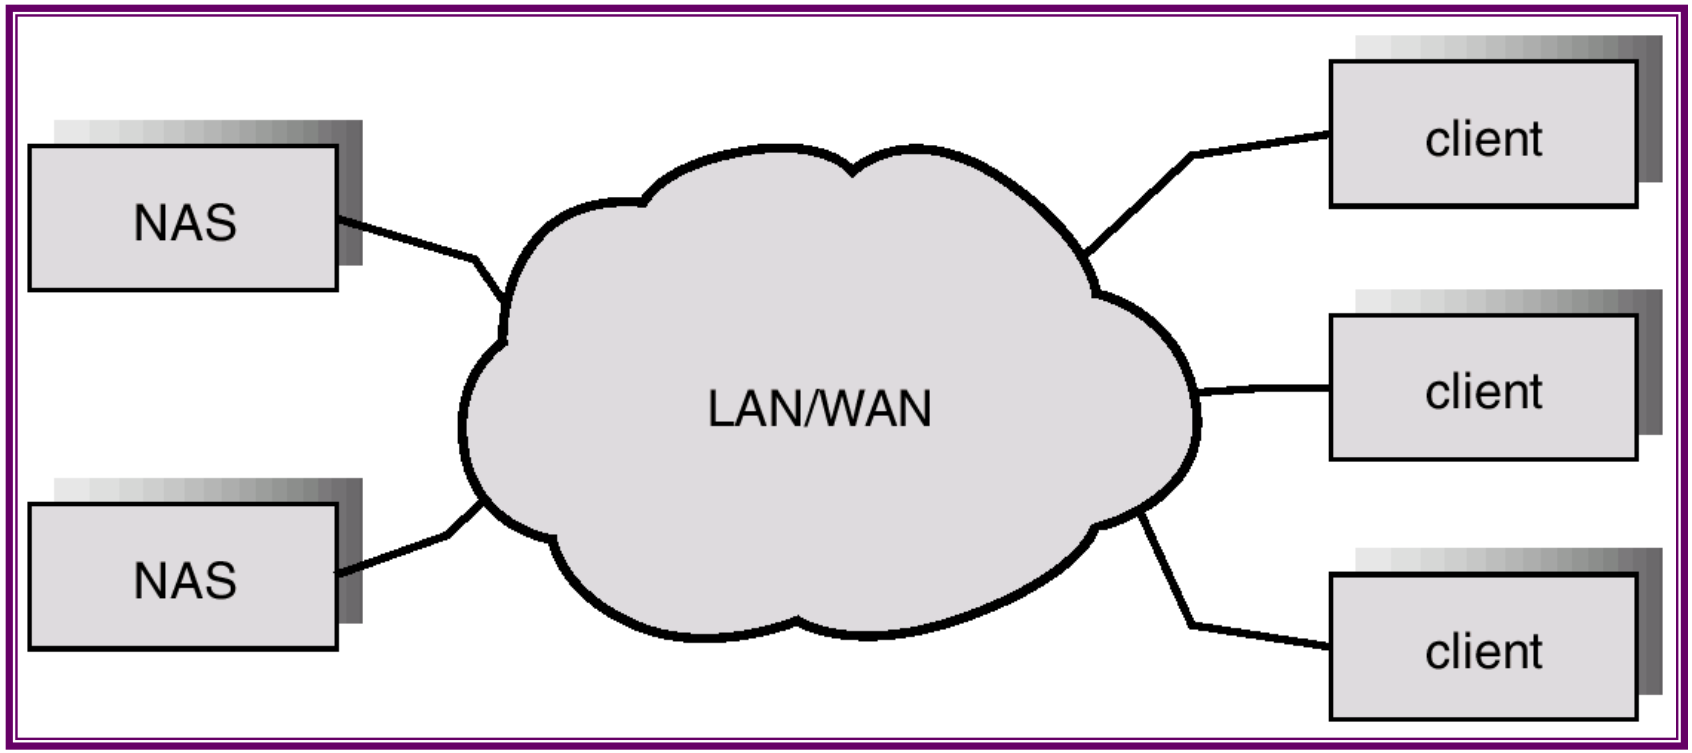

NAS features

- **NAS** is a **file- based** storage architecture.
- **Clients** and **NAS servers** communicate **over** a **standard LAN.**
- The NAS server contains both disk storage and **file system**
- Communicate with the NAS server using standard protocols –
  - NFS (Network File System) or
  - CIFS (Common Internet File System).
- NAS, **ideal** for applications that need to share data at a file level
  - email
  - web hosting
  - database
- NAS is fast and simple to install and manage.
- Additional NAS resources can be easily added to the network.

NAS Appliance

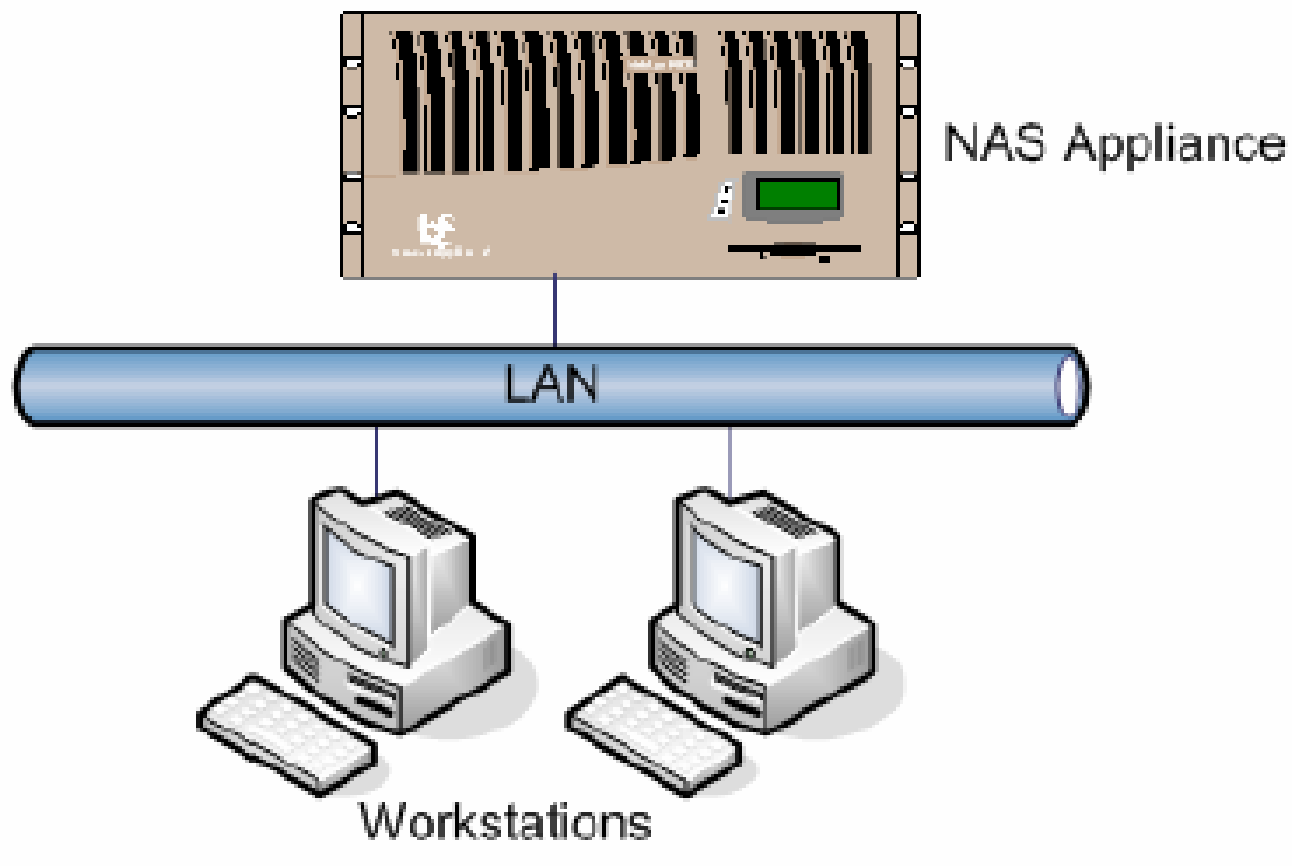

Figure 1.3: A simple NAS deployment.

SAN Architecture

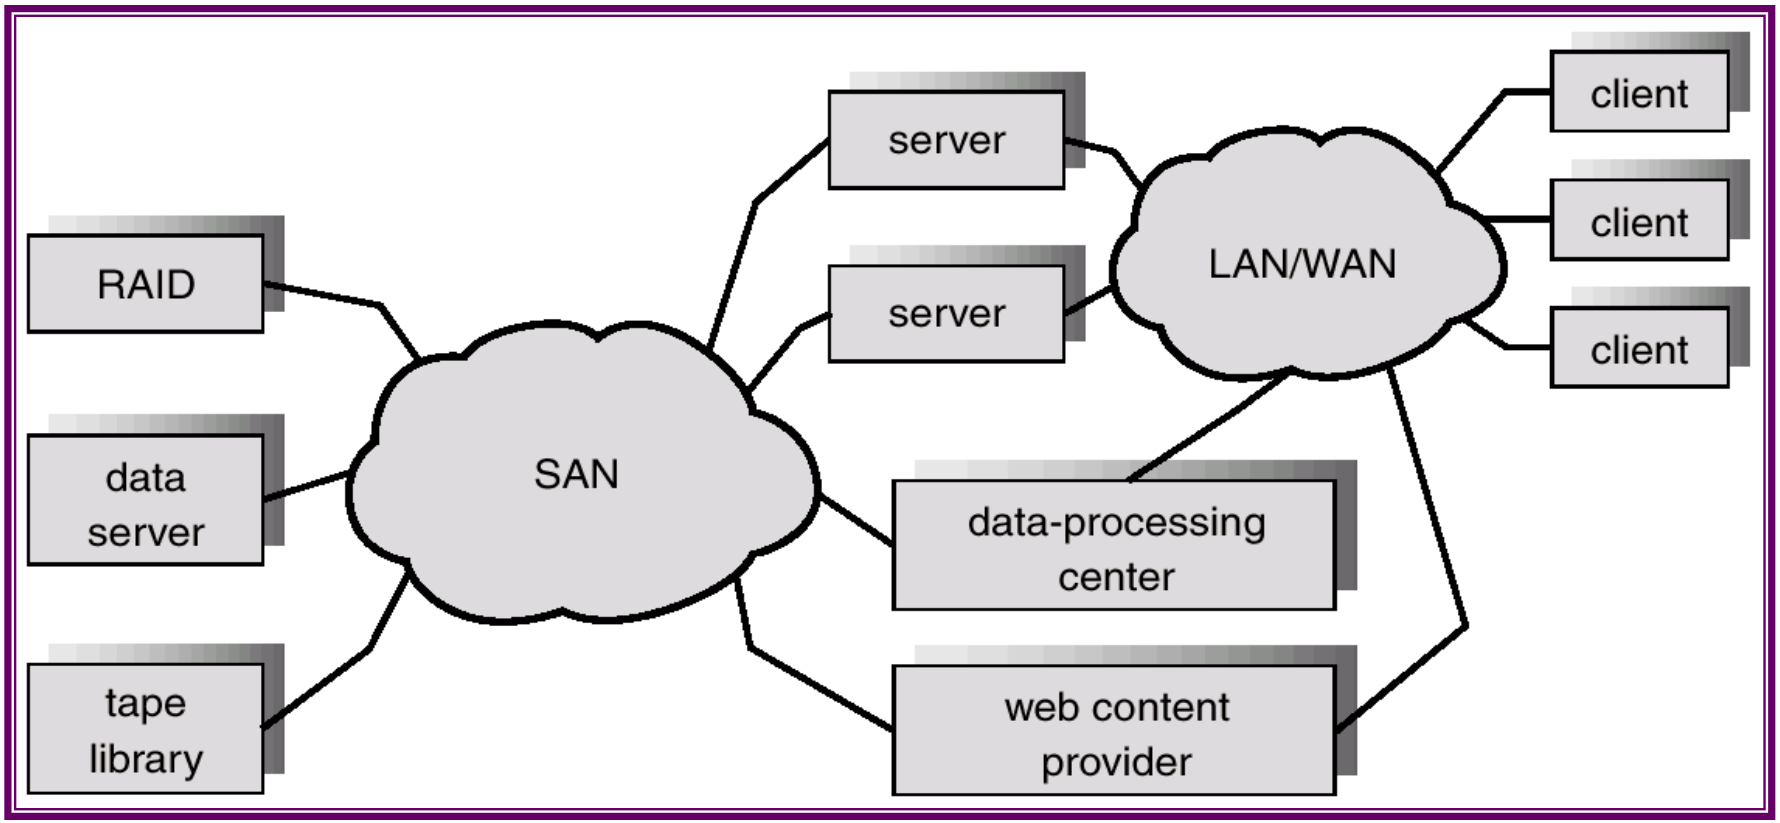

SAN

- FC based SAN
- iSCSI
- FCIP
- iFCP

FC SAN

- Today, FC is the predominant architecture for interconnecting **shared storage devices.**
- Speed—4Gbps (1-2-4-10-12 Gbps)
- Full duplex
- FC SANs support as many as 16 million devices
- FC supports **cable lengths** as long as 10-100km

Typical FC SAN

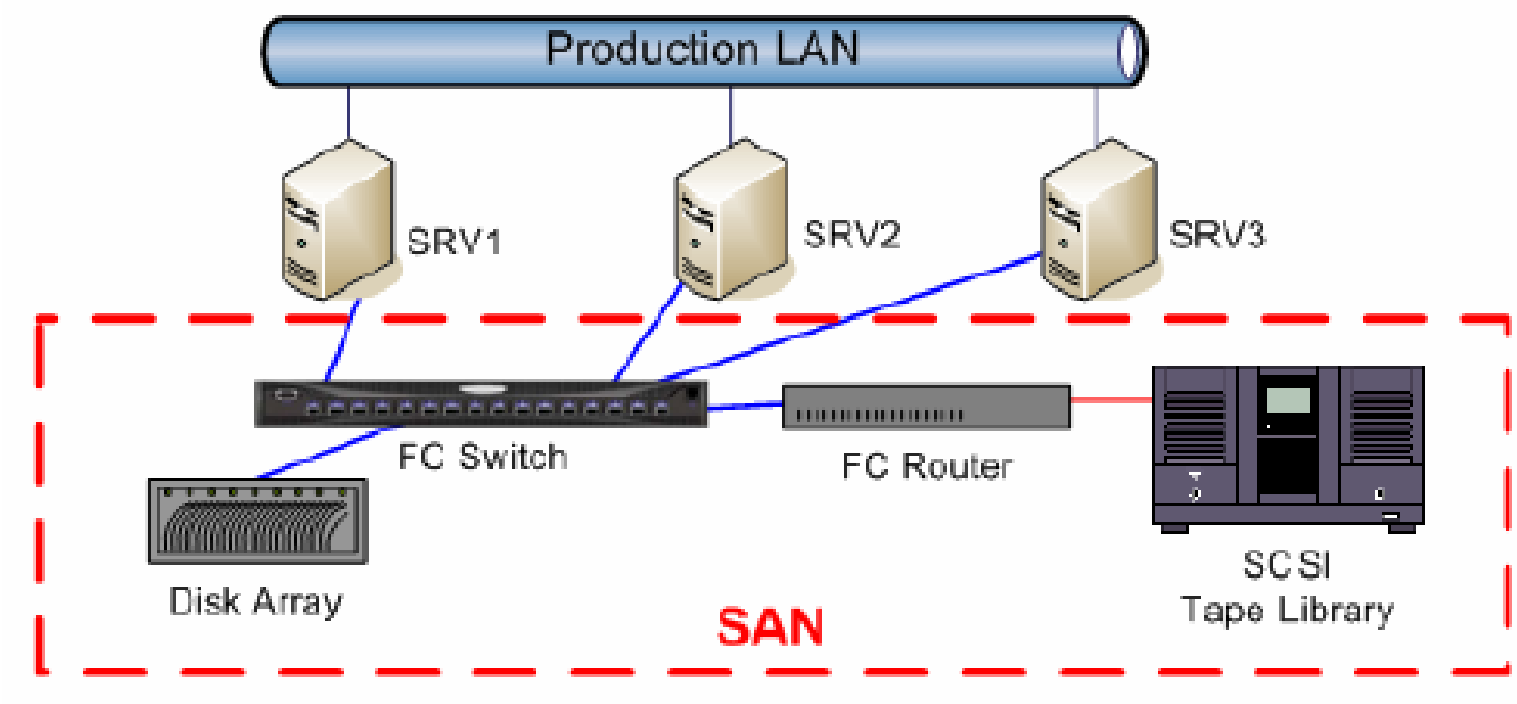

Figure 1.7: A SAN that consists of a switch, router, disk array, and tape library.

FC SAN ideal for cluster architecture

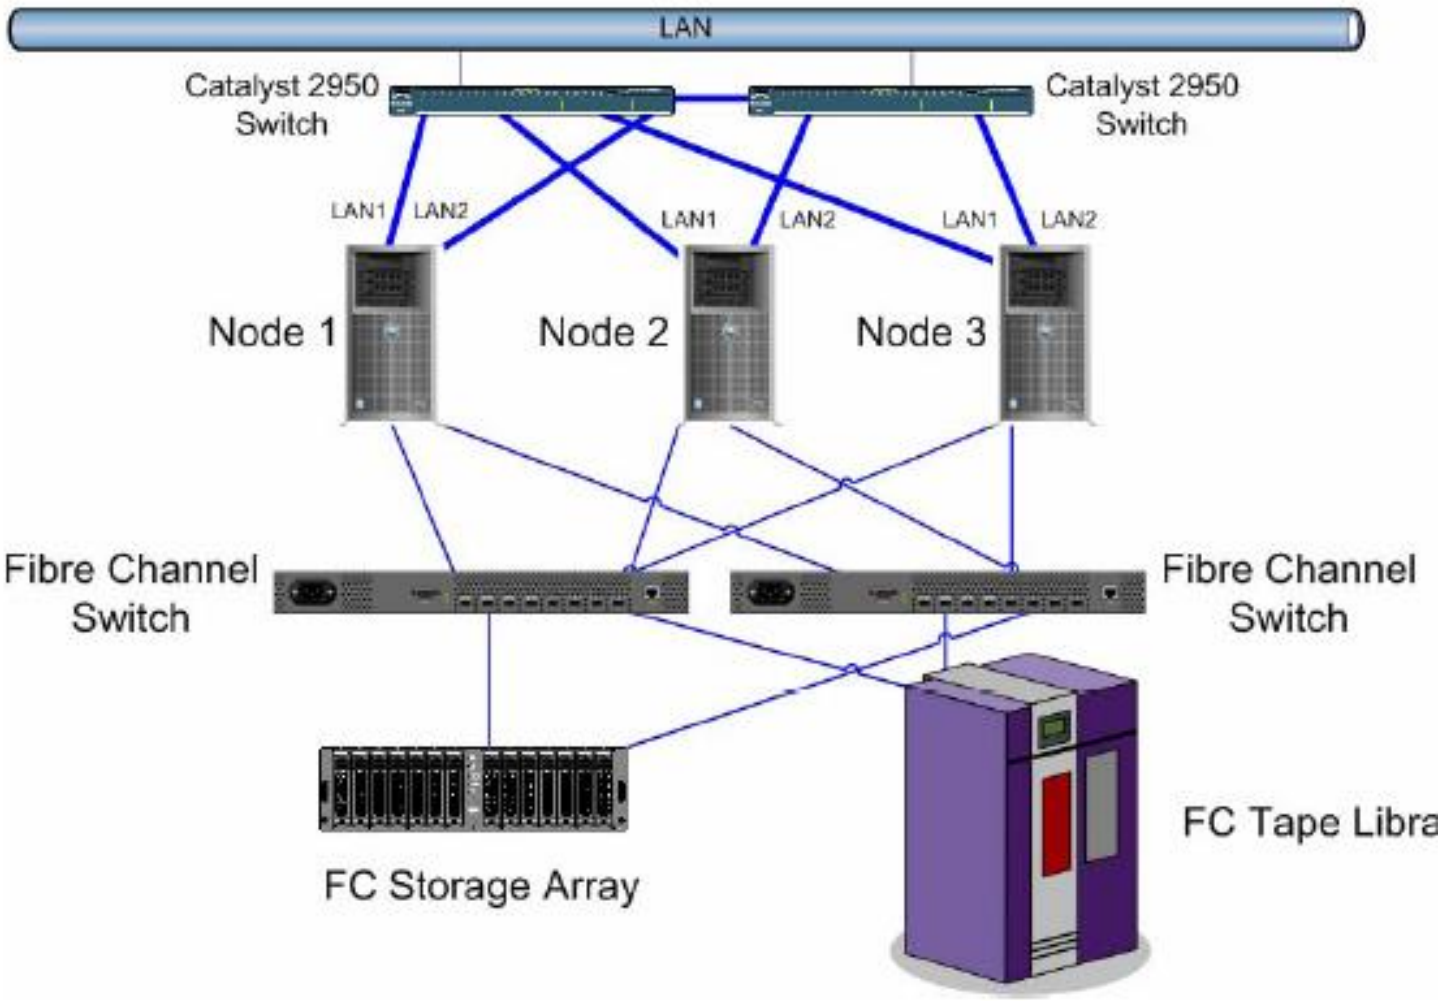

Figure 1.9: An example of a high-availability clustering architecture.

2.2.6 Fibre Channel adapter cable

The LC-SC adapter cable attaches to the end of an LC-LC cable to support SC device connections. A combination of one LC/LC fiber cable and one LC/SC adapter cable is required for each connection. This is used to connect from some of the older 1-Gbps devices to a 2-Gbps capable and LC interface-based SAN.

Shown in Figure 2-9 is a Fibre Channel adapter cable.

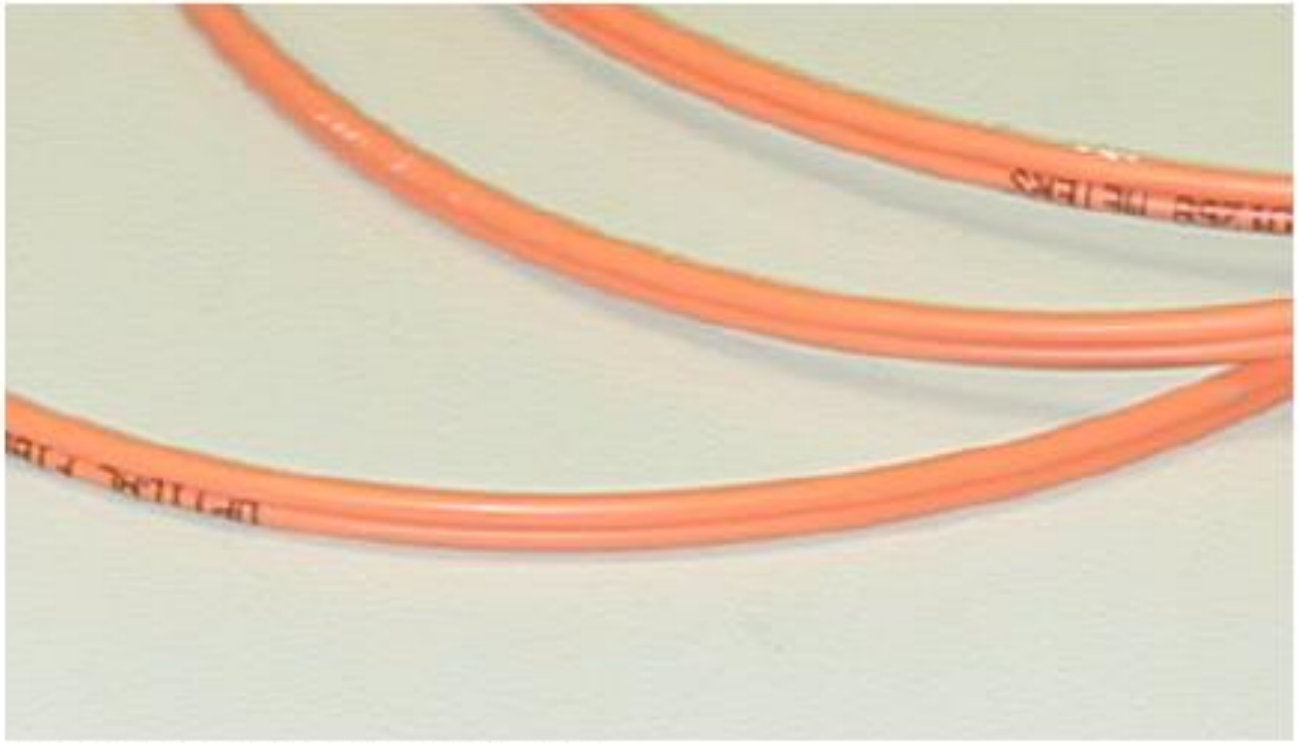

Figure 2-9 Fibre Channel adapter cable

Figure 2-10 Host Bus Adapter

Various cables may be supported by the HBAs, for example:

- Glass fiber
  - Single-mode
  - Multi-mode
- Copper
  - Twisted pair
  - Coaxial

There are several manufacturers of HBAs and an important consideration when planning a SAN, is the choice of HBAs. Some HBAs may have interoperability problems with some other FIbre Channel components.

A server or storage device may have one HBA or it may have many. Depending upon the particular configuration of the SAN, if there are more than one, they might all be identical, or they could be of different types.

The adapters in storage arrays are usually determined by the manufacturer. Factors influencing the choice of HBAs in servers are dealt with in Chapter 4, “SAN design considerations” on page 221.

FC Topology

- There are 3 major Fibre Channel topologies:
- **1. Point-to-Point (FC-P2P).** Two devices are connected back to back. This is the simplest topology, with limited connectivity.

- **2. Arbitrated Loop (FC-AL).** In this design, all devices are in a loop or ring, similar to token ring networking. Adding or removing a device from the loop causes all activity on the loop to be interrupted. The failure of one device causes a break in the ring. **Fibre Channel hubs** exist to connect multiple devices together, or a simple point-to-point connection can be made. An arbitrated loop with two devices degenerates to point-to-point topology.

- **3. Switched Fabric (FC-SW).** All devices are connected to Fibre Channel switches, similar conceptually to modern Ethernet implementations. The switches manage the state of the fabric, providing **optimized interconnections** and also security.

3.1.2 Ports

The basic building block of fibre channel is the port. The following lists the various fibre channel port types and their purpose in switches, servers, and storage.

These are the types of Fibre Channel ports that you are likely to encounter:

- *E_Port* is an expansion port. A port is designated an E_Port when it is used as an inter switch expansion port (ISL) to connect to the E_Port of another switch, to enlarge the switch fabric.
- *F_Port* is a fabric port that is not loop capable. It is used to connect an N_Port point-to-point to a switch.
- *FL_Port* is a fabric port that is loop capable. It is used to connect NL_Ports to the switch in a public loop configuration.
- *G_Port* is a generic port that can operate as either an E_Port or an F_Port. A port is defined as a G_Port after it is connected but has not received response to *loop* initialization or has not yet completed the *link* initialization procedure with the adjacent Fibre Channel device.
- *L_Port* is a loop capable node or switch port.
- *U_Port* is a universal port. A more generic switch port than a G_Port, it can operate as either an E_Port, F_Port, or FL_Port. A port is defined as a U_Port when it is not connected or has not yet assumed a specific function in the fabric.
- *N_Port* is a node port that is not loop capable. It is used to connect an equipment port to the fabric.
- *NL_Port* is a node port that is loop capable. It is used to connect an equipment port to the fabric in a loop configuration through an L_Port or FL_Port.
- *MTx_Port* is a CNT port used as a mirror for viewing the transmit stream of the port to be diagnosed.
- *MRx_Port* is a CNT port used as a mirror for viewing the receive stream of the port to be diagnosed.
- *SD_Port* is a Cisco SPAN port used for mirroring another port for diagnostic purposes.
- *T_Port* was used previously by CNT as a mechanism of connecting directors together. This has been largely replaced by the E_Port.

FC: Point to Point (only two devices)

From Computer Desktop Encyclopedia
© 2001 The Computer Language Co. Inc.

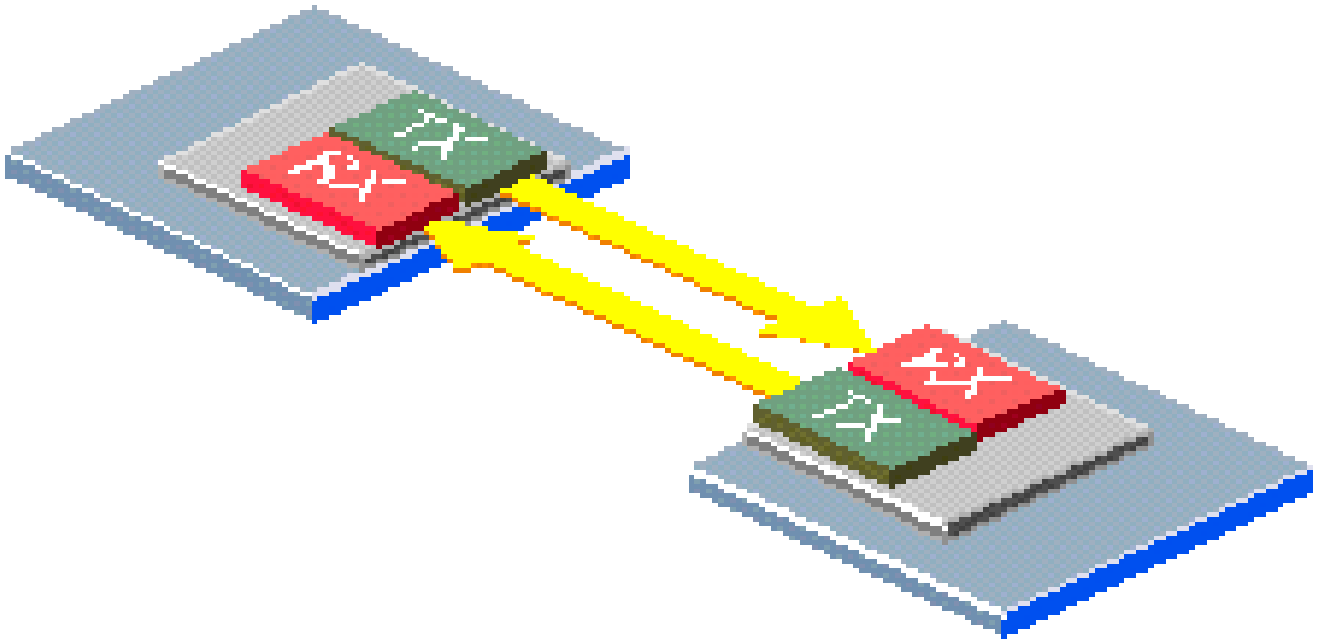

3.1.4 Point-to-point

A *point-to-point* connection is the simplest topology. It is used when there are exactly two nodes, and future expansion is not predicted. There is no sharing of the media, which allows the devices to use the total bandwidth of the link. A simple link initialization is needed before communications can begin.

Fibre Channel is a *full duplex* protocol, which means both paths transmit data simultaneously. Fibre Channel connections based on the 1-Gb standard are able to transmit at 100 MBps and receive at 100 MBps simultaneously. For Fibre Channel connections based on the 2-Gb standard, they can transmit at 200 MBps and receive at 200 MBps simultaneously. This will extend to 4.25 gigabits (Gbps) and 10 Gbps technologies as well.

Illustrated in Figure 3-4 is a simple point-to-point connection.

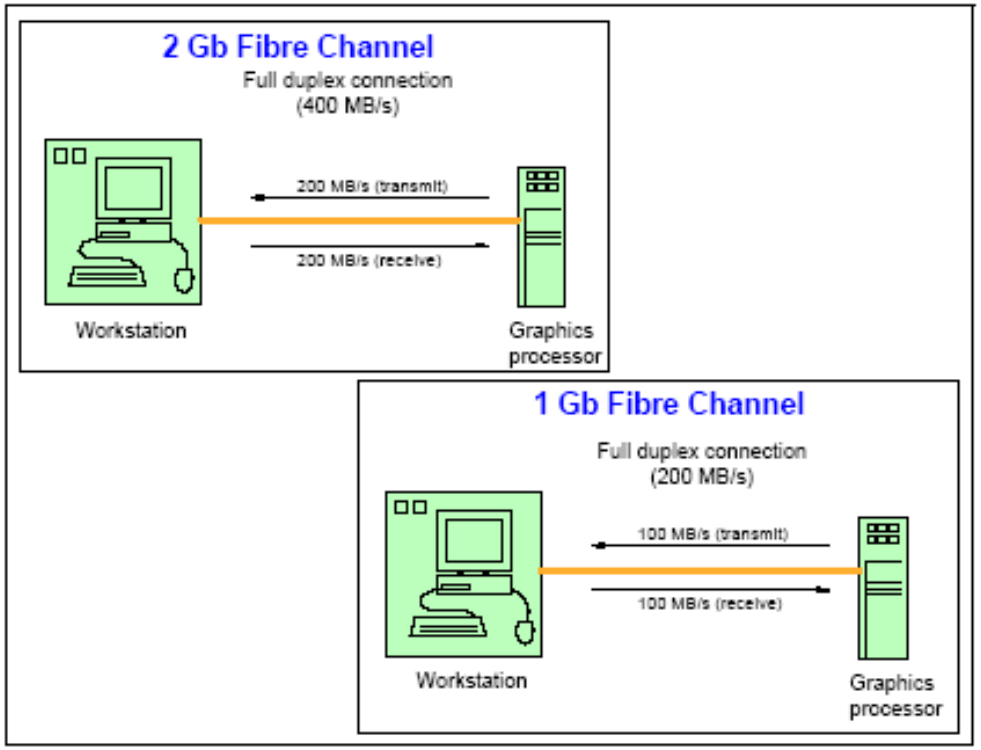

Figure 3-4 Point-to-point

Arbitrated loop

- The second topology is Fibre Channel Arbitrated Loop (FC-AL).
- FC-AL is more useful for storage applications.
- It is a loop of up to **126 nodes** (NL_Ports) that is **managed** as a **shared bus**. Traffic flows in one direction, carrying data frames and primitives around the loop with a total bandwidth of **200 MBps**, or **100 MBps** for a loop based on 1-Gbps technology.

- Using **arbitration protocol**, a single connection is established between a sender and a receiver, and a data frame is transferred around the loop. When the communication comes to an end between the two connected ports, the loop becomes available for arbitration and a new connection may be established.

- Loops can be configured with hubs to make connection management easier. A **distance of up to 10 km** is supported by the Fibre Channel standard for both of these configurations. However, latency on the arbitrated loop configuration is affected by the loop size.

FC-AL

From Computer Desktop Encyclopedia
© 2001 The Computer Language Co. Inc.

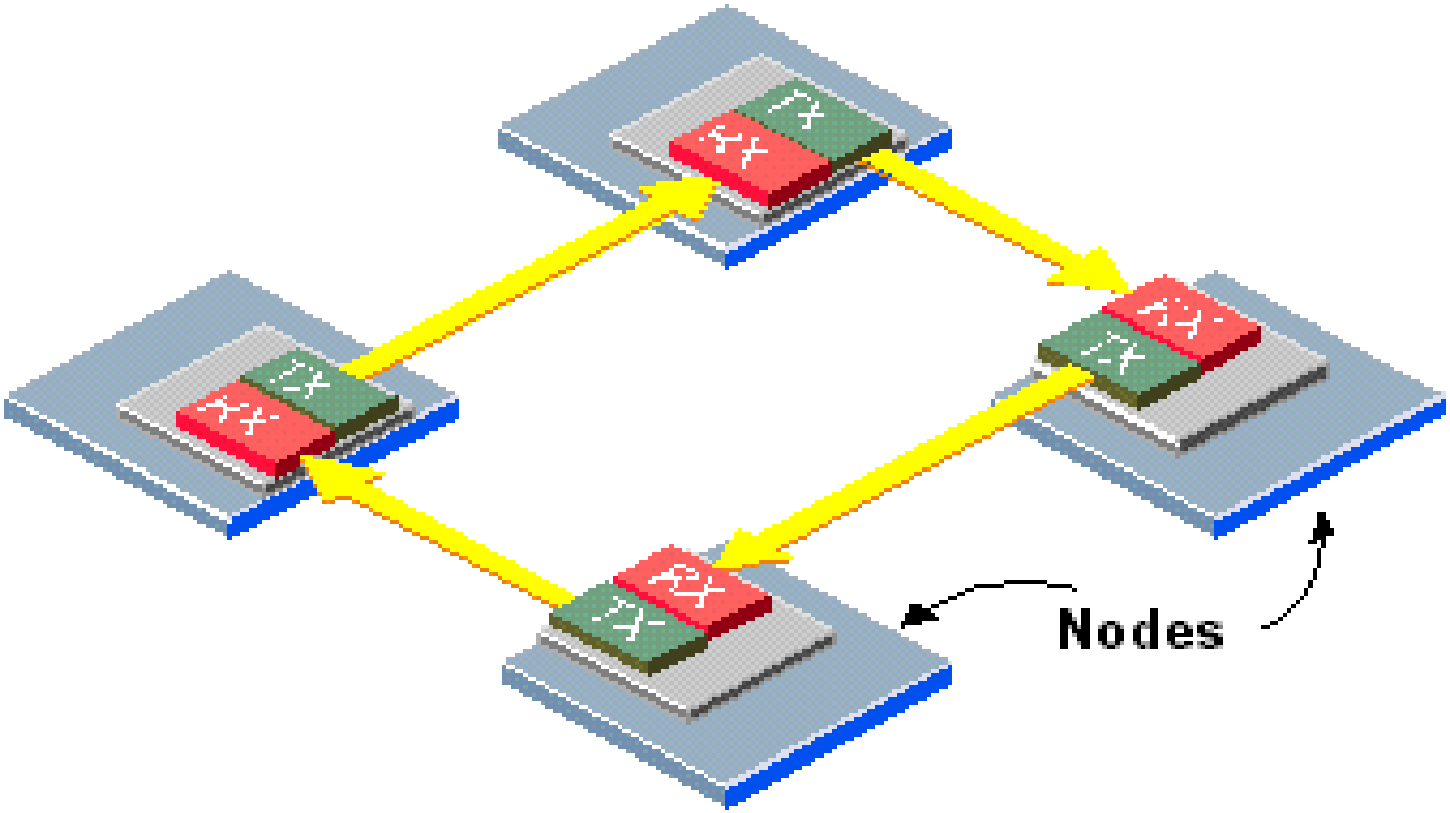

Arbitrated loop

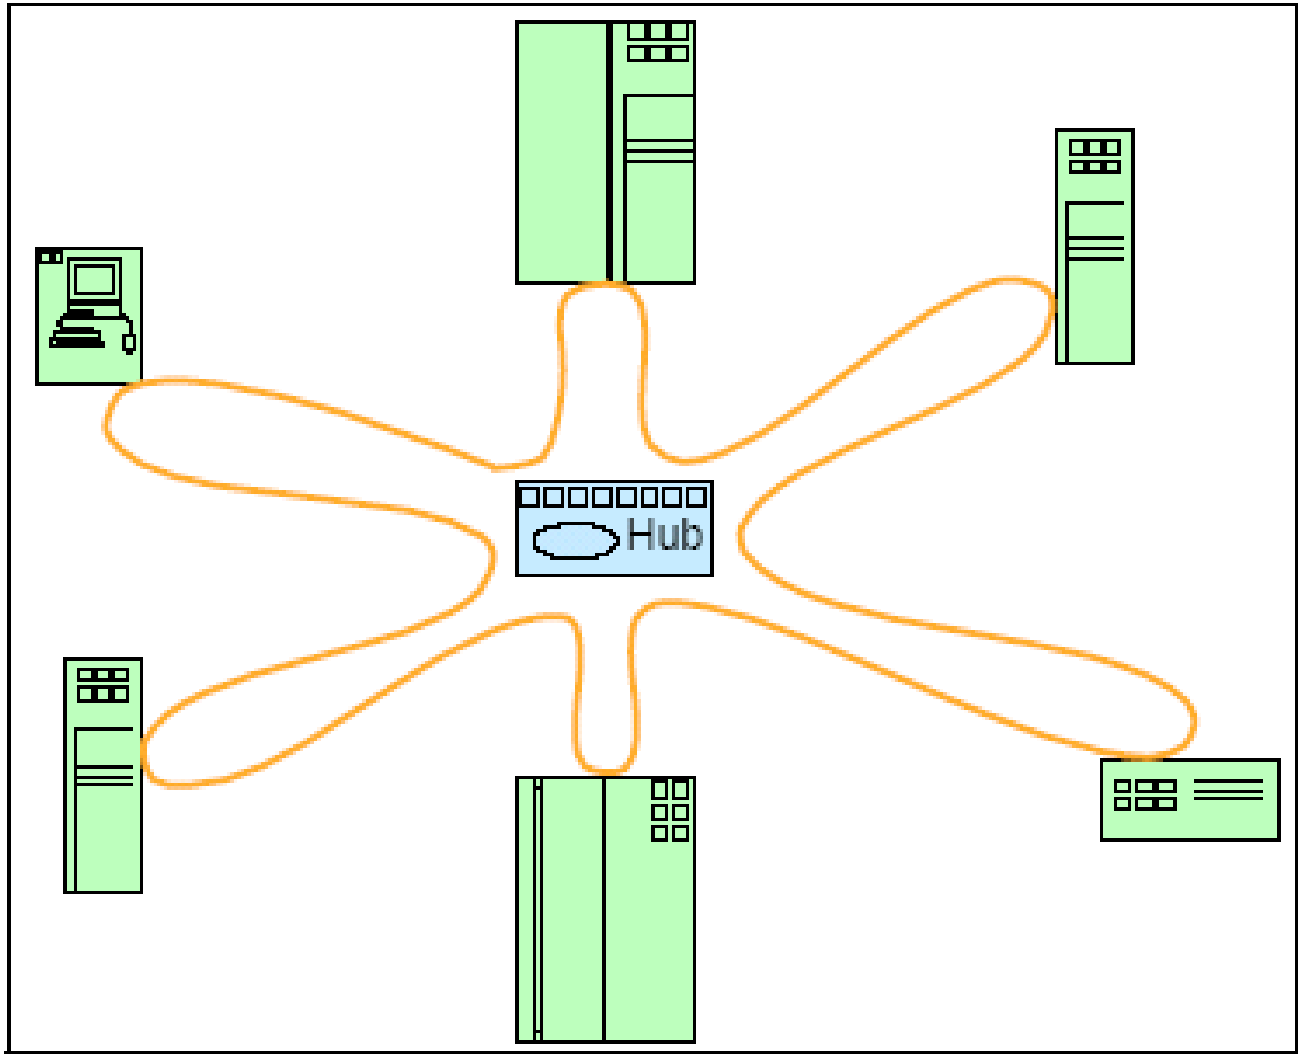

Figure 3-5 Arbitrated loop

FC-AL: Protocols

- **Loop protocols**
- To support the shared behavior of the arbitrated loop, a number of loop-specific protocols are used.

- **These protocols are used to**:
  - **Initialize the loop** and **assign addresses.**
  - **Arbitrate** for access to the loop.
  - **Open** a **loop circuit** with another port in the loop.
  - **Close** a loop circuit when two ports have completed their use of the loop.
  - Implement the access fairness mechanism to ensure that each port has an opportunity to access the loop.

FC-SW

From Computer Desktop Encyclopedia
© 2001 The Computer Language Co. Inc.

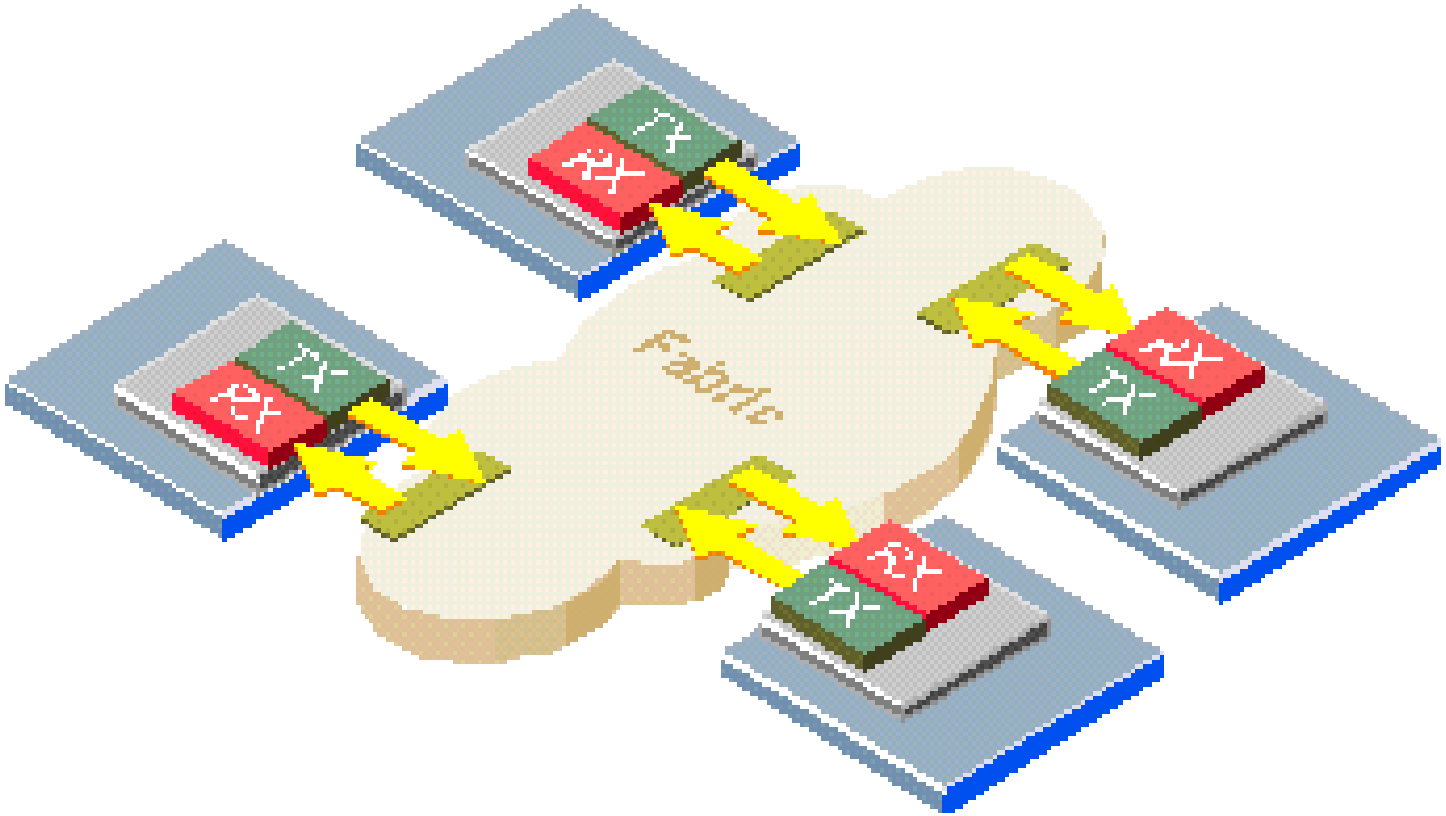

FC-SW: Example

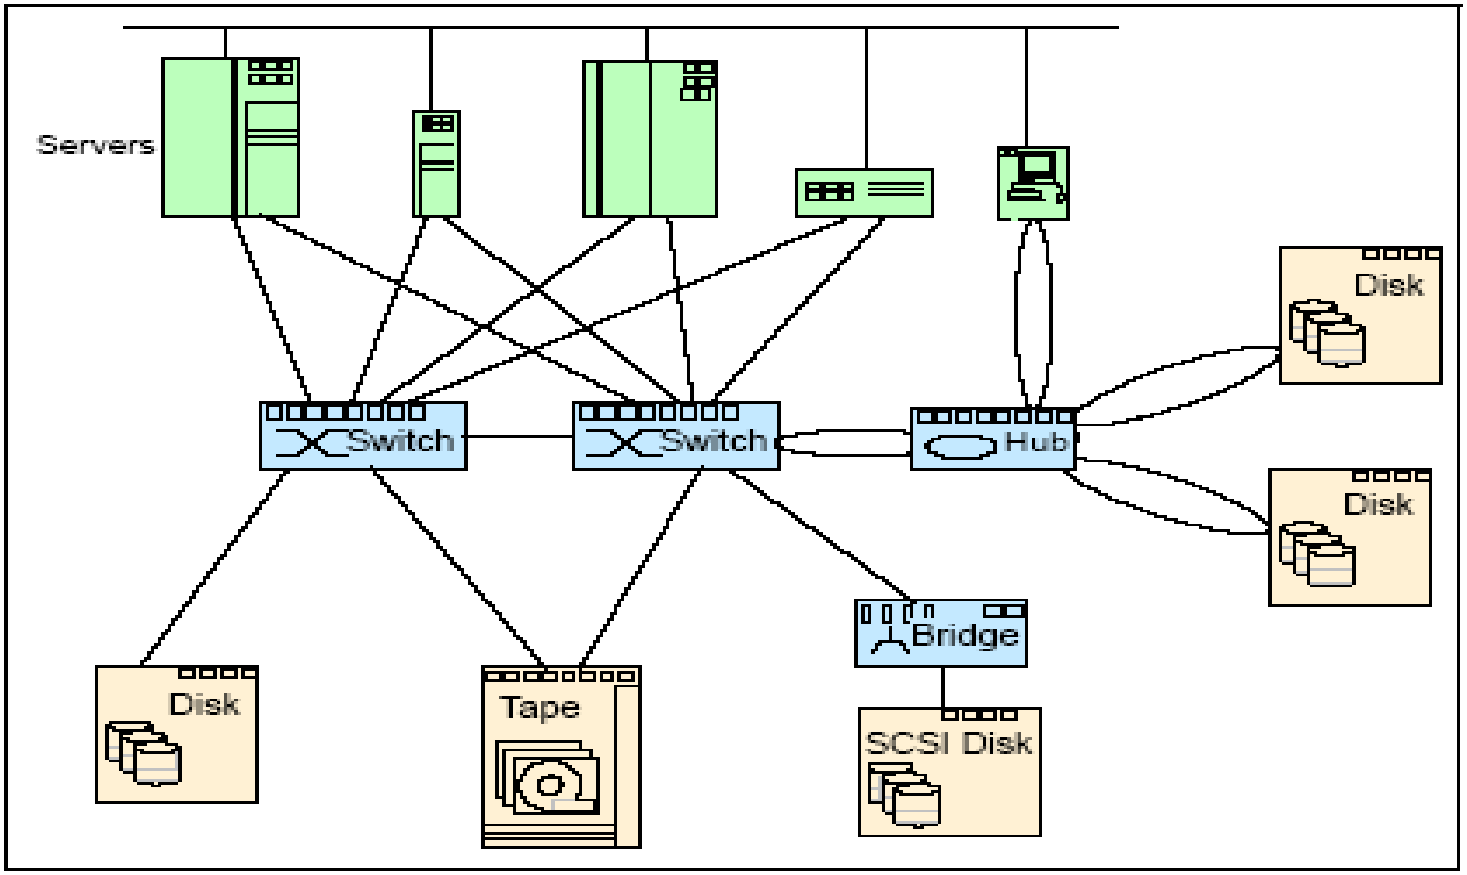

Figure 3-8 Sample switched fabric configuration

FC- Layered protocol model

- **Fibre Channel** is a **serial computer bus**
- intended for connecting
- **high-speed storage devices** to computers

- Fibre Channel, like many protocols, is a layered protocol. It consists of 5 layers, namely:
  - **FC0** The **physical layer**, which includes cables, fiber optics, etc.
  - **FC1** The **data link layer**, which implements the **8b/10b** encoding and decoding of signals.
  - **FC2** The **network layer**, defined by the FC-PH standard, consists of the core of FC (**framing and flow control**).
  - **FC3** A **thin layer** that implements **auxilary functions** that span across multiple ports on a Fibre Channel device.
  - **FC4 Application layers**, or **upper-layer protocol encapsulation**. This layer is responsible for encapsulation of various upper layers over FC.

FC- Layered protocol model

- **Layers**
- Fibre Channel (FC) is broken up into a series of five layers. The concept of layers, starting with the ISO/OSI seven-layer model, allows the development of one layer to remain independent of the adjacent layers.

- Although, FC contains five layers, those layers follow the general principles stated in the ISO/OSI model.

- The 5 layers are divided into 2:
  - Physical and signaling layer
  - Upper layer

FC- Layered protocol model

- The **five layers** are illustrated in Figure 3-30 on page 102

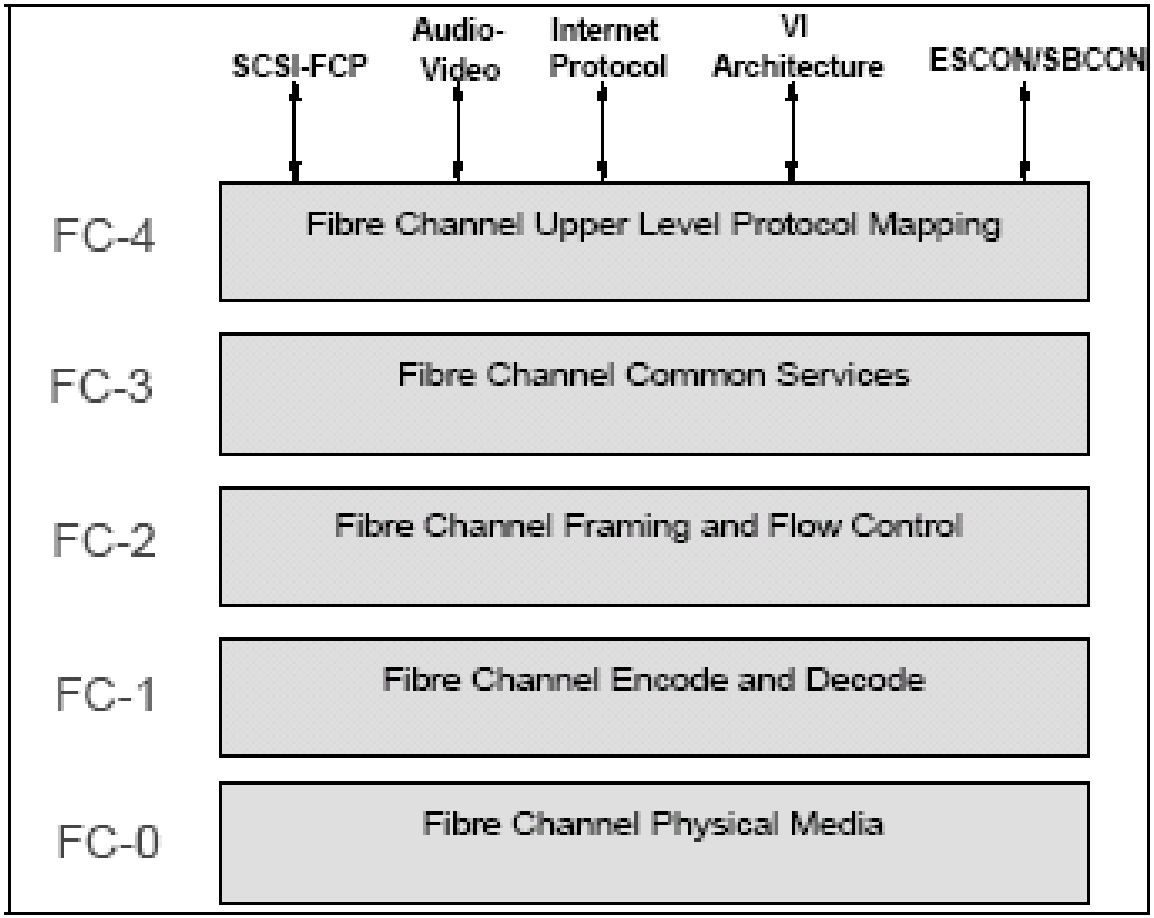

Figure 3-30 Fibre Channel layers

Physical and signaling layers

- The physical and signaling layers include the three lowest layers: FC-0, FC-1, and FC-2.
- **Physical interface and media: FC-0**
- The lowest layer, FC-0, defines the physical link in the system, including the
- cabling, connectors, and electrical parameters for the system at a wide range of data rates. This level is designed for maximum flexibility, and allows the use of a large number of technologies to match the needs of the configuration.
- A communication route between two nodes can be made up of links of different technologies. For example, in reaching its destination, a signal might start out on copper wire and become converted to single-mode fiber for longer distances.

- ***Laser safety***
- Fibre Channel often uses lasers to transmit data, and can, therefore, present an optical health hazard. The FC-0 layer defines an open fiber control (OFC) system, and acts as a safety interlock for point-to-point fiber connections that use semiconductor **laser diodes** as the optical source. If the fiber connection is broken, the ports send a series of pulses until the physical connection is re-established and the necessary handshake procedures are followed.

Transmission protocol: FC-1

- The second layer, FC-1,
- provides the methods
- for adaptive **8B/10B encoding**
- to bind the maximum length of the code,
- maintain DC-balance,
- and provide word alignment.

- This layer is used to integrate the data with the clock information required by serial transmission technologies.

Framing and signaling protocol: FC-2

- FC-2 specifies a data transport mechanism that is independent of upper layer protocols. FC-2 is self-configuring and **supports: point-to-point**, **Arbitrated Loop**, and **switched environments**.

- FC-2, which is the third layer of the FC-PH, provides the **transport methods** to determine:
  - Topologies based on the presence or absence of a fabric
  - Communication models
  - **Classes of service** provided by the fabric and the nodes
  - General fabric model
  - Sequence and exchange identifiers
  - Segmentation and reassembly
- Data is transmitted in 4-byte ordered sets containing data and control characters.
- Ordered sets provide the availability to obtain bit and word synchronization, which also establishes word boundary alignment.

- Together, FC-0, FC-1, and FC-2 form the Fibre Channel physical and signaling interface (FC-PH).

FC-3

- **Upper layers**
- The Upper layer includes two layers: FC-3 and FC-4.
- **Common services: FC-3**
- FC-3 defines functions that **span multiple ports** on a single-node or fabric.

- Functions that are currently supported include:

- **1. A *Hunt Group*** is a set of associated N_Ports attached to a single node.
- This set is assigned an alias identifier that allows **any frames** containing
- the alias to be routed to **any available** N_Port within the set. This
- decreases latency in waiting for an N_Port to become available.

- 2. **Striping**
- – *Striping* is used to multiply bandwidth, using multiple **N_Ports in parallel** to
- **transmit** a **single information unit across multiple links**.

- 3. **Multicast**
- – ***Multicast*** delivers a **single transmission** to **multiple destination ports**.
- Thisincludes the ability to broadcast to all nodes or a subset of nodes.

FC-4

- ***Upper layer protocol mapping (ULP): FC-4***
- The **highest layer, FC-4**, provides the **application-specific protocols**.
- Fibre Channel is equally adept at transporting both network and channel information and allows both protocol types to be concurrently transported over the **same physical interface**.

- Through mapping rules, a specific FC-4 describes how ULP processes of the same FC-4 type interoperate.
- A channel example is **Fibre Channel Protocol (FCP).** This is used to transfer SCSI data over Fibre Channel.
- A **networking example** is sending IP (Internet Protocol) packets between nodes.
- **FICON** is another ULP in use today for **mainframe systems**. FICON is a contraction of *Fibre Connection* and refers to running ESCON traffic over Fibre Channel.

FC-SATA-SAS layers

Overview

- SAS borrows its lower layers from SATA
- SAS borrows its upper layers from Fibre Channel and FCP (the SCSI over Fibre Channel mapping)

FC layer	Standard(s)	Equivalent SAS layer(s)	Standards
ULP	FCP (and various SCSI)	SCSI application	SAS & others
FC-4 Mapping	FCP	Transport	SAS
FC-3	FC-FS	Port layer (roughly)	SAS
FC-2 Protocol	FC-FS	Link and some transport	SAS
FC-1 Code	FC-FS	Phy	SAS
FC-0 Physical	FC-PI	Physical	SAS

- SAS also supports ATA and management upper layers
- FC also supports VI, TCP, and others

Phy layers

Physical/phy layers

FC	SAS
Media: Copper and fiber Up to 100 km with optical	Media: Copper only 1 m internal cable or 18" backplane trace 10 m external cable (4 wide)
Speed: 1.0625, 2.125, 4.25, and 10.51875 or 12.75 Gbps	Speed: 1.50, 3.00 Gbps
Both use 8b10b coding Both use dwords (4 bytes) as base quantities Both are big-endian (SATA is little-endian)	

Topologies

Topologies (part 1)

FC	SAS
Simple or complex fabrics: •Fabric Shortest Path First algorithm •Cross-connected switches •2^24 = 16 M devices in a fabric •Dual fabrics for no-single-points-of-failure	Simple fabrics: •Tree of expanders •No loops allowed (complicates routing) •Initiators perform discovery and program routing tables •Expander sets have maximum of 128 phys •128 * 128 = 16 K maximum addresses in a SAS fabric •Dual fabrics for no-single-points-of-failure
FC supports arbitrated loops •Often used for disk drive connections •127 devices in a loop •Dual loops for no-single-points-of-failure	No loops
Both are point-to-point serial interfaces, not multidrop busses	

Topologies

Addressing (part 1)

- Each SAS port has a SAS address
  - 64-bit NAA=5 format as defined by SPC-3 (and FC-FS)
  - Set by the factory
  - Used for addressing
- Each FC node has a node name and each FC port has a port name
  - Same NAA format as SAS
    - additional NAA types are allowed (e.g. NAA=3)
  - Set by the factory
  - Exchanged during logins
  - Not used for addressing (FC 24-bit fabric address is assigned)

WWN (64-bits)

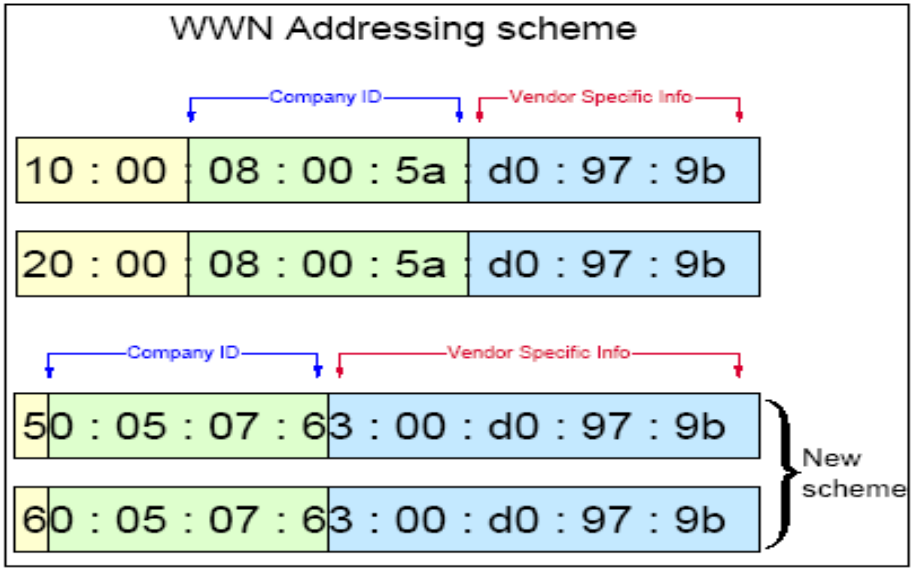

Figure 3-15 World Wide Name addressing scheme

The complete list of vendor identifiers as maintained by the IEEE is available at:
http://standards.ieee.org/regauth/oui/oui.txt

Table 3-1 lists a few of these vendor identifiers.

Table 3-1 WWN company identifiers

WWN (hex)	Company
00-50-76	IBM Corporation
00-60-69	Brocade Communications
08-00-88	McDATA Corporation
00-60-DF	CNT Technologies Corporation

WWPN: 24-bit port address

- Each port in the switched fabric has its own unique 24-bit address. The relationship between this 24-bit address and the 64-bit address associated with World Wide Names is explained in this section.
- The 24-bit address scheme removes manual administration of addresses by allowing the topology itself to assign addresses. This is not like WWN addressing, in which the addresses are assigned to the manufacturers by the IEEE standards committee, and are built in to the device at time of manufacture, similar to naming a child at birth. If the topology itself assigns the 24-bit addresses, then somebody has to be responsible for the addressing scheme from WWN addressing to port addressing.
- In the switched fabric environment, the switch itself is responsible for assigning and maintaining the port addresses. When the device with its WWN logs into the switch on a specific port, the switch will assign the port address to that port and the switch will also maintain the correlation between the port address and the WWN address of the device on that port. This function of the switch is implemented by using a Name Server.
- The **Name Server** is a component of the fabric operating system, **which runs inside the switch**. It is essentially a database of objects in which fabric-attached devices registers their values.
- Dynamic addressing also removes the potential element of human error in address maintenance, and provides more flexibility in additions, moves, and changes in the SAN.

WWPN: 24-bit port address

- A 24-bit port address consists of 3 parts:
- Domain (bits from 23 to 16)
- Area (bits from 15 to 08)
- Port or Arbitrated Loop physical address: AL_PA (bits from 07 to 00)

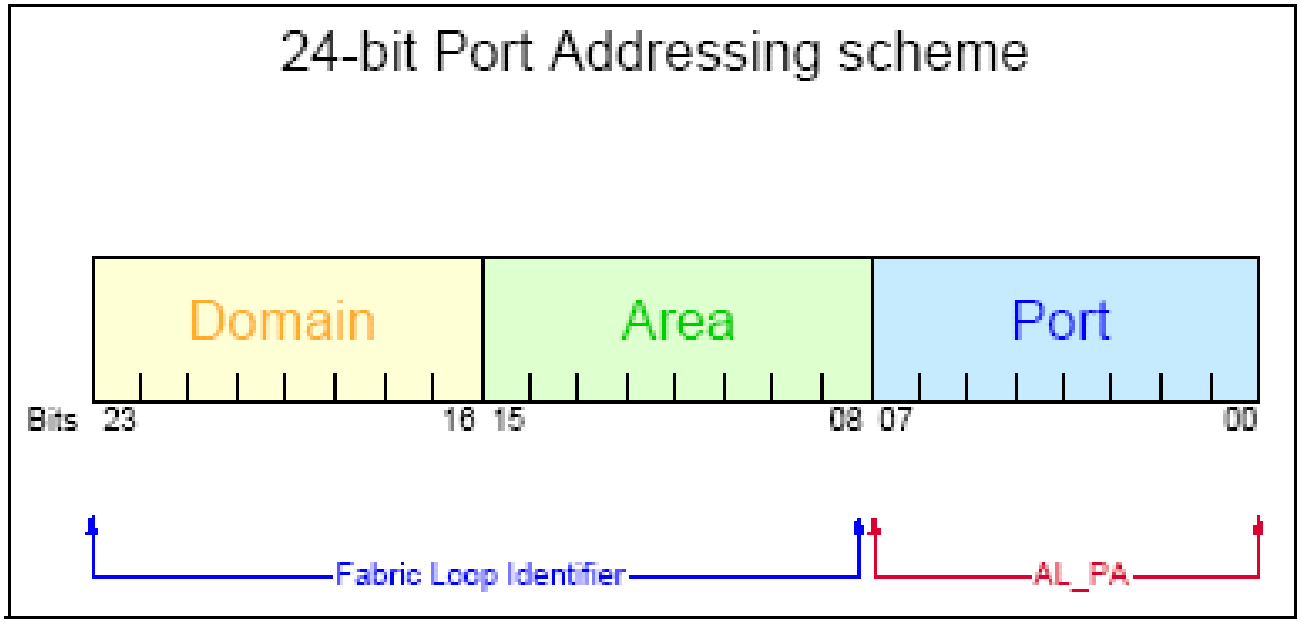

Figure 3-18 Fabric port address

WWN:WWPN

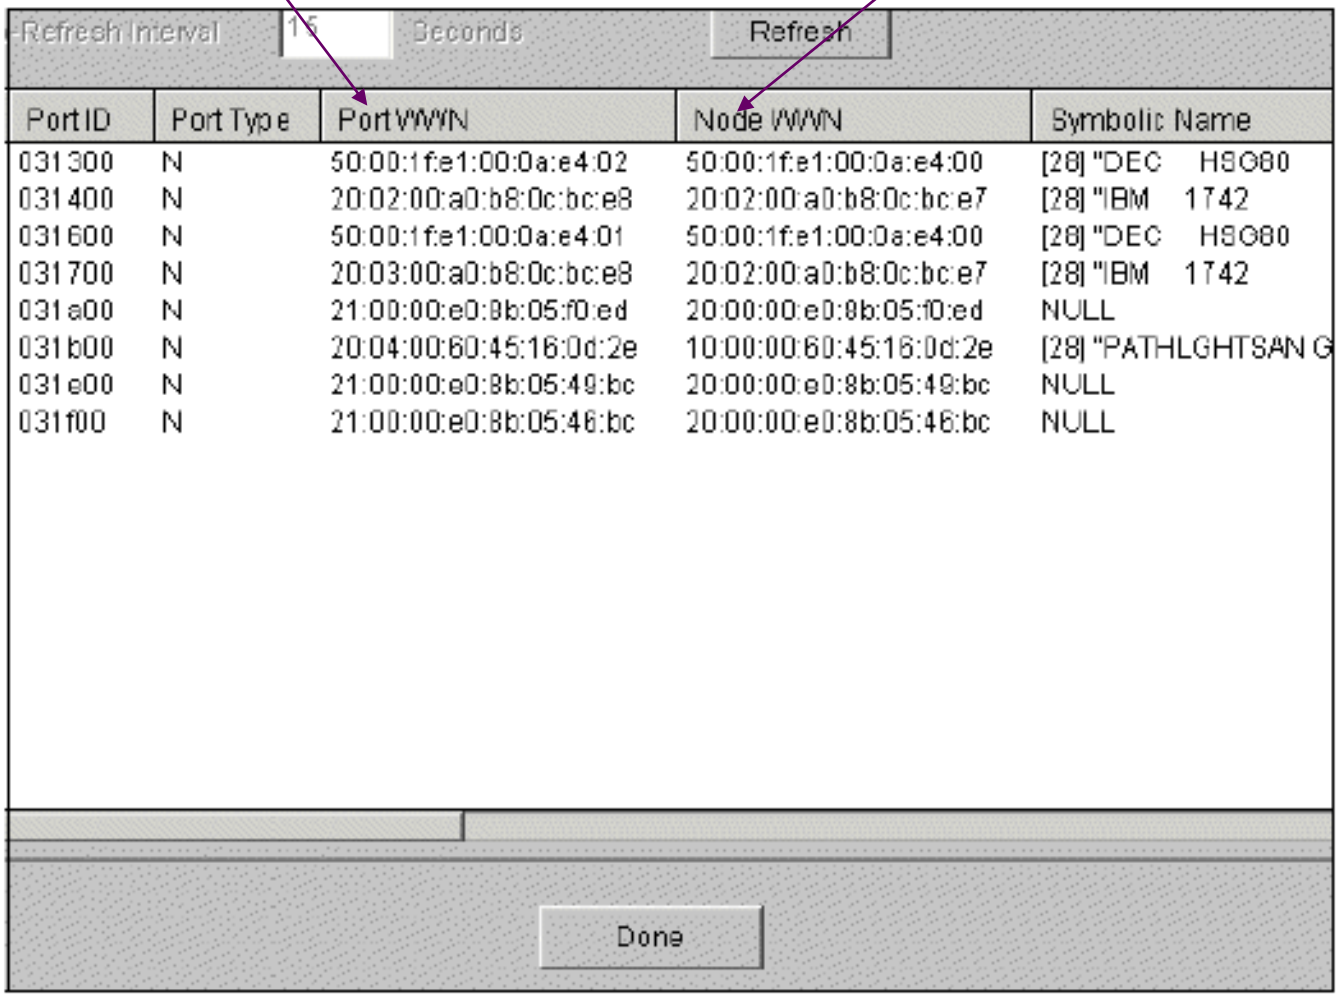

Figure 3-17 WWN and WWPN entries in a name server table

The switch must make a correlation between the WWN and the 24-bit port address.

Addressing part-2

Addressing (part 2)

- FC uses a port address for frame routing
  - 24-bit Port address assigned during initialization
  - 8-bit Arbitrated Loop Physical Address (AL_PA) is lower eight bits of a 24-bit address; used for loop arbitration
- FC has many well-known addresses for fabric services
  - Broadcast, fabric login, fabric controller, directory server, time server, management server, quality of service facilitator, alias server, security key distribution server, clock synchronization server, multicast server
    - Services can be provided by one or more switches or by management appliances
  - In SAS, expanders are addressable like any other device
    - No special addresses

Frames

Frames

- Both use a 24 byte frame header
- SAS has no optional frame headers
  - FC defines several optional headers
    - expiration/security, network, association, and device
- SAS maximum data size 1 KB
  - FC maximum 2 KB data
- SAS uses primitives for link control (e.g. ACK primitive)
  - FC uses frames for link control (e.g. ACK_0 frame)

FC login

Login/logout

- SAS has no logins
- FC has three layers of logins
  1. Fabric login (FLOGI): announce presence to the fabric
  2. Port login/logout (PLOGI/LOGO): establish a session with another port
  3. Process login/logout (PRLI/PRLO): establish a session from an upper level process with another upper lever process
     - Third-party process logout (TPRLO) also available

Class of service

Classes of service

- SAS only has one class of service, similar to Fibre Channel class 1

FC class of service	Description
1	Virtual circuit, full bandwidth (how a telephone network seems to operate)
2	Frame-routed, confirmed delivery (ACKs), out-of-order frames possible
3	Frame-routed, unconfirmed delivery, out-of-order and lost frames possible (though they cause problems in practice; most switches provide in-order delivery to be safe). Almost all FC implementations use class 3.
4	Virtual circuit, fractional bandwidth
6	Multicast

iSCSI

- Modern low cost SAN
- Low costs through the re-use of **existing LAN's** and LAN skills.
- Based on Gb Ethernet
- iSCSI= IP based SAN=SCSI over IP
- **Block** based protocol **like FC**
- iSCSI combine:
  - SCSI protocol
  - IP transport protocol
  - Ethernet Infrastructure

Classical SCSI v iSCSI

- Local machine

- SCSI initiator (controller)
- SCSI targets (disks: JBOD, RAID)

- Phases
  - **Command** I-T over local SCSI bus
  - **Data** I-T or T-I over local SCSI
  - **Status** T-I over local SCSI bus

Classical SCSI v iSCSI

- iSCSI
- SCSI Targets and Initiators are everywhere on the network
- SCSI initiator (host with iSCSI support)
- SCSI targets (host with storage, iSCSI suport)

- Phases
  - Command I-T — SCSI command into IP frame, over IP network
  - Data I-T or T-I — data into IP frames, over IP network
  - Status T-I — SCSI status into IP frame, over IP network

iSCSI

- iSCSI architecture involves a host configured as an
  - **iSCSI target.**
  - **iSCSI initator**

- The **iSCSI target** can be
  - a **server with locally connected storage** or
  - a storage device that natively supports iSCSI.

- **Clients** that access the storage over the network using the iSCSI protocol are known as **initiators**.

- **Initiators** need to have iSCSI client software
- installed in order
- to access the iSCSI target.

iSCSI Example

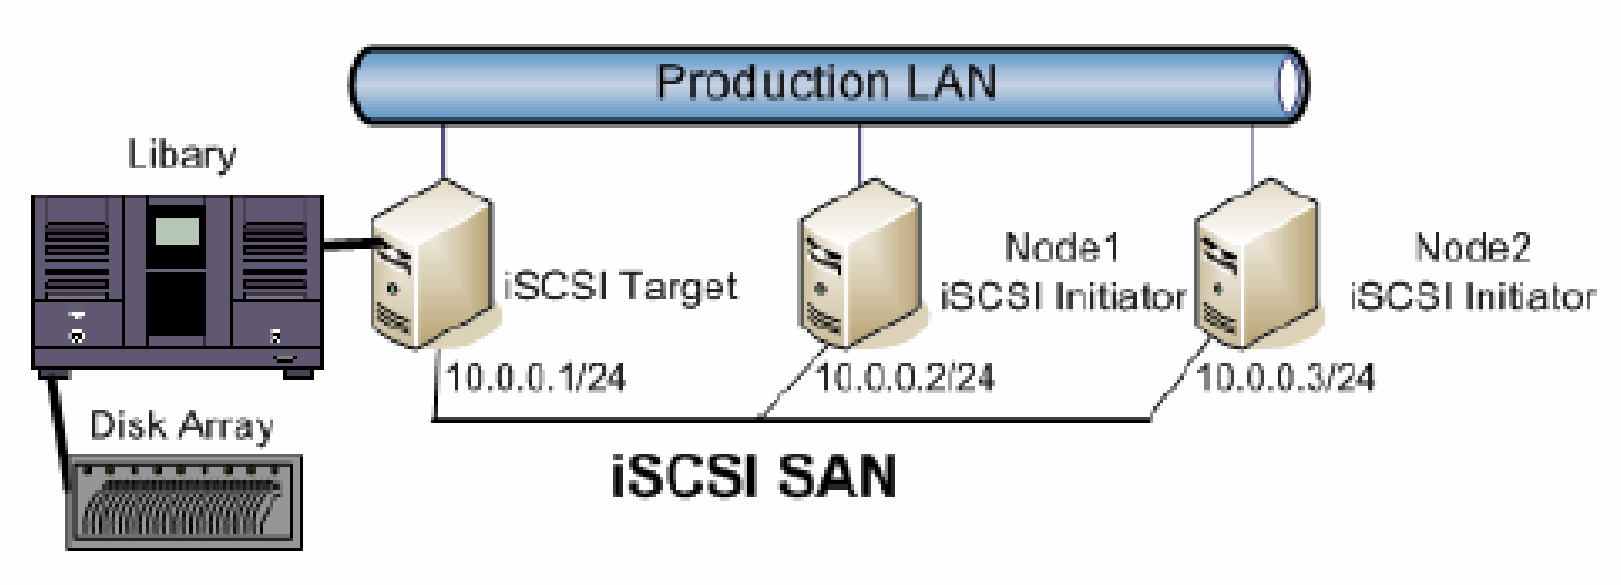

Figure 1.8: A small iSCSI SAN.

iSCSI

- In iSCSI networks, SCSI data blocks are transported within IP packets over Ethernet or other WAN networks.
- **iSCSI protocol** is managed either:
  - in memory resident iSCSI drivers,
  - in dedicated **iSCSI HBA's**
  - NIC (Network Interface Cards).
- The major benefits of iSCSI are the very low storage networks are likely to set-up and operational costs through the **re-use** of existing LAN's and LAN skills.
- iSCSI looks set to become established as a strong player in particular within **Windows workgroup environments**.
- **Larger Unix environments** and large enterprises are likely to continue to be **predominantly FC SAN-based**.

iSCSI Implementation

- Finally, how does my **server connect**, there are **3 approaches:**
- 1. A standard NIC with an iSCSI driver
- 2. A TOE (TCP Offload Engine) NIC with an iSCSI driver
- 3. An HBA (Host Bus Adapter) designed for iSCSI from the traditional Fibre Channel HBA vendors

TOE

- **TCP offload engines** (TOEs).
- This breakthrough moves
- much of the computational overhead
- of the TCP/IP software stack into hardware.

- Thus it largely eliminates the heavy burden the application CPU normally has when TCP traffic is high.

- Although intelligent TOE host bus adapters/network interface cards (HBA/NICs) will add to cost,
- IDC believes they will be needed on only the high-end portion of the market.

Main Benefits

- Cost
- Distance is not a limiting factor for the iSCSI SAN.
- Security is built into iSCSI and uses standard TCP and GigE technologies such as IPSEC
- Backward 10/100eclinet + 1/10 GbitE eClients

Emerging technologies

- Many hot technologies emerge. Some stay around and some go the way of the dinosaurs. Here is what is hot at the time of this writing.
- **iSCSI**
- **iFCP**
- **FCIP**

iSCSI

- ***Internet SCSI* (iSCSI)** is a transport protocol that carries SCSI commands from an initiator to a target.
- It is a data storage networking protocol
- that transports
- standard Small Computer System Interface (SCSI) requests
- **over**
- the standard Transmission Control Protocol/Internet Protocol (TCP/IP) networking technology.
- iSCSI enables the implementation of IP-based storage area networks (SANs), enabling customers to use the same networking technologies, from the box level to the Internet, for both storage and data networks.
- As it uses TCP/IP, iSCSI is also well-suited to run over almost **any physical network**. By eliminating the need for a second network technology just for storage, **iSCSI will lower the costs** of deploying networked storage and increase its potential market.

iSCSI

- One of the major advantages is that as iSCSI carries SCSI commands over existing IP networks, it has an innate and important ability to facilitate the transfer of data over both inter- and intra-nets, and to manage storage over long distances.

- iSCSI, in simple terms, works in this way: when an end user or application sends a request, the operating system generates the appropriate SCSI commands and data request. These then go through encapsulation procedures. A packet header is added before the resulting IP packets are transmitted over a TCP/IP connection. At the receiving end when the packets are received they are unravelled and the SCSI commands are separated from the data request. The SCSI commands are sent on to the target storage controller and, ultimately, the SCSI storage device.
- iSCSI is a bidirectional protocol which means it can also be used to return data in response if required.

FCIP

- *Fibre Channel over IP* (FCIP) is also known as ***Fibre Channel tunneling*** or ***storage tunneling***.

- It is a method for allowing the transmission of Fibre Channel information to be tunnelled through the IP network. Because most organizations already have an existing IP infrastructure, the attraction of being able to link geographically dispersed SANs, at relatively low cost, is enormous.

- **FCIP** encapsulates **Fibre Channel block data** and subsequently transports it over a TCP socket. TCP/IP services are utilized to **establish connectivity between remote SANs**. Any congestion control and management, as well as data error and data loss recovery, is handled by TCP/IP services, and does not affect FC fabric services.

FCIP

- The major point with FCIP is that is does not replace FC with IP. It simply allows deployments of **FC fabrics using IP tunnelling**. The assumption that this might lead to is that the industry has decided that FC-based SANs are more than appropriate. The only need for the IP connection is to facilitate any distance requirement that is beyond the current scope of an FCP SAN.

- As a starting point for an IP storage discussion, the Storage Networking Industry Association's IP Storage Forum (IPS Forum) is a vendor-neutral environment for end users to become informed on the current and future directions of IP-based storage technology:
- http://www.snia.org/ipstorage/home/

iFCP

- Internet FC Protocol (iFCP) is **much more robust than FCIP**.
- **Like FCIP**, iFCP can also be used to bridge FC switches over an IP network.
- However, this protocol also provides the ability to **network native IP storage devices** and **FC devices together** on the **same IP-based storage network**. With the rise of gigabit Ethernet networks, consider iFCP as a way to provide full integration between an FC and IP network. Another rising protocol that provides the same level of hardware integration over gigabit Ethernet is iSCSI.
- Internet Fibre Channel Protocol (iFCP) **is** a **mechanism for transmitting data** to and from **Fibre Channel storage devices in a SAN**, or on the **Internet using TCP/IP.**
- FCP gives the ability to incorporate already **existing SCSI** and **Fibre Channel networks into the Internet**.
- iFCP is able to be used in tandem with existing Fibre Channel protocols, such as FCIP, or it can replace them.
- Whereas FCIP is a **tunneled solution**,
- **iFCP** is an **FCP-routed solution**.

iFCP = iSCSI + FCIP

- The appeal of iFCP is that for customers
- that have a wide range of FC devices,
- and who want to be able to connect these with the IP network, iFCP gives the ability to permit this.

- iFCP can interconnect FC SANs with IP networks, and also allows customers to use the TCP/IP network in place of the SAN.

- iFCP is a gateway-to-gateway protocol, and does not simply encapsulate FC block data.

- Gateway devices are used as the medium **between the FC initiators and targets.** As these gateways can either replace or be used in tandem with existing FC fabrics, iFCP could be used to help migration from a Fibre Channel SAN to an IP SAN, or allow a combination of both.