IT8001
Information Storage and Management

Professional Elective

Topic

Storage Area Networks

Varun C.M.

Assistant Professor, Department of Information Technology,
St. Xavier's Catholic College of Engineering

Expected Outcome

CO2

Evaluate storage architectures, including storage subsystems, DAS, SAN, NAS, and CAS.

Fiber Channel

Fiber Channel is a high-speed network technology that runs on high-speed optical fiber cables and serial copper cables.
The FC technology was developed to meet the demand for increased speeds of data transfer between servers and mass storage systems.

SAN Evolution

A SAN carries data between servers and storage devices through Fiber Channel network.
A SAN enables storage consolidation and enables storage to be shared across multiple servers. This improves the utilization of storage resources compared to direct-attached storage architecture and reduces the total amount of storage an organization needs to purchase and manage.

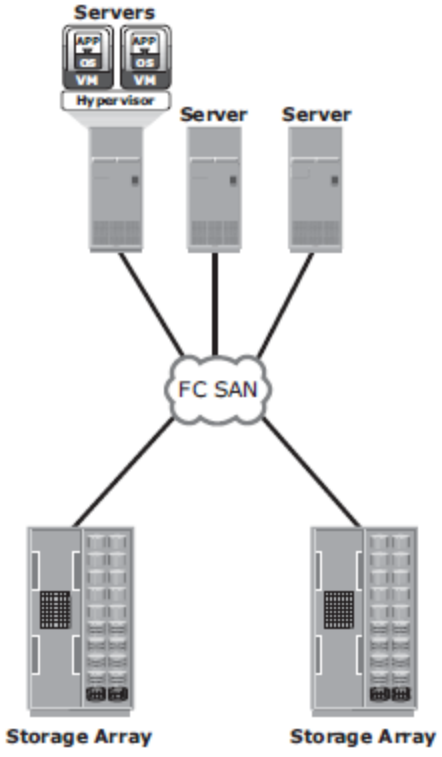

Image courtesy EMC Corporation, "Information Storage and Management"

SAN Evolution contd…

In its earliest implementation, the FC SAN was a simple grouping of hosts and storage devices connected to a network using an FC hub as a connectivity device.

This configuration of an FC SAN is known as a Fibre Channel Arbitrated Loop (FC-AL).

The inherent limitations associated with hubs gave way to high-performance FC switches.

Use of switches in SAN improved connectivity and performance and enabled FC SANs to be highly scalable.

This enhanced data accessibility to applications across the enterprise.

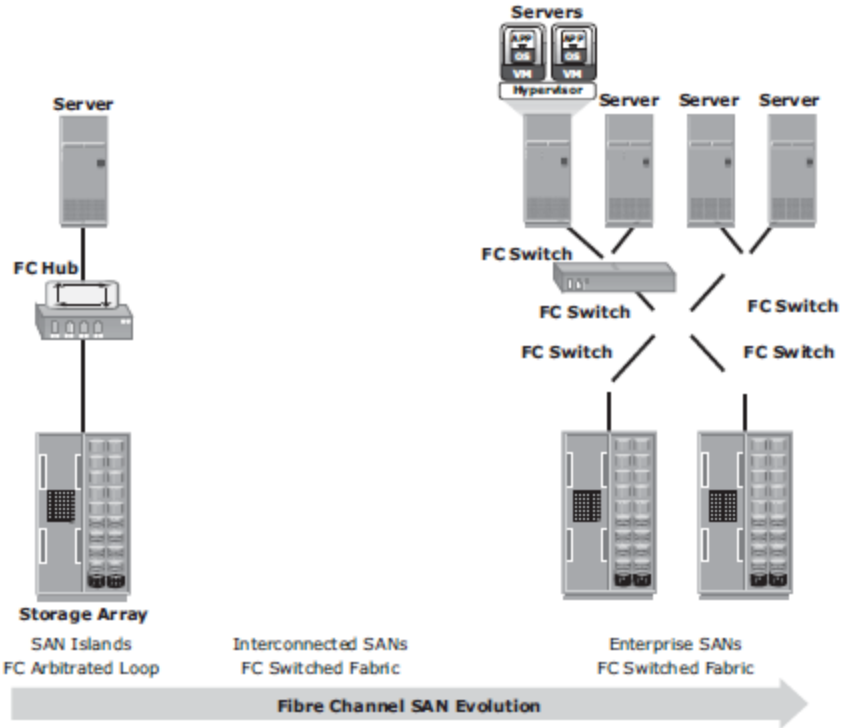

Image courtesy EMC Corporation, "Information Storage and Management"

SAN Components

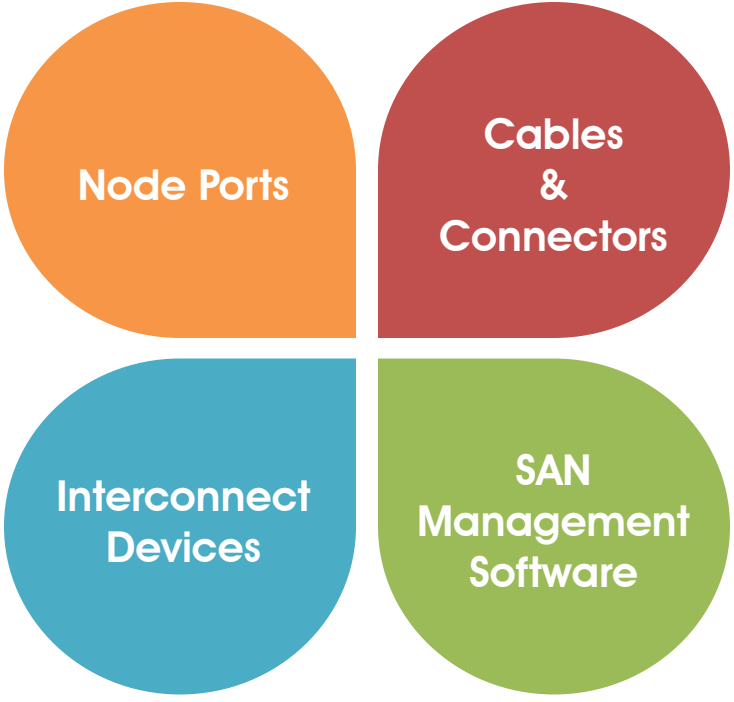

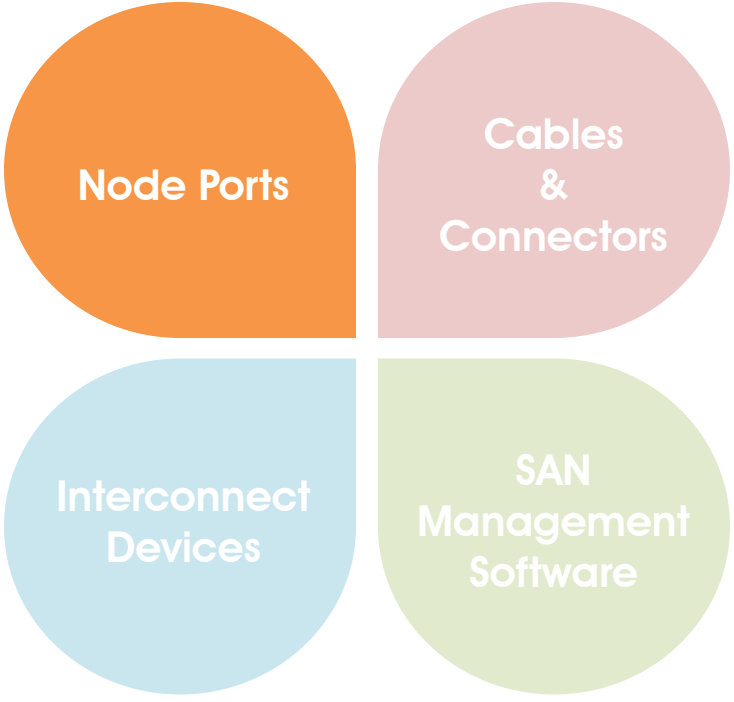

The end devices, such as hosts, storage arrays, and tape libraries, are all referred to as nodes.

Each node requires one or more ports to provide a physical interface for communicating with other nodes.

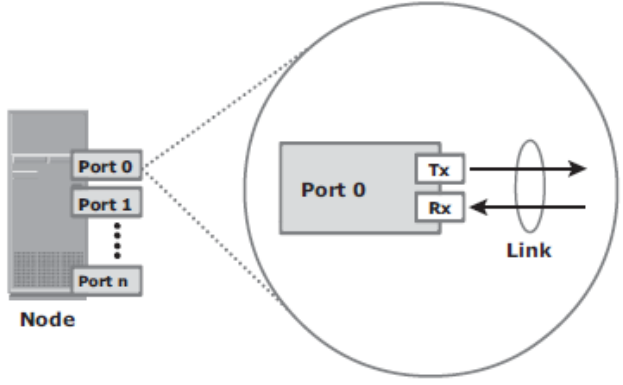

Image courtesy EMC Corporation, "Information Storage and Management"

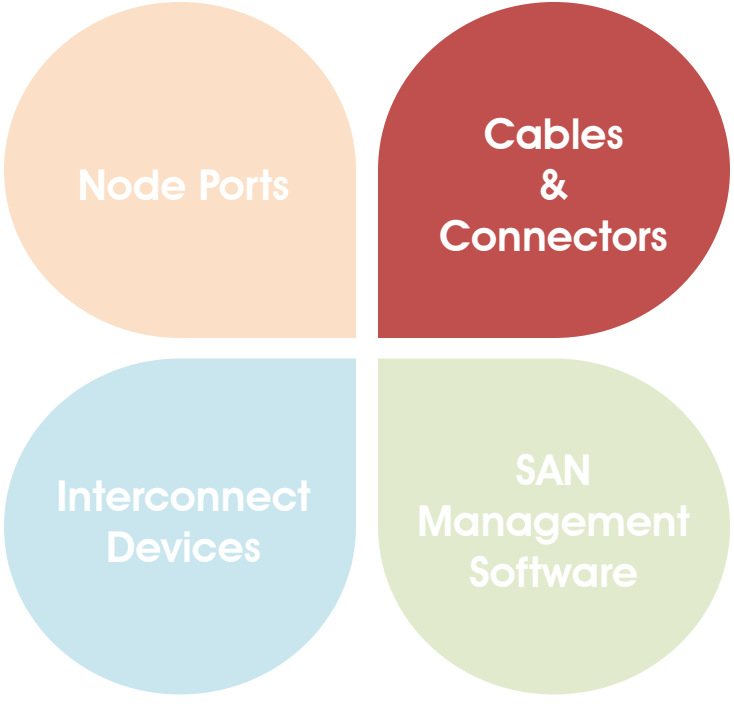

There are two types of optical cables: multimode and single-mode.

Multimode fiber (MMF) cable carries multiple beams of light projected at different angles simultaneously onto the core of the cable

Single-mode fiber (SMF) carries a single ray of light projected at the center of the core

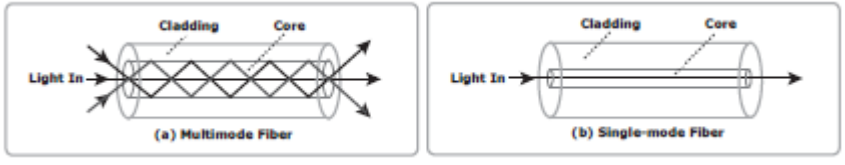

Image courtesy EMC Corporation, "Information Storage and Management"

SAN Components contd...

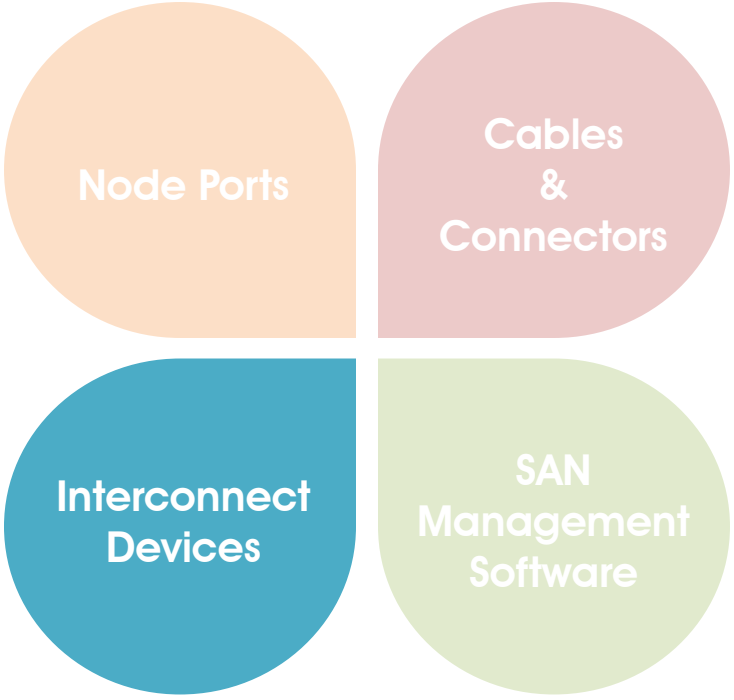

Hubs physically connect nodes in a logical loop or a physical star topology.
Switches are more intelligent than hubs and directly route data from one physical port to another.
Directors are high-end switches with a higher port count and better fault tolerance capabilities.
In a modular switch, the port count is increased by installing additional port cards to open slots.
A port card or blade has multiple ports for connecting nodes and other FC switches.

SAN Components contd...

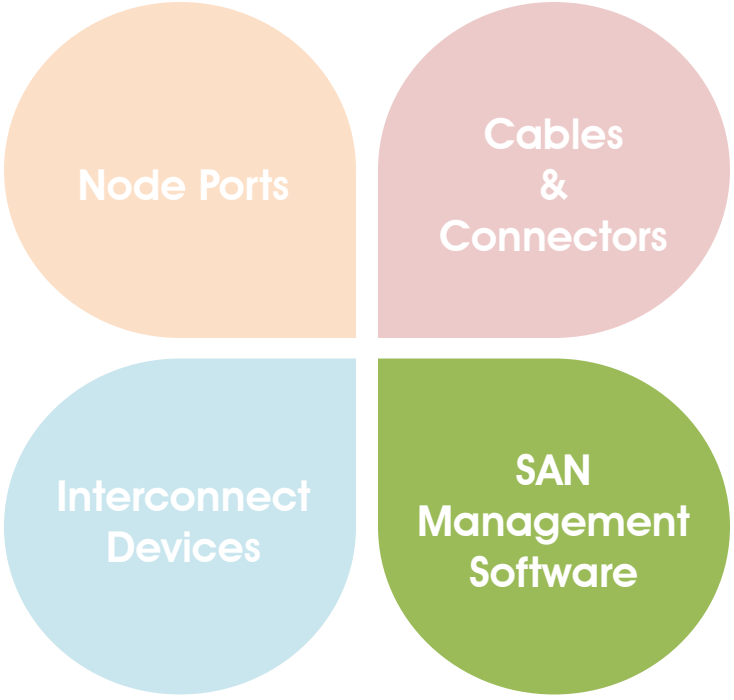

SAN management software manages the interfaces between hosts, interconnect devices, and storage arrays.

The software provides a view of the SAN environment and enables management of various resources from one central console.

It provides key management functions, including mapping of storage devices, switches, and servers, monitoring and generating alerts for discovered devices, and zoning

Fiber Channel Connectivity

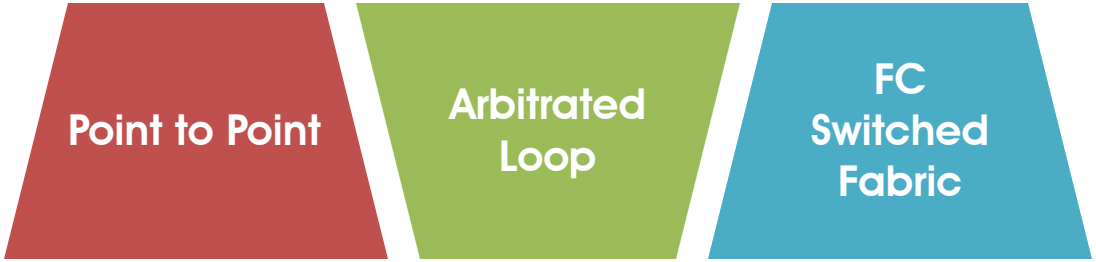

Fiber Channel Connectivity contd...

Point-to-point is the simplest FC configuration - two devices are connected directly to each other. This configuration provides a dedicated connection for data transmission between nodes.
However, the point-to-point configuration offers limited connectivity, because only two devices can communicate with each other at a given time.
Moreover, it cannot be scaled to accommodate a large number of nodes.

Fiber Channel Connectivity contd...

In the FC-AL configuration, devices are attached to a shared loop.
FC-AL has the characteristics of a token ring topology and a physical star topology.
In FC-AL, each device contends with other devices to perform I/O operations. Devices on the loop must "arbitrate" to gain control of the loop.
At any given time, only one device can perform I/O operations on the loop.

Fiber Channel Connectivity contd...

FC-SW is also referred to as fabric connect. A fabric is a logical space in which all nodes communicate with one another in a network. This virtual space can be created with a switch or a network of switches.

Each switch in a fabric contains a unique domain identifier, which is part of the fabric's addressing scheme.

In FC-SW, nodes do not share a loop; instead, data is transferred through a dedicated path between the nodes.

In a switched fabric, the link between any two switches is called an Inter switch link (ISL). ISLs enable switches to be connected together to form a single, larger fabric.

Fiber Channel Ports

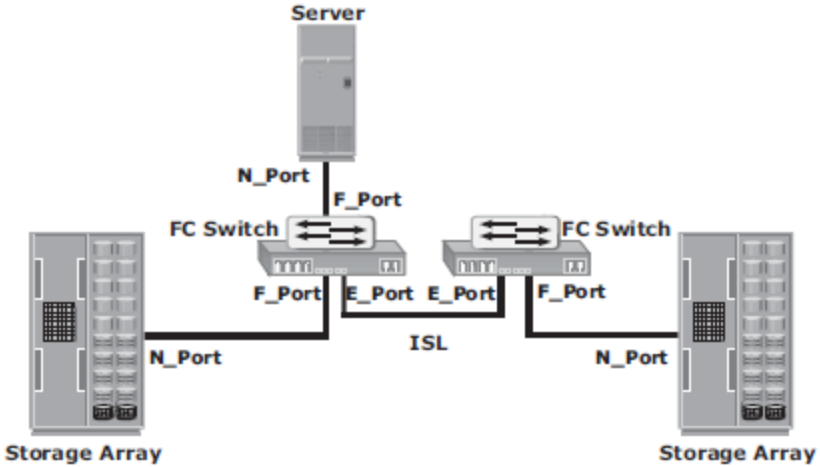

Image courtesy EMC Corporation, "Information Storage and Management"

N_Port

An end point in the fabric. This port is also known as the *node port*. Typically, it is a host port or a storage array port connected to a switch in a switched fabric.

Fiber Channel Ports contd...

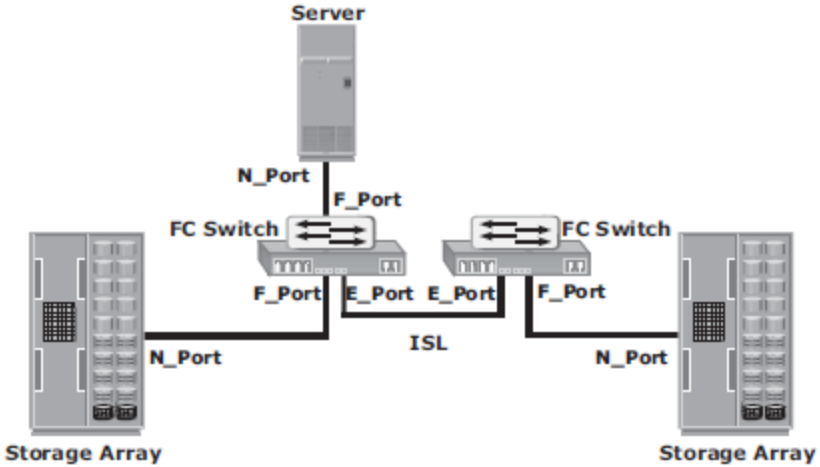

Image courtesy EMC Corporation, "Information Storage and Management"

N_Port

An end point in the fabric. This port is also known as the *node port*. Typically, it is a host port or a storage array port connected to a switch in a switched fabric.

Fiber Channel Ports contd...

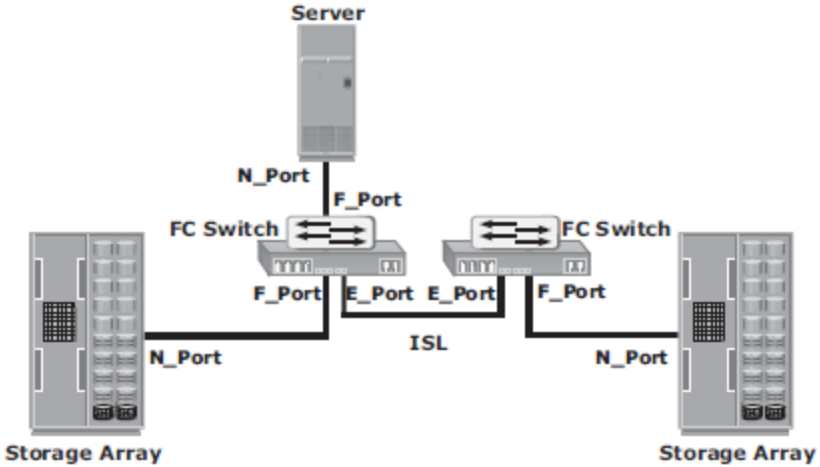

Image courtesy EMC Corporation, "Information Storage and Management"

E_Port

A port that forms the connection between two FC switches.
This port is also known as the expansion port.
The E_Port on an FC switch connects to the E_Port of another FC switch in the fabric through ISLs.

Fiber Channel Ports contd...

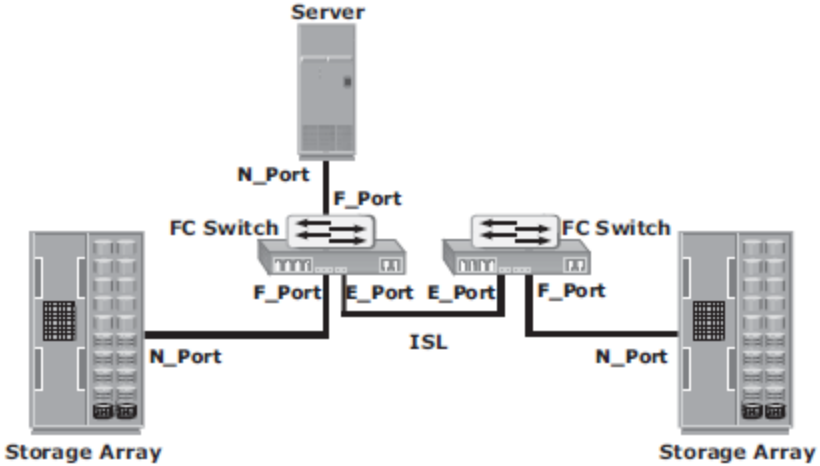

Image courtesy EMC Corporation, "Information Storage and Management"

E_Port

A port that forms the connection between two FC switches.
This port is also known as the expansion port.
The E_Port on an FC switch connects to the E_Port of another FC switch in the fabric through ISLs.

Fiber Channel Ports contd...

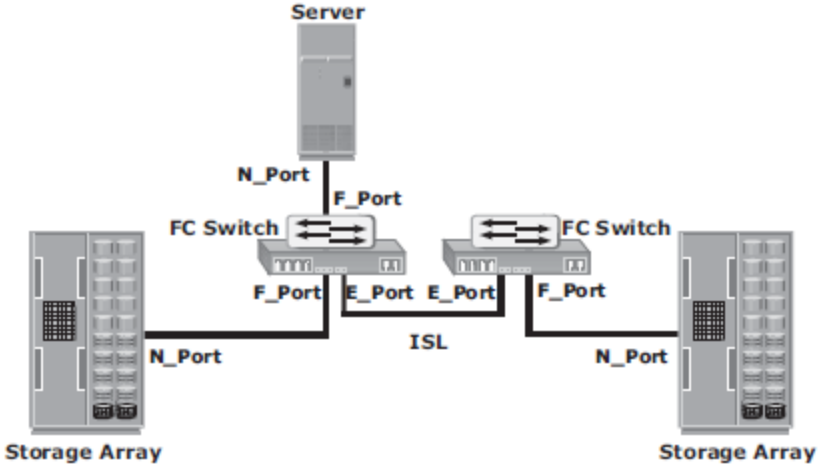

Image courtesy EMC Corporation, "Information Storage and Management"

F_Port

A port on a switch that connects an N_Port.
It is also known as a fabric port.

Fiber Channel Ports contd...

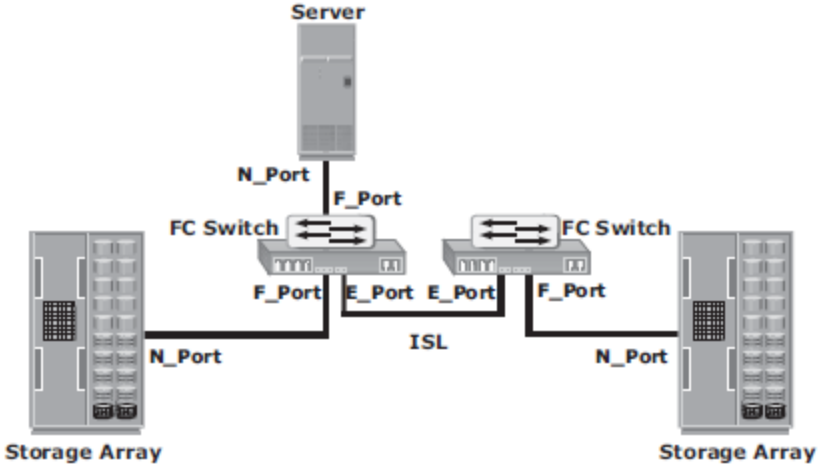

Image courtesy EMC Corporation, "Information Storage and Management"

F_Port

A port on a switch that connects an N_Port.
It is also known as a fabric port.

Fiber Channel Ports contd...

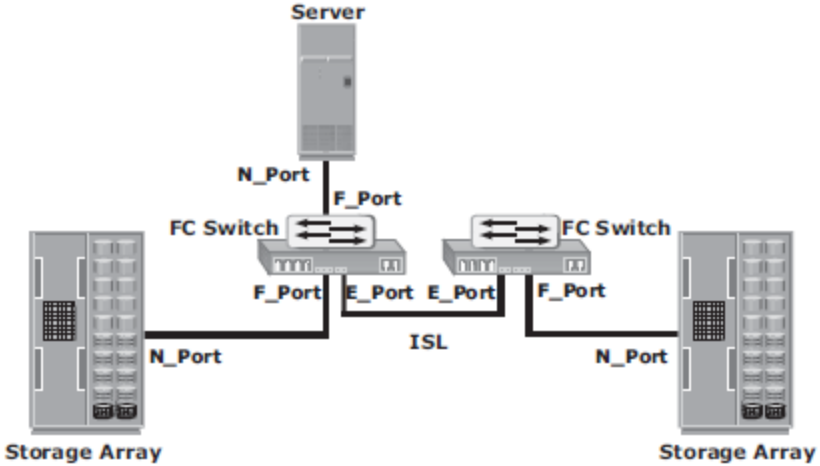

Image courtesy EMC Corporation, "Information Storage and Management"

G_Port

A generic port on a switch that can operate as an E_Port or an F_Port and determines its functionality automatically during initialization.

Fiber Channel Ports contd...

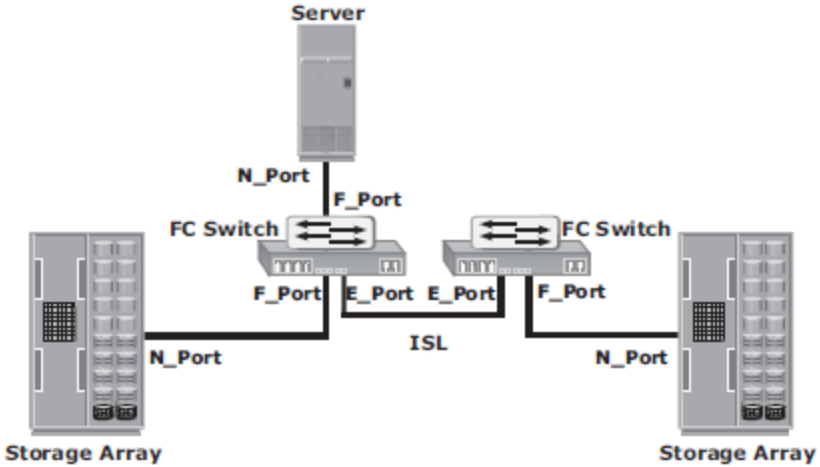

Image courtesy EMC Corporation, "Information Storage and Management"

G_Port

A generic port on a switch that can operate as an E_Port or an F_Port and determines its functionality automatically during initialization.

Fiber Channel Architecture

The FC architecture represents true channel/network integration and captures some of the benefits of both channel and network technology.
FC SAN uses the Fiber Channel Protocol (FCP) that provides both channel speed for data transfer with low protocol overhead and scalability of network technology.
FCP forms the fundamental construct of the FC SAN infrastructure.
Fiber Channel provides a serial data transfer interface that operates over copper wire and optical fiber.
The key advantages of FCP are,

- Sustained transmission bandwidth over long distances.
- Support for a larger number of addressable devices over a network. Theoretically, FC can support more than 15 million device addresses on a network.
- Support speeds up to 16 Gbps (16 GFC).

Fiber Channel Protocol Stack

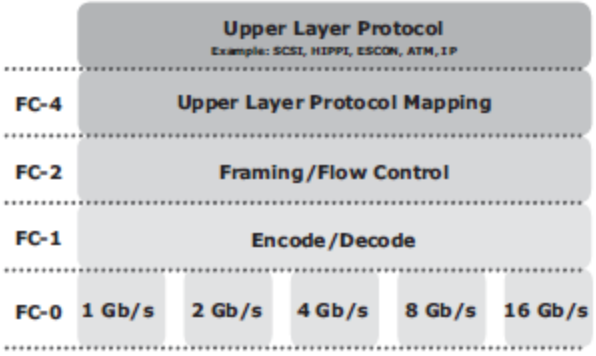

Image courtesy EMC Corporation, "Information Storage and Management"

FC-4 Layer

FC-4 is the uppermost layer in the FCP stack. This layer defines the application interfaces and the way *Upper Layer Protocols* (ULPs) are mapped to the lower FC layers.

The FC standard defines several protocols that can operate on the FC-4 layer.

Some of the protocols include SCSI, High Performance Parallel Interface (HIPPI) Framing Protocol, Enterprise Storage Connectivity (ESCON), Asynchronous Transfer Mode (ATM), and IP.

Fiber Channel Protocol Stack

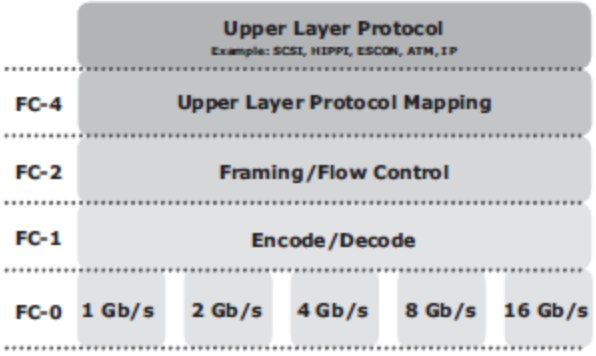

Image courtesy EMC Corporation, "Information Storage and Management"

FC-4 Layer

FC-4 is the uppermost layer in the FCP stack. This layer defines the application interfaces and the way *Upper Layer Protocols* (ULPs) are mapped to the lower FC layers.

The FC standard defines several protocols that can operate on the FC-4 layer.

Some of the protocols include SCSI, High Performance Parallel Interface (HIPPI) Framing Protocol, Enterprise Storage Connectivity (ESCON), Asynchronous Transfer Mode (ATM), and IP.

Fiber Channel Protocol Stack

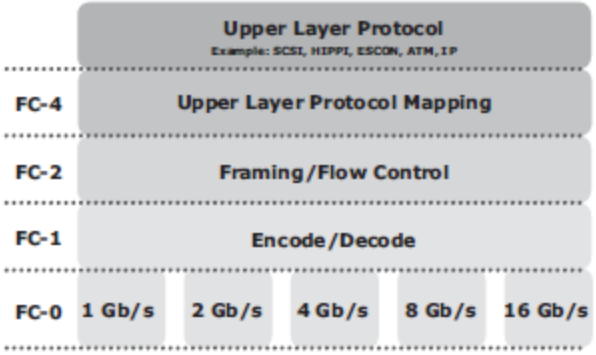

Image courtesy EMC Corporation, "Information Storage and Management"

FC-2 Layer

The FC-2 layer provides Fiber Channel addressing, structure, and organization of data (frames, sequences, and exchanges).

It also defines fabric services, classes of service, flow control, and routing.

Fiber Channel Protocol Stack

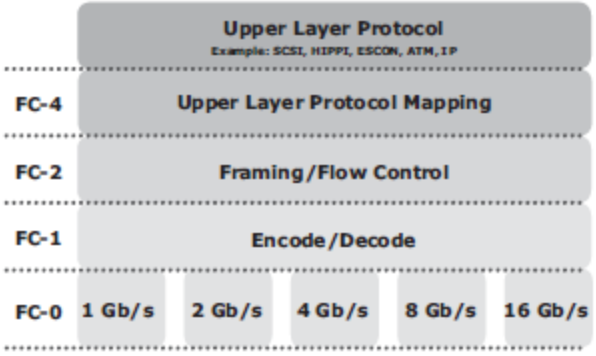

Image courtesy EMC Corporation, "Information Storage and Management"

FC-2 Layer

The FC-2 layer provides Fiber Channel addressing, structure, and organization of data (frames, sequences, and exchanges).

It also defines fabric services, classes of service, flow control, and routing.

Fiber Channel Protocol Stack

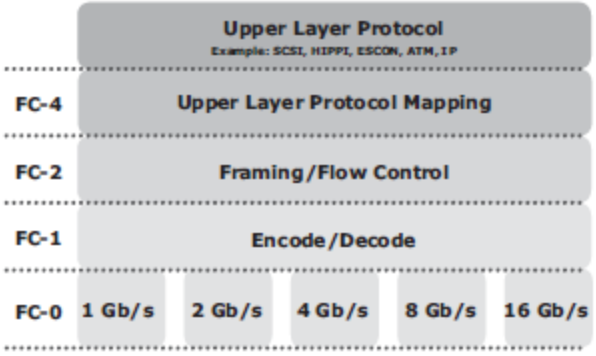

Image courtesy EMC Corporation, "Information Storage and Management"

FC-1 Layer

The FC-1 layer defines how data is encoded prior to transmission and decoded upon receipt.

At the transmitter node, an 8-bit character is encoded into a 10-bit transmissions character.

This character is then transmitted to the receiver node.

At the receiver node, the 10-bit character is passed to the FC-1 layer, which decodes the 10-bit character into the original 8-bit character.

Fiber Channel Protocol Stack

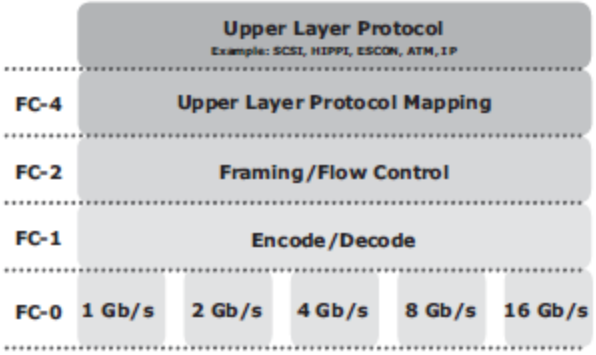

Image courtesy EMC Corporation, "Information Storage and Management"

FC-1 Layer

The FC-1 layer defines how data is encoded prior to transmission and decoded upon receipt.

At the transmitter node, an 8-bit character is encoded into a 10-bit transmissions character.

This character is then transmitted to the receiver node.

At the receiver node, the 10-bit character is passed to the FC-1 layer, which decodes the 10-bit character into the original 8-bit character.

Fiber Channel Protocol Stack

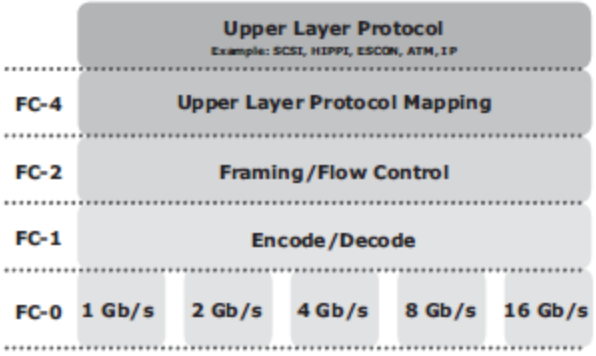

Image courtesy EMC Corporation, "Information Storage and Management"

FC-0 Layer

FC-0 is the lowest layer in the FCP stack.

This layer defines the physical interface, media, and transmission of bits. The FC-0 specification includes cables, connectors, and optical and electrical parameters for a variety of data rates.

The FC transmission can use both electrical and optical media.

Fiber Channel Protocol Stack

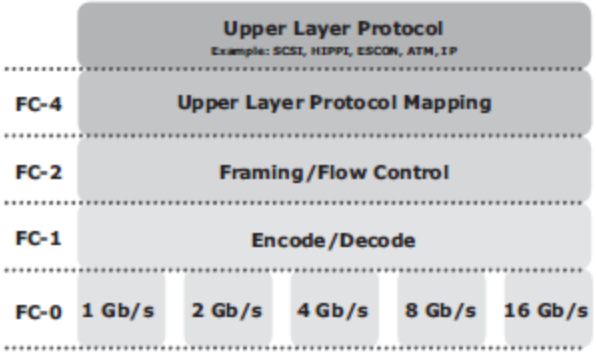

Image courtesy EMC Corporation, "Information Storage and Management"

FC-0 Layer

FC-0 is the lowest layer in the FCP stack.

This layer defines the physical interface, media, and transmission of bits. The FC-0 specification includes cables, connectors, and optical and electrical parameters for a variety of data rates.

The FC transmission can use both electrical and optical media.

Fiber Channel Addressing

An FC address is dynamically assigned when a node port logs on to the fabric. The addressing mechanism provided here corresponds to the fabric with the switch as an interconnecting device.

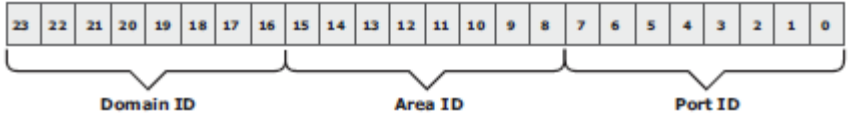

Image courtesy EMC Corporation, "Information Storage and Management"

Fiber Channel Addressing

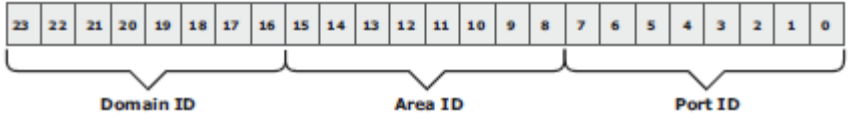

Image courtesy EMC Corporation, "Information Storage and Management"

Domain ID

A domain ID is a unique number provided to each switch in the fabric. Although this is an 8-bit field, there are only 239 available addresses for domain ID because some addresses are deemed special and reserved for fabric management services.

Area ID

The area ID is used to identify a group of switch ports used for connecting nodes.

Port ID

the port ID, identifies the port within the group.

World Wide Names

Each device in the FC environment is assigned a 64-bit unique identifier called the World Wide Name (WWN).

The Fiber Channel environment uses two types of WWNs,

1. World Wide Node Name (WWNN)
2. World Wide Port Name (WWPN)

WWNs are similar to the Media Access Control (MAC) addresses used in IP networking. WWNs are burned into the hardware or assigned through software.

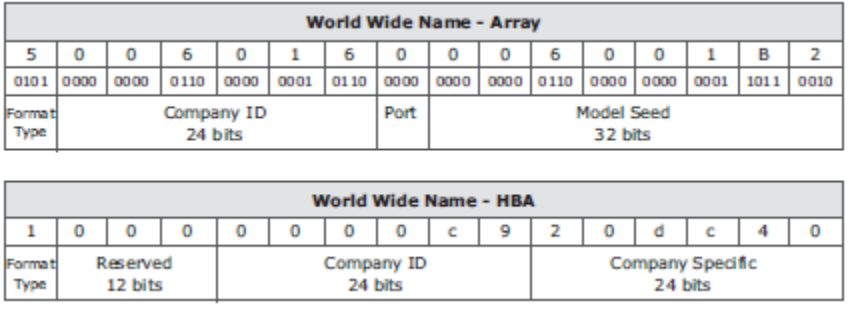

World Wide Name - Array															
5	0	0	6	0	1	6	0	0	0	6	0	0	1	B	2
0101	0000	0000	0110	0000	0001	0110	0000	0000	0000	0110	0000	0000	0001	1011	0010
Format Type	Company ID 24 bits						Port	Model Seed 32 bits							

World Wide Name - HBA															
1	0	0	0	0	0	0	0	c	9	2	0	d	c	4	0
Format Type	Reserved 12 bits			Company ID 24 bits						Company Specific 24 bits					

Image courtesy EMC Corporation, "Information Storage and Management"

FC Frame

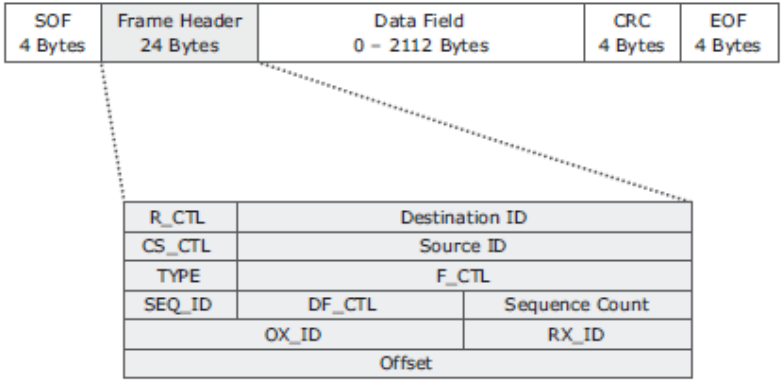

Image courtesy EMC Corporation, "Information Storage and Management"

An FC frame consists of five parts,

1. Start of Frame (SOF)
2. Frame Header
3. Data Field
4. Cyclic Redundancy Check (CRC)
5. End of Frame (EOF)

FC Frame

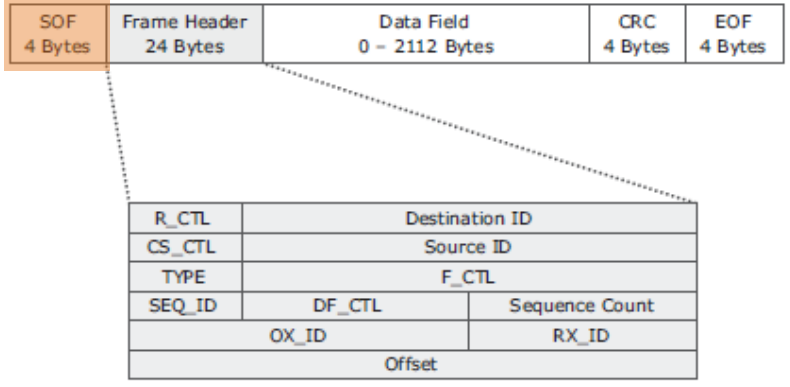

Image courtesy EMC Corporation, "Information Storage and Management"

1. Start of Frame (SOF)

It acts as an delimiter, also it indicates whether the frame is the first frame in a sequence of frames.

FC Frame

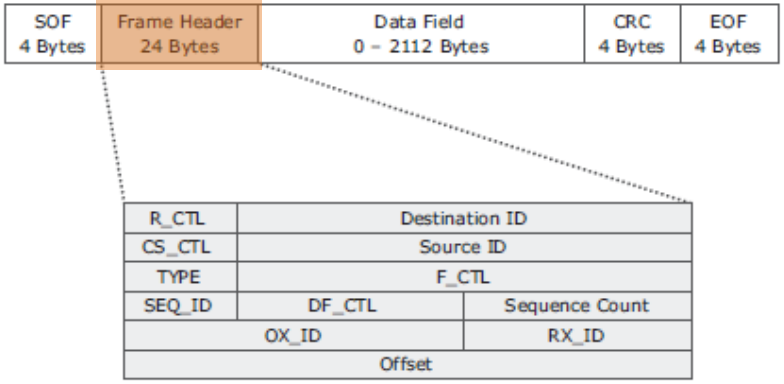

Image courtesy EMC Corporation, "Information Storage and Management"

2. *Frame Header*

The frame header also defines

i. Routing Control (R_CTL)

This field denotes whether the frame is a link control frame or a data frame.

ii. Class Specific Control (CS_CTL)

This field specifies link speeds for class 1 and class 4 data transmission.

iii. TYPE

This field describes the upper layer protocol (ULP) to be carried on the frame if it is a data frame.

iv. Data Field Control (DF_CTL)

A 1-byte field that indicates the existence of any optional headers at the beginning of the data payload.

v. Frame Control (F_CTL)

A 3-byte field that contains control information related to frame content.

FC Frame

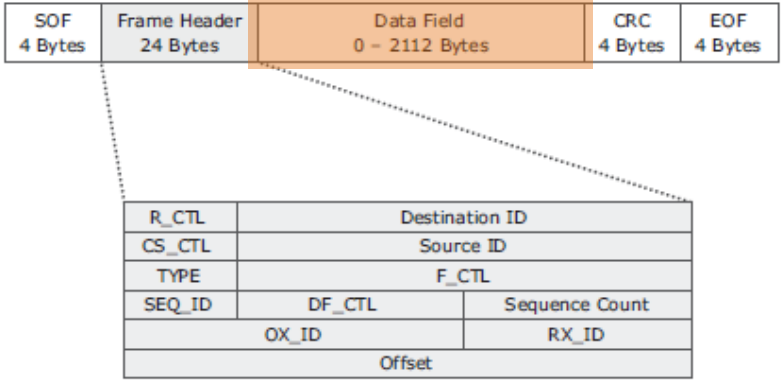

Image courtesy EMC Corporation, "Information Storage and Management"

3. *Data Field*

The data field in an FC frame contains the data payload, up to 2,112 bytes of actual data with 36 bytes of fixed overhead.

FC Frame

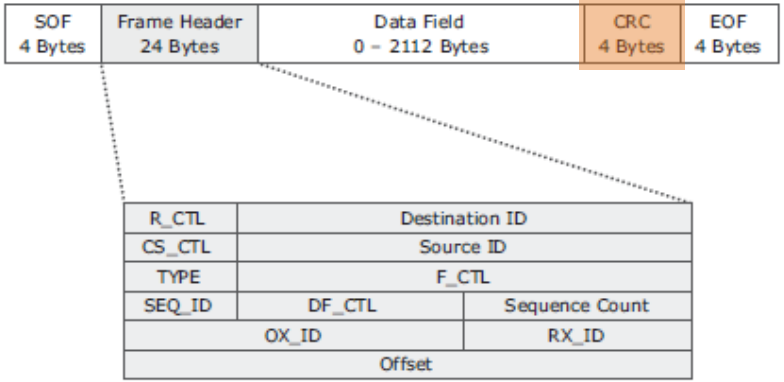

Image courtesy EMC Corporation, "Information Storage and Management"

4. Cyclic Redundancy Check (CRC)

The CRC checksum facilitates error detection for the content of the frame. This checksum verifies data integrity by checking whether the content of the frames are received correctly.

FC Frame

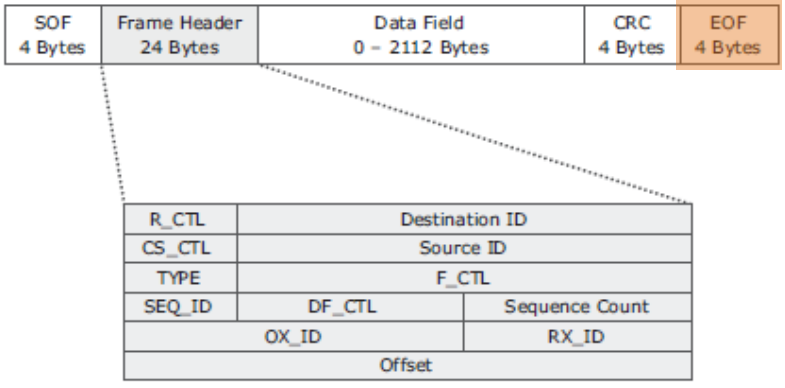

Image courtesy EMC Corporation, "Information Storage and Management"

5. End of Frame (EOF)

It acts as an delimiter, also it indicates end of the frame.

Fiber Channel Login Types

Fiber Channel Login Types contd...

Performed between an N_Port and an F_Port.
To log on to the fabric, a node sends a FLOGI frame with the WWNN and WWPN parameters to the login service at the predefined FC address FFFFFE.
In turn, the switch accepts the login and returns an Accept (ACC) frame with the assigned FC address for the node.
Immediately after the FLOGI, the N_Port registers itself with the local Name Server on the switch, indicating its WWNN, WWPN, port type, class of service, assigned FC address and so on.

Fiber Channel Login Types contd...

Performed between two N_Ports to establish a session.
The initiator N_Port sends a PLOGI request frame to the target N_Port, which accepts it.
The target N_Port returns an ACC to the initiator N_Port.
Next, the N_Ports exchange service parameters relevant to the session.

Fiber Channel Login Types contd...

Performed between two N_Ports.
This login relates to the FC-4 ULPs, such as SCSI.
If the ULP is SCSI, N_Ports exchange SCSI-related service parameters.

Zoning

Zoning is an FC switch function that enables node ports within the fabric to be logically segmented into groups and to communicate with each other within the group

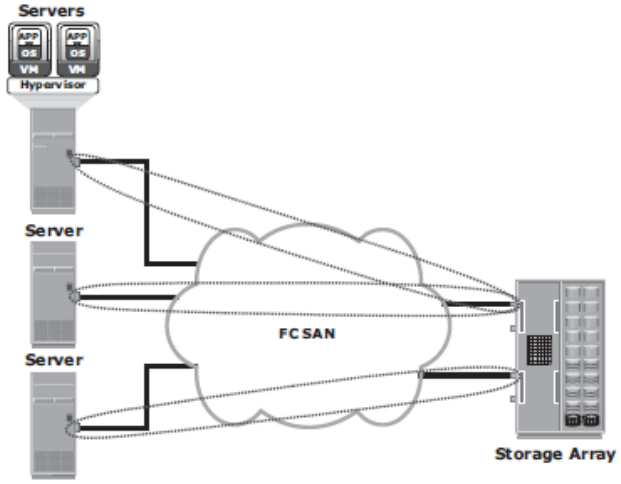

Image courtesy EMC Corporation, "Information Storage and Management"

Zoning contd...

Whenever a change takes place in the name server database, the fabric controller sends a Registered State Change Notification (RSCN) to all the nodes impacted by the change.
If zoning is not configured, the fabric controller sends an RSCN to all the nodes in the fabric.
Involving the nodes that are not impacted by the change results in increased fabric-management traffic.
In the presence of zoning, a fabric sends the RSCN to only those nodes in a zone where the change has occurred.

Zoning contd...

A zone set is composed of a group of zones that can be activated or deactivated as a single entity in a fabric.
Multiple zone sets may be defined in a fabric, but only one zone set can be active at a time.
Members are nodes within the SAN that can be included in a zone.
A port or node can be a member of multiple zones. Nodes distributed across multiple switches in a switched fabric may also be grouped into the same zone.
Zone sets are also referred to as zone configurations.

Types of Zoning

Port Zoning

In port zoning, access to node is determined by the physical switch port to which a node is connected. The zone members are the port identifier to which HBA and its targets are connected.

WWN Zoning

Uses World Wide Names to define zones. The zone members are the unique WWN addresses of the HBA and its targets.

Mixed Zoning

Combines the qualities of both WWN zoning and port zoning. Using mixed zoning enables a specific node port to be tied to the WWN of another node.

1. Mesh Topology

In a ***full mesh***, every switch is connected to every other switch in the topology. A full mesh topology may be appropriate when the number of switches involved is small.

In a ***partial mesh*** topology, several hops or ISLs may be required for the traffic to reach its destination. Partial mesh offers more scalability than full mesh topology. However, without proper placement of host and storage devices, traffic management in a partial mesh fabric might be complicated and ISLs could become overloaded due to excessive traffic aggregation.

2. Core-Edge Fabric

The ***edge tier*** is usually composed of switches and offers an inexpensive approach to adding more hosts in a fabric. Each switch at the edge tier is attached to a switch at the core tier through ISLs.

The ***core tier*** is usually composed of enterprise directors that ensure high fabric availability.

2. Core-Edge Fabric

In core-edge topology, the edge-tier switches are not connected to each other. The core-edge fabric topology increases connectivity within the SAN while conserving the overall port utilization.

If fabric expansion is required, additional edge switches are connected to the core. The core of the fabric is also extended by adding more switches or directors at the core tier.

Based on the number of core-tier switches, this topology has different variations, such as, single-core topology and dual-core topology.

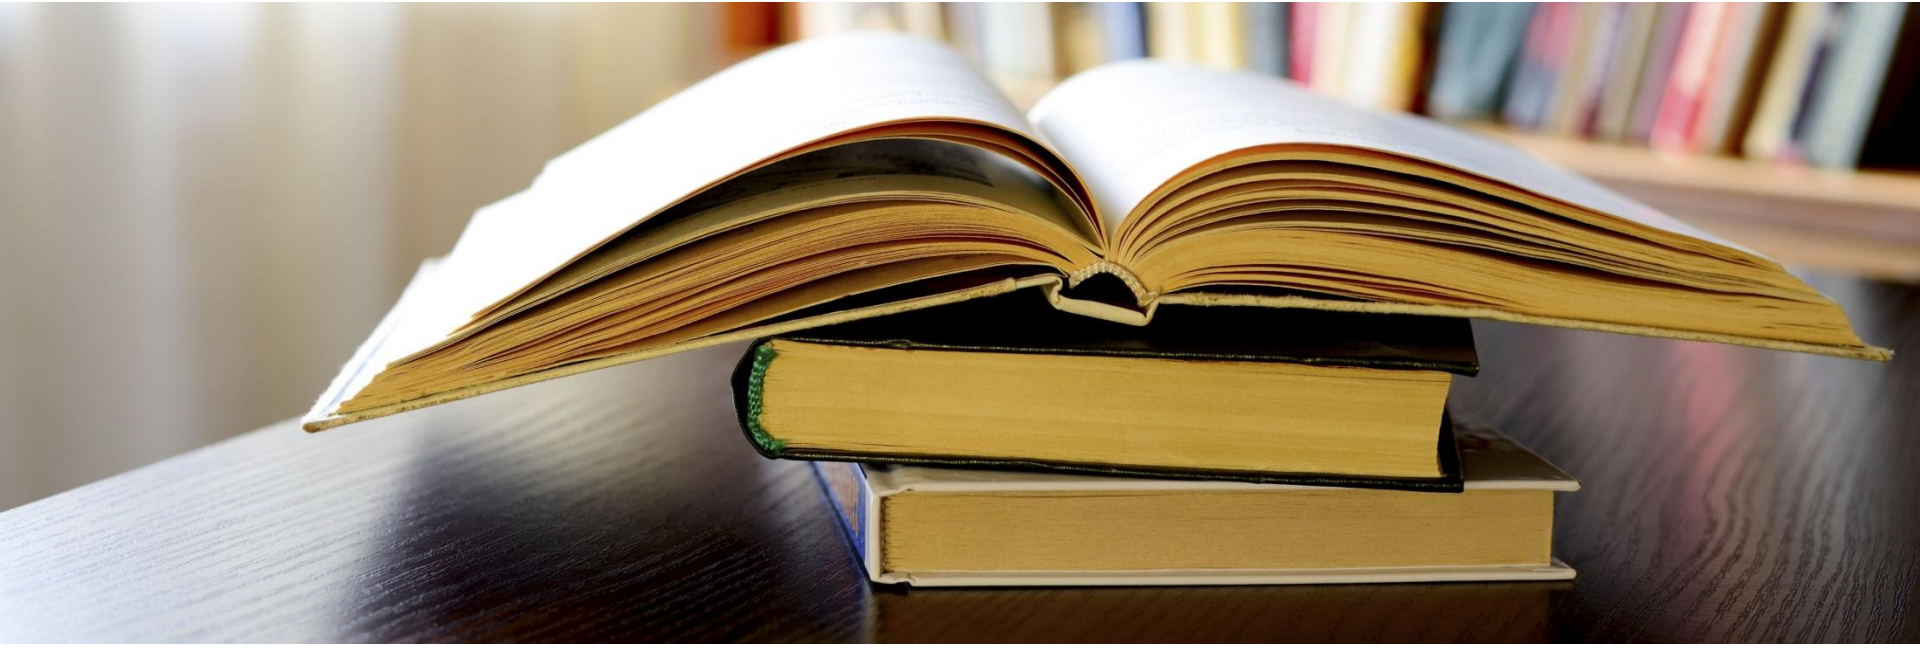

References

EMC Corporation, "Information Storage and Management" , Wiley, India.
Robert Spalding, "Storage Networks: The Complete Reference", Tata McGraw Hill, Osborne, 2003.

Varun C.M.
Assistant Professor, Department of Information Technology, St. Xavier's Catholic College of Engineering