CLOUD MANAGEMENT (UNIT -2)

Q.25 What is Fibre Channel? What is storage area network? Discuss the evolution of Fibre channel SAN from arbitrated loop to enterprise SAN.
Ans:

- **Fibre Channel**
  - Fibre Channel is a high-speed network technology that runs on high-speed optical fiber cables (preferred for front-end SAN connectivity) and serial copper cables (preferred for back-end disk connectivity).
  - The FC technology was created to meet the demand for increased speeds of data transfer among computers, servers, and mass storage subsystems.
  - Higher data transmission speeds are an important feature of the FC networking technology.
  - FC in full-duplex mode could sustain throughput of 200 MB/s. In comparison with Ultra-SCSI, FC is a significant leap in storage networking technology.
  - The FC architecture is highly scalable and single FC network can accommodate approximately 15 million nodes.
- **SAN and Its Evolution**
  - A storage area network (SAN) carries data between servers and storage devices through fibre channel switches.
  - A SAN enables storage consolidation and allows storage to be shared across multiple servers.

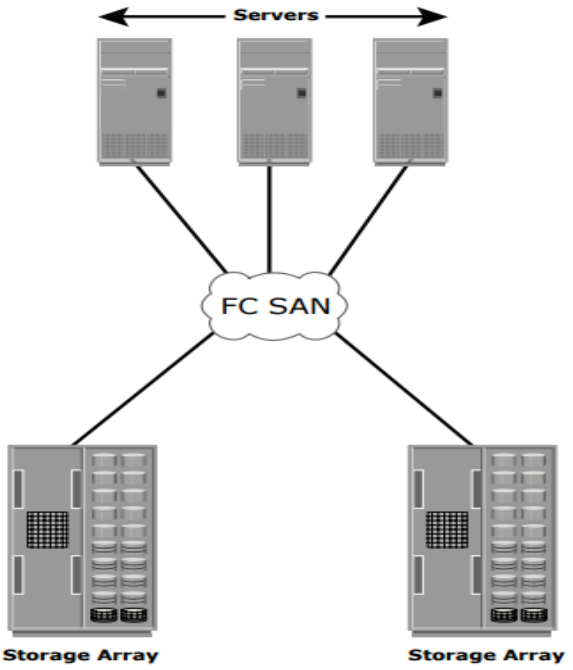

Figure: SAN Implementation

- A SAN provides the physical communication infrastructure and enables secure and robust communication between host and storage devices.
- In its earliest implementation, the SAN was a simple grouping of hosts and the associated storage that was connected to a network using a hub as a connectivity device.
- This configuration of a SAN is known as **a Fibre Channel Arbitrated Loop (FC-AL).**
- Use of hubs resulted in isolated FC-AL SAN islands because hubs provide limited connectivity and bandwidth.
- The switched fabric topologies improved connectivity and performance, which enabled SANs to be highly scalable.
- This enhanced data accessibility to applications across the enterprise.

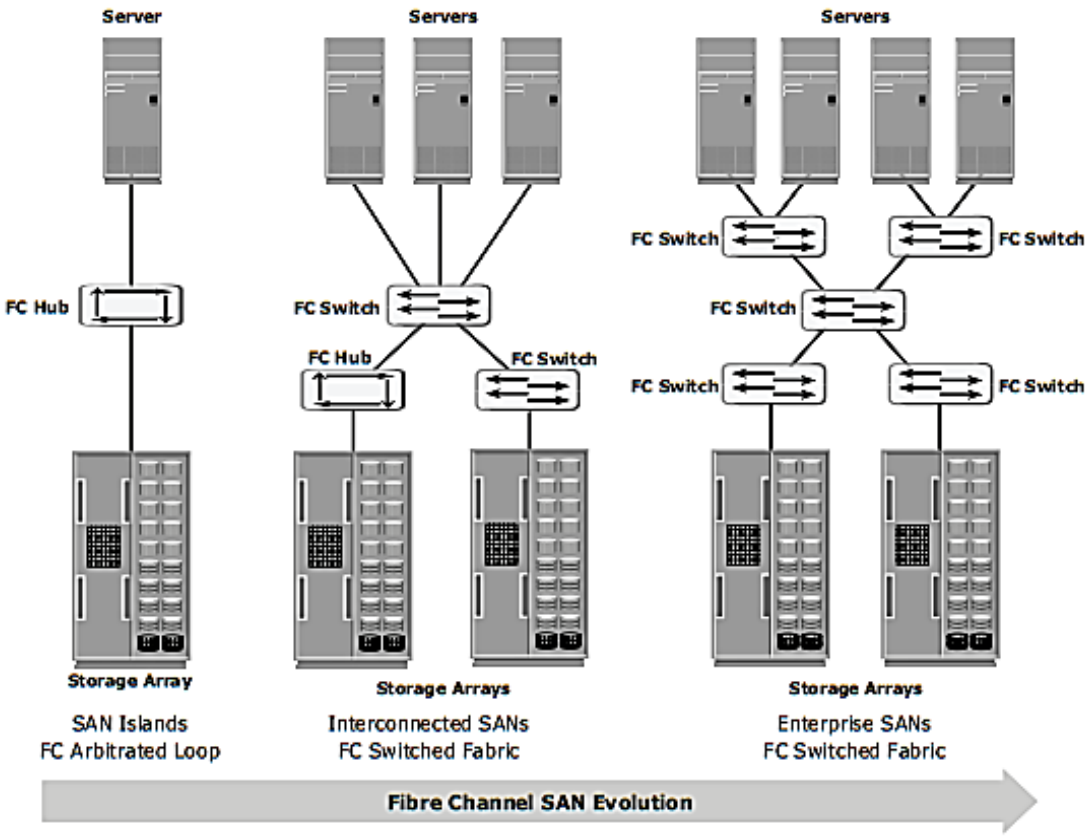

Q.26 What are the components of SAN? Explain in detail.
Ans:

- **Component of SAN**
  - A SAN consists of three basic components: servers, network infrastructure, and storage.
  - These components can be further broken down into the following key elements: node ports, cabling, interconnecting devices, storage arrays, and SAN management software.

1. Node Ports

- In fibre channel, devices such as hosts, storage and tape libraries are all referred to as nodes.
- Each node is a source or destination of information for one or more nodes.
- Each node requires one or more ports to provide a physical interface for communicating with other nodes.
- A port operates in full-duplex data transmission mode with a transmit (Tx) link and a receive (Rx) link.

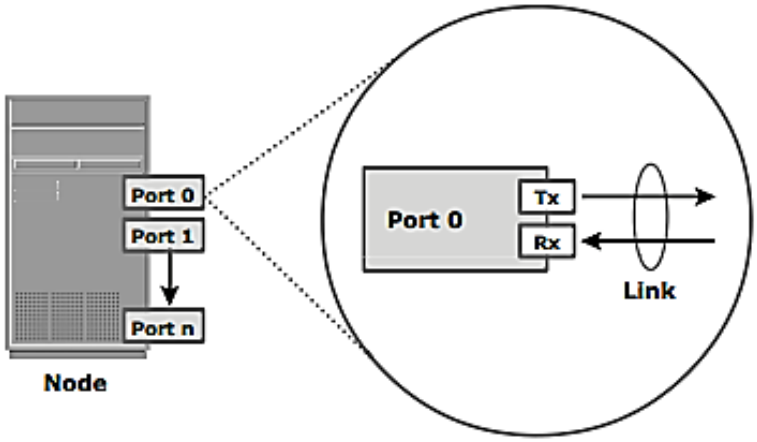

Figure: nodes, ports and links

2. Cabling

- SAN implementations use optical fiber cabling.
- Copper can be used for shorter distances for back-end connectivity, as it provides a better signal-to-noise ratio for distances up to 30 meters.
- Optical fiber cables carry data in the form of light. There are two types of optical cables, **multi-mode and single-mode.**
- **Multi-mode fiber (MMF)** cable carries multiple beams of light projected at different angles simultaneously onto the core of the cable.
- Based on the bandwidth, multi-mode fibers are classified as OM1 (62.5µm), OM2 (50µm) and laser optimized OM3 (50µm).

- **Single-mode fiber (SMF)** carries a single ray of light projected at the center of the core.
- These cables are available in diameters of 7–11 microns; the most common size is 9 microns.
- Among all types of fibre cables, single-mode provides minimum signal attenuation over maximum distance (up to 10 km).

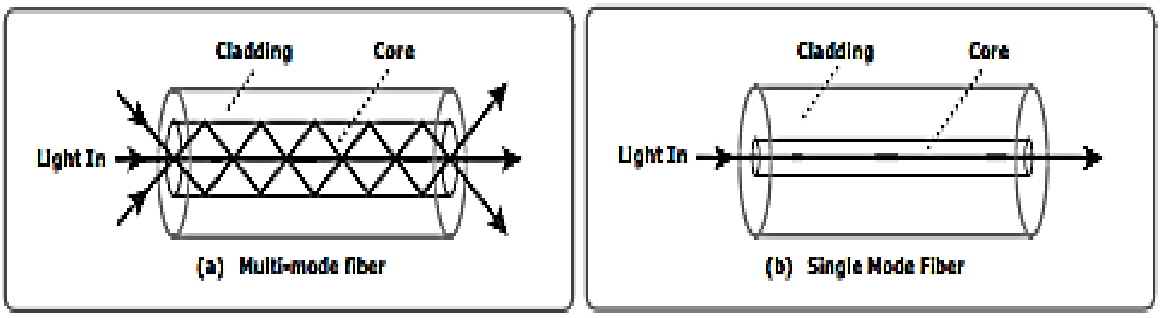

Figure: Multi-mode fiber and Single-mode fiber

3. Interconnect Devices

- Hubs, switches, and directors are the interconnect devices commonly used in SAN.
- **Hubs** are used as communication devices in FC-AL implementations.
- Hubs physically connect nodes in a logical loop or a physical star topology.
- All the nodes must share the bandwidth because data travels through all the connection points.
- **Switches** are more intelligent than hubs and directly route data from one physical port to another. Therefore, nodes do not share the bandwidth.
- **Director**s are larger than switches and are deployed for data center implementations.
- The function of directors is similar to that of FC switches, but directors have higher port count and fault tolerance capabilities.

4. Storage Arrays

- The fundamental purpose of a SAN is to provide host access to storage resources.
- The large storage capacities offered by modern storage arrays have been exploited in SAN environments for storage consolidation and centralization.
- SAN implementations complement the standard features of storage arrays by providing high availability and redundancy, improved performance, business continuity, and multiple host connectivity.

5. SAN Management Software

- 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 logical partitioning of the SAN, called zoning.

Q.27 Explain the three basic connectivity options for Fibre Channel. (or Explain Fibre Channel arbitrated loop.) (or Explain Fibre Channel switched fabric.)

Ans:

The FC architecture supports three basic interconnectivity options: point-to-point, arbitrated loop (FC-AL), and fabric connect.

1. Point-to-Point

- 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, as only two devices can communicate with each other at a given time.
- Standard DAS usess pointto-point connectivity.

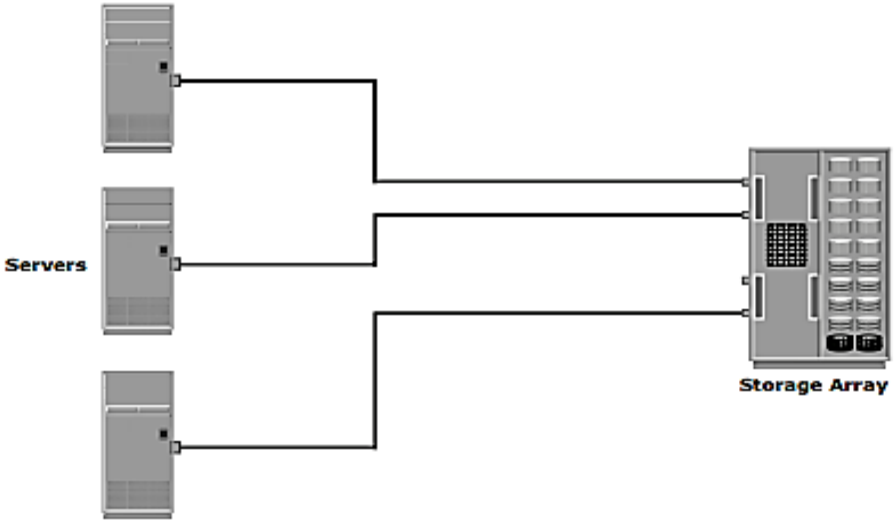

Figure: Point-to-point topology

2. Fibre Channel Arbitrated Loop

- In the FC-AL configuration, devices are attached to a shared loop.
- 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.

- FC-AL can be implemented without any interconnecting devices by directly connecting one device to another in a ring through cables.
- FC-AL can be implemented without any interconnecting devices by directly connecting one device to another in a ring through cables.
- However, FC-AL implementations may also use hubs whereby the arbitrated loop is physically connected in a star topology.

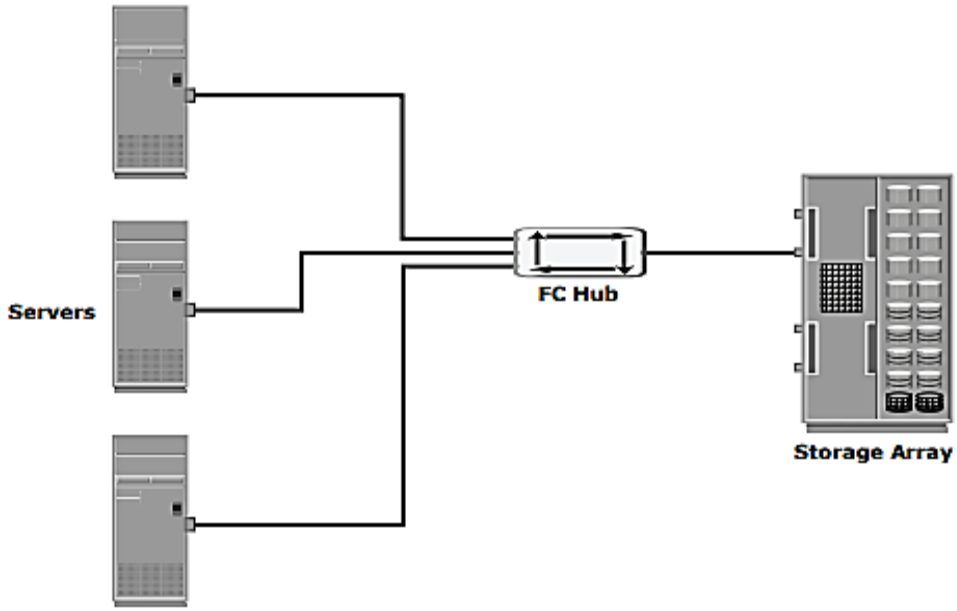

Figure: Fibre Channel arbitrated loop

3. Fibre Channel Switched Fabric

- Unlike a loop configuration, a Fibre Channel switched fabric (FC-SW) network provides interconnected devices, dedicated bandwidth, and scalability.
- 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.
- 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.
- A fabric topology can be described by the number of tiers it contains.
- The number of tiers in a fabric is based on the number of switches traversed between two points that are farthest from each other.

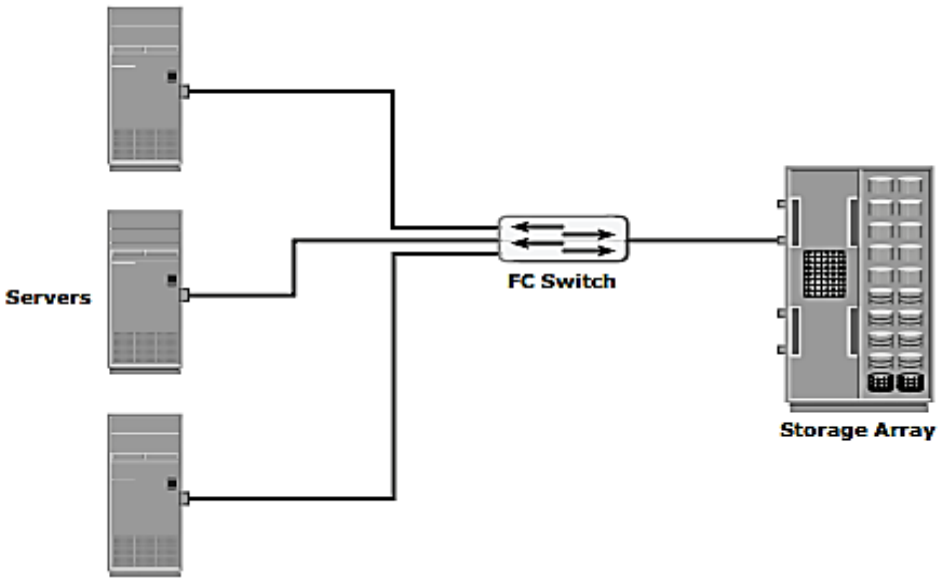

Figure: Fibre Channel switched fabric

Q.28 What are the different types of fibre channel ports? Explain.
Ans:

- Ports are the basic building blocks of an FC network.
- Ports on the switch can be one of the following types:

- **N_port:** An end point in the fabric. This port is also known as the node port. Typically, it is a host port (HBA) or a storage array port that is connected to a switch in a switched fabric.
- **NL_port:** A node port that supports the arbitrated loop topology. This port is also known as the **node loop port.**
- **E_port:** An FC 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 a link, which is called an **Inter-Switch Link (ISL).**
- **F_port:** A port on a switch that connects an N_port. It is also known as a fabric port.
- **FL_port**: A fabric port that participates in FC-AL. This port is connected to the NL_ports on an FC-AL loop. A FL_port also connects a loop to a switch in a switched fabric. As a result, all NL_ports in the loop can participate in FC-SW.
- **G_port:** A generic port that can operate as an E_port or an F_port and determines its functionality automatically during initialization.

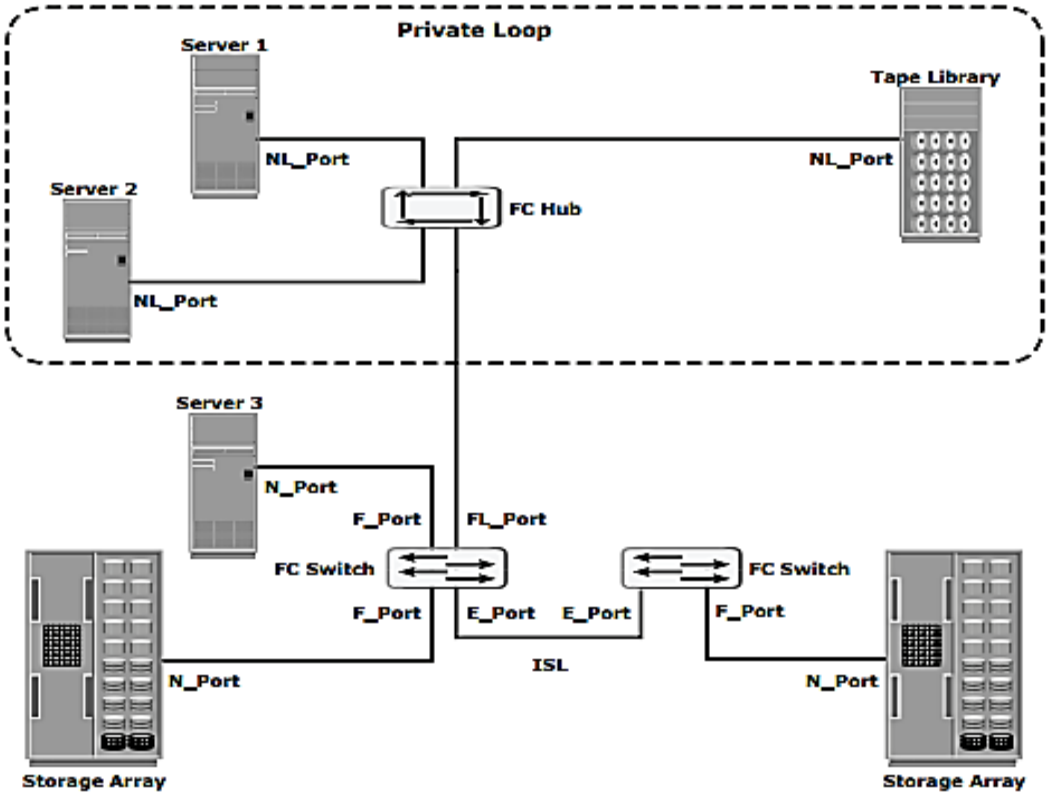

Figure: Fibre channel ports

Q.29 Explain the Fibre channel architecture. (or What is fibre channel protocol? Explain the fibre channel protocol stack.)

Ans:

- **Fibre channel architecture**
  - The FC architecture represents true channel/network integration with standard interconnecting devices. Connections in a SAN are accomplished using FC.
  - Traditionally, transmissions from host to storage devices are carried out over channel connections such as a parallel bus.
  - Channel technologies provide high levels of performance with low protocol overheads.
  - Such performance is due to the static nature of channels and the high level of hardware and software integration provided by the channel technologies.
- **Fibre channel protocol**
  - Fibre Channel Protocol (FCP) is the implementation of serial SCSI-3 over an FC network.

- In the FCP architecture, all external and remote storage devices attached to the SAN appear as local devices to the host operating system.
- FCP Sustained transmission bandwidth over long distances.

➢ **Fibre Channel Protocol Stack**

- FCP defines the communication protocol in five layers: FC-0 through FC-4 except FC-3 layer, which is not implemented.

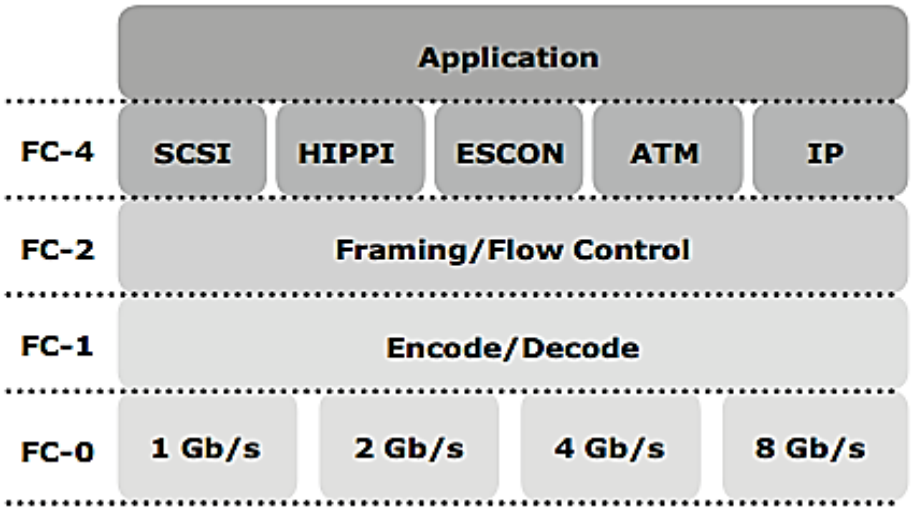

Figure: Fibre channel protocol stack

1. FC-4 Upper Layer Protocol

- 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, HIPPI Framing Protocol, Enterprise Storage Connectivity (ESCON), ATM, and IP.

2. FC-2 Transport Layer

- The FC-2 is the transport layer that contains the payload, addresses of the source and destination ports, and link control information.
- The FC-2 layer provides Fibre Channel addressing, structure, and organization of data (frames, sequences, and exchanges).
- It also defines fabric services, classes of service, flow control, and routing.

3. FC-1 Transmission Protocol

- This layer defines the transmission protocol that includes serial encoding and decoding rules, special characters used, and error control.

- 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.

4. FC-0 Physical Interface

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

Q.30 Explain fibre channel addressing.

Ans:

- An FC address is dynamically assigned when a port logs on to the fabric.
- These ports can be an N_port and an NL_port in a public loop, or an NL_port in a private loop.

➢ **FC address of N_port**

- The first field of the FC address of an N_port contains the **domain ID** of the switch.
- This is an 8-bit field. Out of the possible 256 domain IDs, 239 are available for use; the remaining 17 addresses are reserved for specific services.
- The **area ID** is used to identify a group of F_ports. An
- The last field in the FC address identifies the F_port within the group.

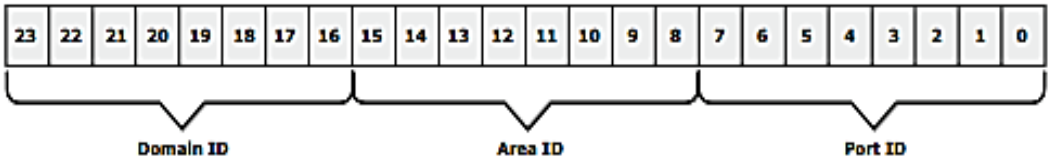

Figure: 24-bit FC address of N_port

➢ **FC Address of an NL_port**

- The FC addressing scheme for an NL_port differs from other ports.
- The two upper bytes in the FC addresses of the NL_ports in a private loop are assigned zero values.
- However, when an arbitrated loop is connected to a fabric through an FL_port, it becomes a public loop.
- The last field in the FC addresses of the NL_ports, in both public and private loops, identifies the AL-PA. There are 127 allowable AL-PA addresses.

- **World Wide Names**
  - Each device in the FC environment is assigned a 64-bit unique identifier called the World Wide Name (WWN).
  - The Fibre Channel environment uses two types of WWNs: World Wide Node Name (WWNN) and World Wide Port Name (WWPN).

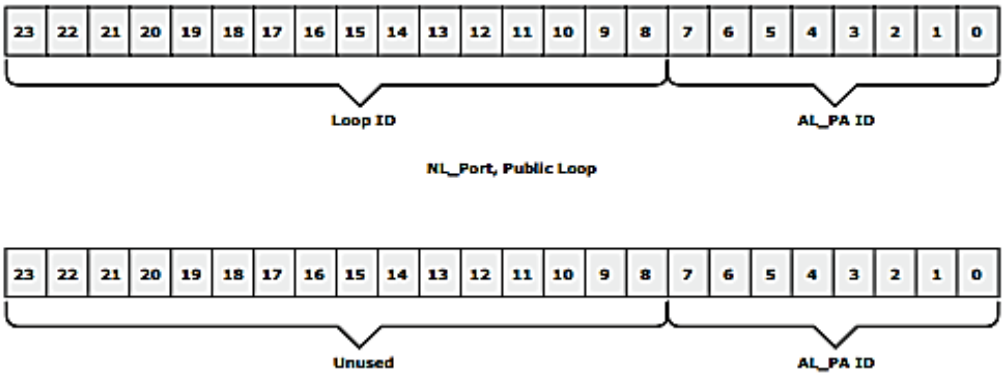

World Wide Name - Array															
5	0	0	6	0	1	6	0	0	0	6	0	0	1	B	2
0101	0000	0000	0110	0000	0001	0110	0000	0000	0000	0110	0000	0000	0001	1011	0010
	Company ID 24 bits						Port	Model Seed 32 bits							

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

Figure: World Wide Names

Q.31 Explain the fibre channel frame.

Ans:

- An FC frame consists of five parts: start of frame (SOF), frame header, data field, cyclic redundancy check (CRC), and end of frame (EOF).

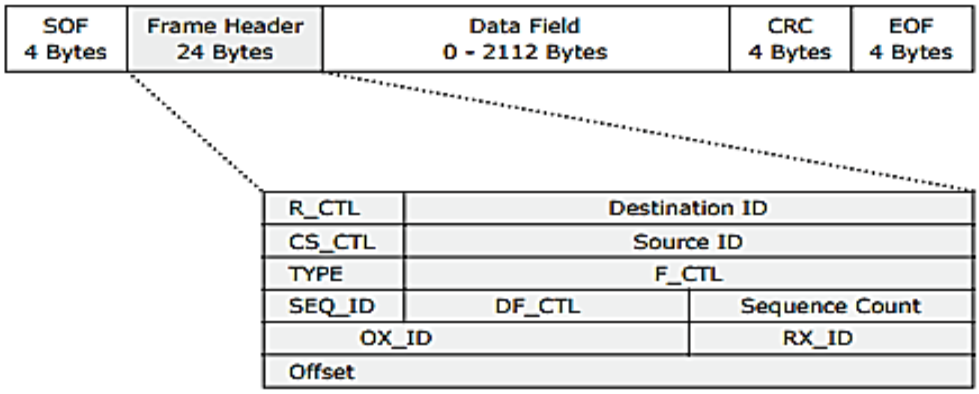

Figure: FC frame

- The SOF and EOF act as delimiters
- The **SOF** is a flag that indicates whether the frame is the first frame in a sequence of frames.
- The **frame header** is 24 bytes long and contains addressing information for the frame.
- It includes the following information: Source ID (S_ID), Destination ID (D_ID), Sequence ID (SEQ_ID), Sequence Count (SEQ_CNT), Originating Exchange ID (OX_ID), and Responder Exchange ID (RX_ID).
- The S_ID and D_ID are standard FC addresses for the source port and the destination port, respectively.
- The SEQ_ID and OX_ID identify the frame as a component of a specific sequence and exchange, respectively.
- The **data field** in an FC frame contains the data payload, up to 2,112 bytes of original data — in most cases, SCSI data.
- 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 was received correctly.

Q.32 Discuss the data transport in fibre channel network.

Ans:

- In an FC network, data transport is analogous to a conversation between two people, whereby a frame represents a word, a sequence represents a sentence, and an exchange represents a conversation.

➢ **Exchange operation**

- An exchange operation enables two N_ports to identify and manage a set of information units. This unit maps to a sequence. Sequences can be both unidirectional and bidirectional depending upon the type of data sequence exchanged between the initiator and the target.

➢ **Sequence**

- A sequence refers to a contiguous set of frames that are sent from one port to another. A sequence corresponds to an information unit, as defined by the ULP.

➢ **Frame**

- A frame is the fundamental unit of data transfer at Layer 2. Each frame can contain up to 2,112 bytes of payload.

Q.33 What are the two flow control mechanisms in fibre channel technology? Explain.

Ans:

- Flow control defines the pace of the flow of data frames during data transmission.
- FC technology uses two flow-control mechanisms: buffer-to-buffer credit (BB_Credit) and end-to-end credit (EE_Credit).

1. BB_Credit

- FC uses the BB_Credit mechanism for hardware-based flow control.
- BB_Credit controls the maximum number of frames that can be present over the link at any given point in time.
- In a switched fabric, BB_Credit management may take place between any two FC ports.
- The BB_Credit mechanism provides frame acknowledgment through the Receiver Ready (R_RDY) primitive.

2. EE_Credit

- The function of end-to-end credit, known as EE_Credit, is similar to that of BB_ Credit.
- When an initiator and a target establish themselves as nodes communicating with each other, they exchange the EE_Credit parameters.

Q.34 Explain the different classes of service defined by fibre channel standard.

Ans:

- The FC standards define different classes of service to meet the requirements of a wide range of applications.
- The table below shows three classes of services and their features.

	CLASS 1	CLASS 2	CLASS 3
Communication type	Dedicated connection	Nondedicated connection	Nondedicated connection
Flow control	End-to-end credit	End-to-end credit B-to-B credit	B-to-B credit
Frame delivery	In order delivery	Order not guaranteed	Order not guaranteed
Frame acknowl-edgement	Acknowledged	Acknowledged	Not acknowledged
Multiplexing	No	Yes	Yes
Bandwidth utilization	Poor	Moderate	High

Figure: FC Class of Services

- Another class of services is class F, which is intended for use by the switches communicating through ISLs.
- Class F is similar to Class 2, and it provides notification of nondelivery of frames.
- Other defined Classes 4, 5, and 6 are used for specific applications.

Q.35 What is zoning? What are different categories of zoning? Explain.
Ans:

- Zoning is an FC switch function that enables nodes within the fabric to be logically segmented into groups that can communicate with each other.

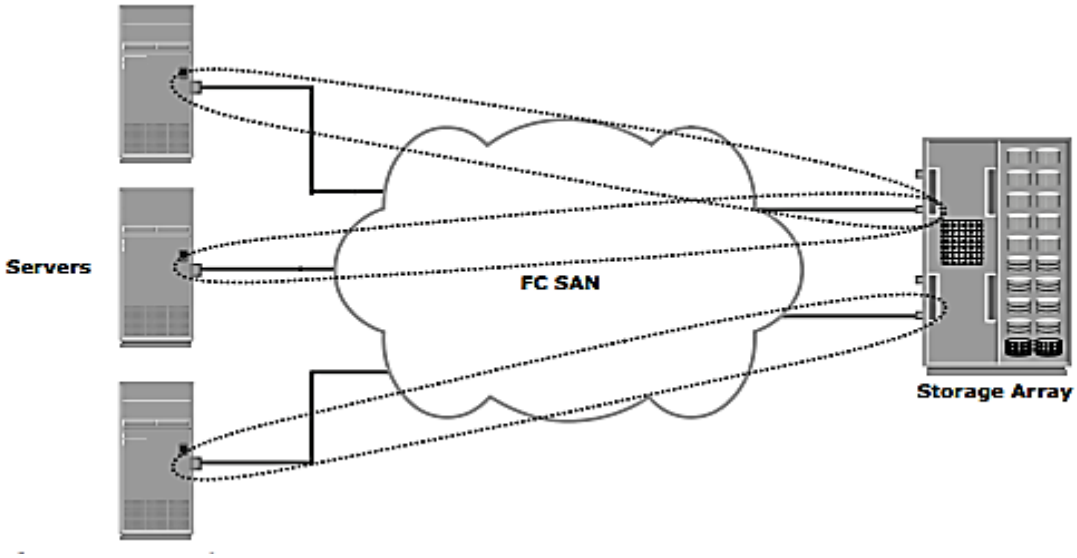

Figure: Zoning

- Multiple zone sets may be defined in a fabric, but only one zone set can be active at a time.
- A zone set is a set of zones and a zone is a set of members.
- A member may be in multiple zones. Members, zones, and zone sets form the hierarchy defined in the zoning process.

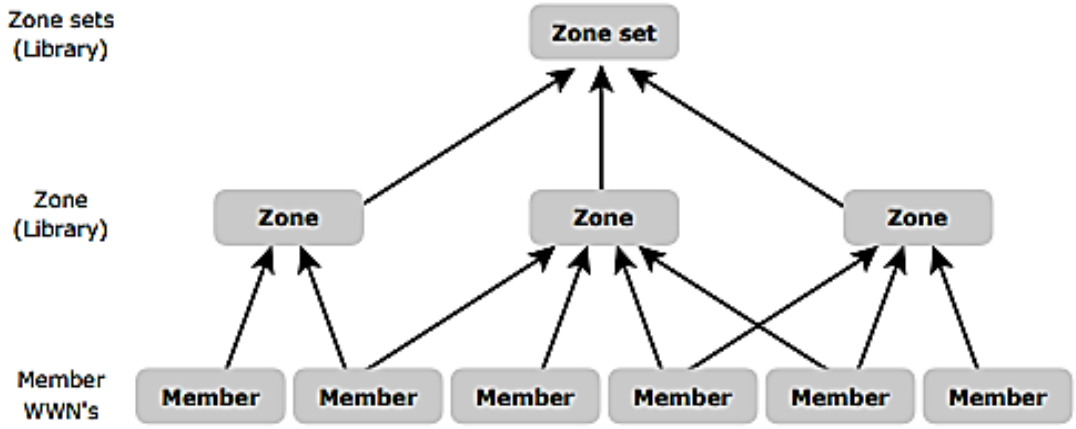

Figure: Members, zones, and zone sets

- **Types of Zoning**
  - Zoning can be categorized into three types:

1. Port zoning

- It uses the FC addresses of the physical ports to define zones.
- The FC address is dynamically assigned when the port logs on to the fabric.
- Therefore, any change in the fabric configuration affects zoning.
- Port zoning is also called **hard zoning**, although this method is secure.

2. WWN zoning

- It uses World Wide Names to define zones.
- WWN zoning is also referred to as soft zoning.
- A major advantage of WWN zoning is its flexibility.

3. Mixed zoning

- It combines the qualities of both WWN zoning and port zoning.
- Using mixed zoning enables a specific port to be tied to the WWN of a node.

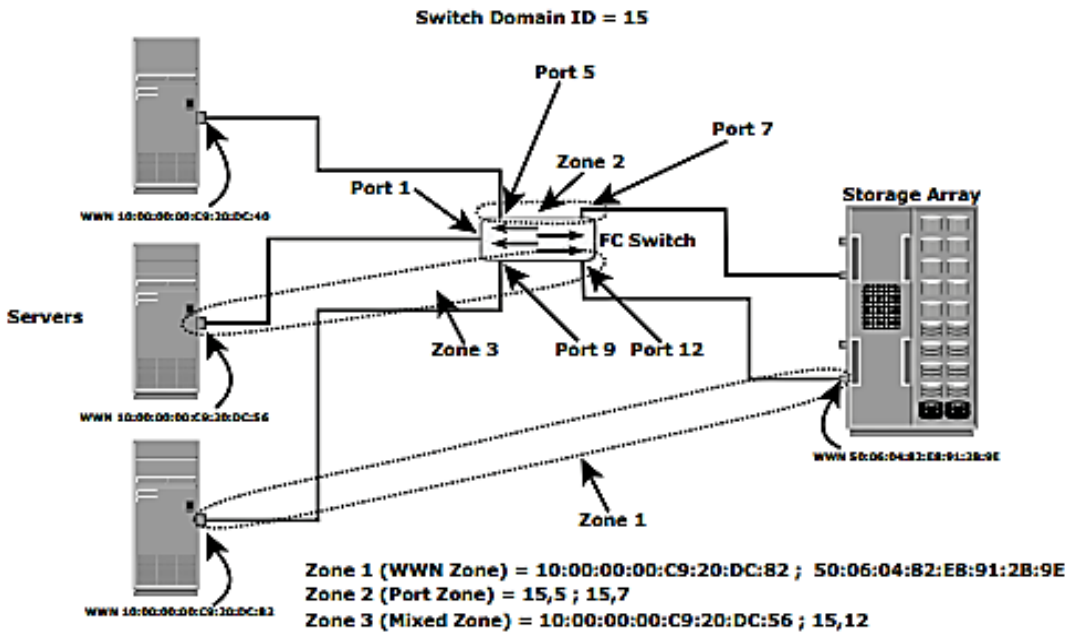

Figure: Types of zoning

Q.36 What are the three login types in Fibre channel? Explain.

Ans: Fabric services define three login types:

1. Fabric login (FLOGI)

- It is performed between an N_port and an F_port.
- To log on to the fabric, a device sends a FLOGI frame with the World Wide Node Name (WWNN) and World Wide Port Name (WWPN) parameters to the login service at the well-known FC address FFFFFE.
- Immediately after the FLOGI, the N_port registers itself with the local name server on the switch, indicating its WWNN, WWPN, and assigned FC address.

2. Port login (PLOGI)

- It is performed between an N_port and another N_port 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.

3. Process login (PRLI)

- It is also performed between an N_port and another N_ port.
- This login relates to the FC-4 ULPs such as SCSI. N_ports exchange SCSI-3-related service parameters.

Q.37 Explain the different topologies of Fibre Channel.
Ans:

- Fabric design follows standard topologies to connect devices.
- Core-edge fabric is one of the popular topology designs.
- Variations of core-edge fabric and mesh topologies are most commonly deployed in SAN implementations.

➢ **Core-Edge Fabric**

- In the core-edge fabric topology, there are two types of switch tiers in this fabric.
- The edge tier usually comprises switches and offers an inexpensive approach to adding more hosts in a fabric.
- The nodes on the edge can communicate with each other.
- The core tier usually comprises enterprise directors that ensure high fabric availability.
- The core-edge fabric topology increases connectivity within the SAN while conserving overall port utilization.
- This topology can have different variations.
- In a single-core topology, all hosts are connected to the edge tier and all storage is connected to the core tier.

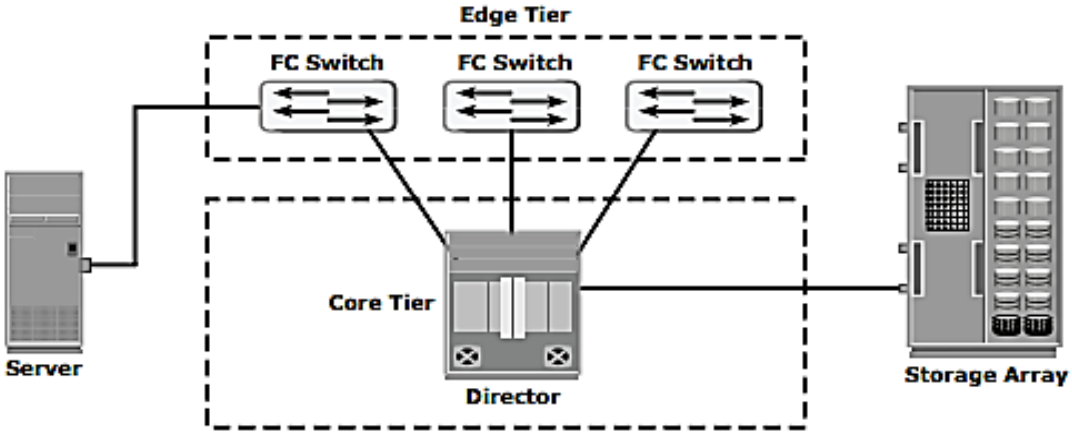

Figure: Single core topology

- A dual-core topology can be expanded to include more core switches. However, to maintain the topology, it is essential that new ISLs are created to connect each edge switch to the new core switch that is added.

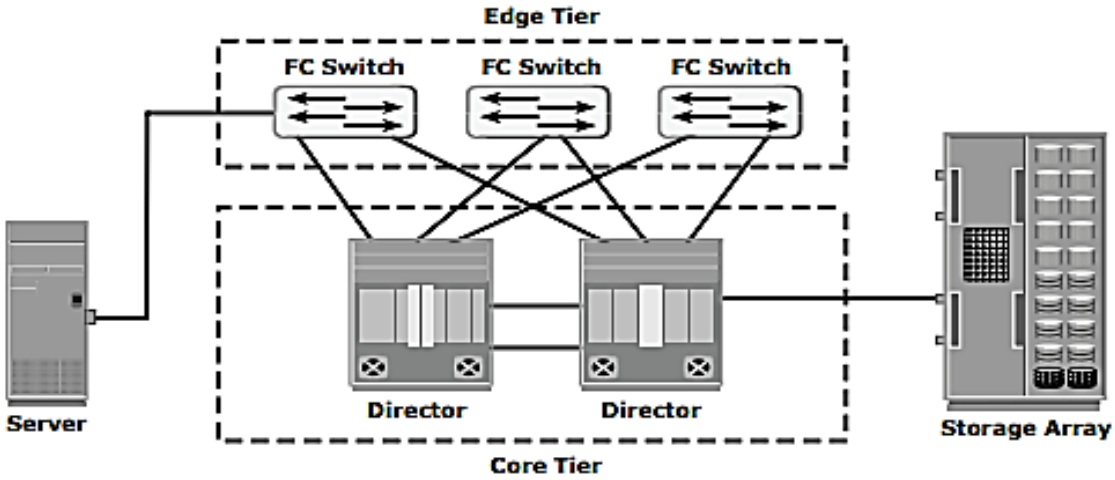

Figure: Dual-core topology

- **Mesh Topology**
  - In a mesh topology, each switch is directly connected to other switches by using ISLs.
  - This topology promotes enhanced connectivity within the SAN.
  - When the number of ports on a network increases, the number of nodes that can participate and communicate also increases.
  - A mesh topology may be one of the two types: full mesh or partial mesh.
  - In a full mesh, every switch is connected to every other switch in the topology.
  - In a partial mesh topology, several hops or ISLs may be required for the traffic to reach its destination.

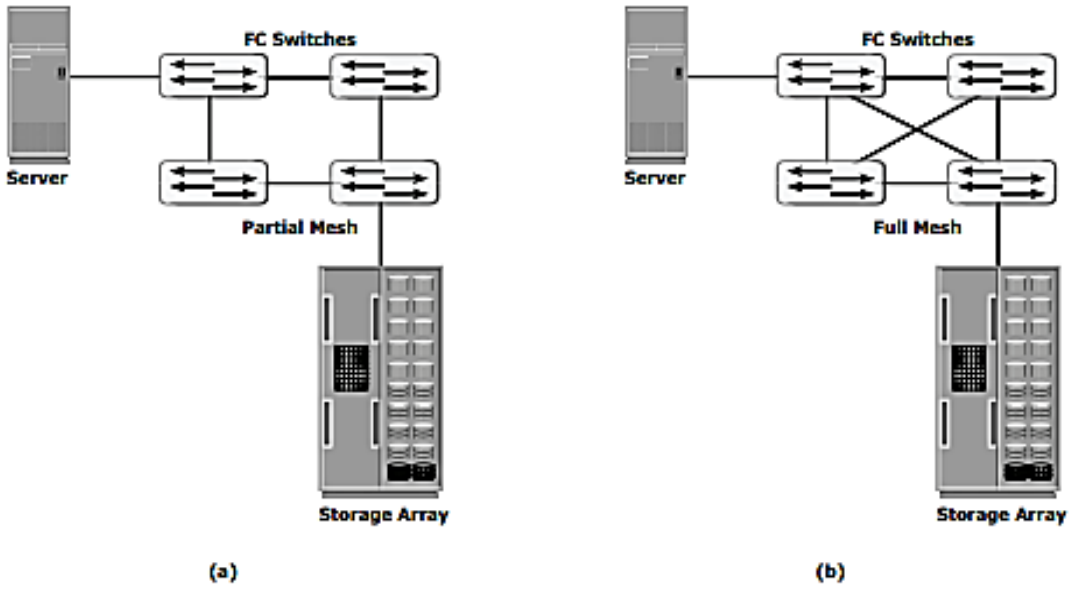

Figure: Partial mesh and full mesh topologies

Q.38 What is network attached storage? What are its benefits?

Ans:

- Network-attached storage (NAS) is an IP-based file-sharing device attached to a local area network.
- NAS provides the advantages of server consolidation by eliminating the need for multiple file servers.
- It provides storage consolidation through file-level data access and sharing.
- NAS uses network and file-sharing protocols to perform filing and storage functions.
- These protocols include TCP/IP for data transfer and CIFS and NFS for remote file service.
- To enable data sharing, NAS typically uses NFS for UNIX, CIFS for Windows, and File Transfer Protocol (FTP) and other protocols for both environments.
- A NAS device uses its own operating system and integrated hardware, software components to meet specific file service needs.
- A NAS device is a dedicated, high-performance, high-speed, single-purpose file serving and storage system.
- NAS serves a mix of clients and servers over an IP network.

➢ Benefits of NAS

1. Supports comprehensive access to information

- Enables efficient file sharing and supports many-to-one and one-to-many configurations.
- The many-to-one configuration enables a NAS device to serve many clients simultaneously.
- The one-to-many configuration enables one client to connect with many NAS devices simultaneously.

2. Improved efficiency

- Eliminates bottlenecks that occur during file access from a general-purpose file server because NAS uses an operating system specialized for file serving.

3. Improved flexibility

- Compatible for clients on both UNIX and Windows platforms using industry-standard protocols.
- NAS is flexible and can serve requests from different types of clients from the same source.

4. Centralized storage

- Centralizes data storage to minimize data duplication on client workstations, simplify data management, and ensures greater data protection.

5. Simplified management

- Provides a centralized console that makes it possible to manage file systems efficiently.

6. Scalability

- Scales well in accordance with different utilization profiles and types of business applications because of the high performance and low-latency design.

7. High availability

- Offers efficient replication and recovery options, enabling high data availability.
- NAS uses redundant networking components that provide maximum connectivity options.

8. Security

- Ensures security, user authentication, and file locking in conjunction with industry-standard security schemas.

Q.39 Explain the NAS I/O process with the help of a diagram.

Ans:

- NAS uses file-level access for all of its I/O operations.

1. File Systems and Remote File Sharing

- A file system is a structured way of storing and organizing data files.
- Many file systems maintain a file access table to simplify the process of finding and accessing files.

2. Accessing a File System

- A file system must be mounted before it can be used.
- In most cases, the operating system mounts a local file system during the boot process.
- The mount process creates a link between the file system and the operating system. When mounting a file system, the operating system organizes files and directories in a tree-like structure and grants the user the privilege of accessing this structure.
- The tree is rooted at a mount point that is named using operating system conventions.
- Users and applications can traverse the entire tree from the root to the leaf nodes.
- Files are located at leaf nodes, and directories and subdirectories are located at intermediate roots.
- The relationship between the user and the file system terminates when the file system is unmounted.

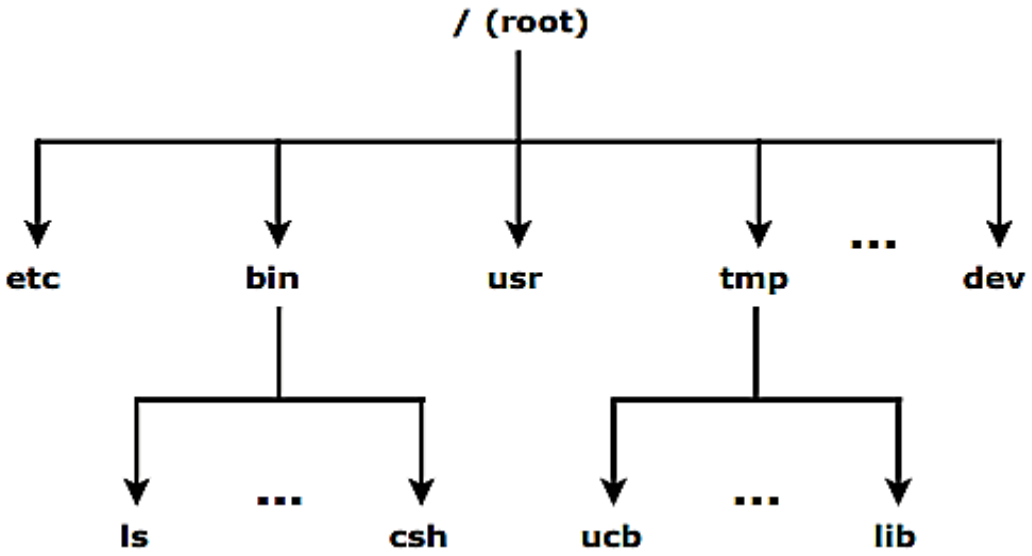

Figure: UNIX directory structure

3. File Sharing

- File sharing refers to storing and accessing data files over a network.
- In a file sharing environment, a user who creates the file determines the type of access to be given to other users (read, write, execute, append, delete, and list) and controls changes to the file.

Q.40 What are the components of network attached storage? Explain with the help of a diagram.

Ans: A NAS device has the following components:

- NAS head (CPU and Memory)
- One or more network interface cards (NICs), which provide connectivity to the network. Examples of NICs include Gigabit Ethernet, Fast Ethernet, ATM, and Fiber Distributed Data Interface (FDDI).
- An optimized operating system for managing NAS functionality.
- NFS and CIFS protocols for file sharing
- Industry-standard storage protocols to connect and manage physical disk resources, such as ATA, SCSI, or FC.

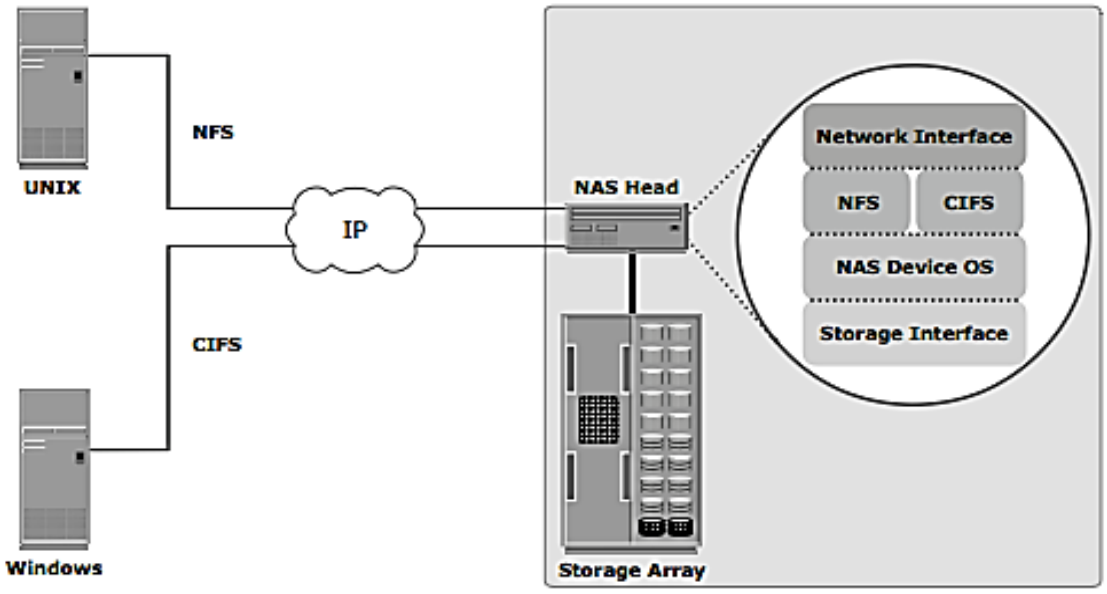

Figure: Components of NAS

Q.41 What are the two types of NAS implementation? Explain in detail

Ans: There are two types of NAS implementations: integrated and gateway.

- **Integrated NAS**
  - An integrated NAS device has all the components of NAS, such as the NAS head and storage, in a single enclosure, or frame.
  - The NAS head connects to the IP network to provide connectivity to the clients and service the file I/O requests.
  - The storage consists of a number of disks that can range from low-cost ATA to high throughput FC disk drives.
  - Management software manages the NAS head and storage configurations.
  - An integrated NAS solution ranges from a low-end device, which is a single enclosure, to a high-end solution that can have an externally connected storage array.
  - A low-end appliance-type NAS solution is suitable for applications that a small department may use.
  - In a high-end NAS solution, external and dedicated storage can be used.
- **Gateway NAS**
  - A gateway NAS device consists of an independent NAS head and one or more storage arrays.

- The NAS head performs the same functions that it does in the integrated solution; while the storage is shared with other applications that require block-level I/O.
- The gateway NAS is the most scalable because NAS heads and storage arrays can be independently scaled up when required.
- Adding processing capacity to the NAS gateway is an example of scaling.
- When the storage limit is reached, it can scale up, adding capacity on the SAN independently of the NAS head.
- Gateway NAS enables high utilization of storage capacity by sharing it with SAN environment.

Q.42 Explain the file sharing protocols supported by NAS devices.
Ans:

- Most NAS devices support multiple file service protocols to handle file I/O requests to a remote file system.
- NFS and CIFS are the common protocols for file sharing.

➢ **NFS**

- NFS is a client/server protocol for file sharing that is most commonly used on UNIX systems.
- NFS was originally based on the connectionless User Datagram Protocol (UDP).
- The NFS protocol provides a set of RPCs to access a remote file system for the following operations:
  - Searching files and directories
  - Opening, reading, writing to, and closing a file
  - Changing file attributes
  - Modifying file links and directories
- NFS uses the mount protocol to create a connection between the client and the remote system to transfer data.
- Currently, three versions of NFS are in use:

1. NFS version 2 (NFSv2)

- Uses UDP to provide a stateless network connection between a client and a server.

2. NFS version 3 (NFSv3)

- The most commonly used version, it uses UDP or TCP, and is based on the stateless protocol design.
- It includes some new features, such as a 64-bit file size, asynchronous writes, and additional file attributes to reduce re-fetching.

3. NFS version 4 (NFSv4)

- This version uses TCP and is based on a stateful protocol design. It offers enhanced security.

➢ **CIFS**

- CIFS is a client/server application protocol that enables client programs to make requests for files and services on remote computers over TCP/IP.
- The CIFS protocol enables remote clients to gain access to files that are on a server.
- CIFS enables file sharing with other clients by using special locks.
- File names in CIFS are encoded using unicode characters.
- CIFS provides the following features to ensure data integrity:
  - It uses file and record locking to prevent users from overwriting the work of another user on a file or a record.
  - It runs over TCP.
  - It supports fault tolerance and can automatically restore connections and reopen files that were open prior to interruption.

Q.43 Enumerate the steps required to host files and permit users to access hosted files on NAS devices.

Ans: Following are the steps required to host files and permit users to access the hosted files on a NAS device:

1. Create storage array volumes

- Create volumes on the storage array and assign Logical Unit Numbers (LUN) to the volumes. Present the newly created volumes to the NAS device.

2. Create NAS Volumes

- Perform a discovery operation on the NAS device, to recognize the new array-volumes and create NAS Volumes (logical volumes).
- Multiple volumes from the storage array may be combined to form large NAS volumes

3. Create NAS file systems

- Create NAS file systems on the NAS volumes.

4. Mount file systems

- Mount the created NAS file system on the NAS device.

5. Access the file systems

- Publish the mounted file systems on the network using NFS or CIFS for client access.

Q.44 What are the factors affecting the NAS performance and availability? Explain.

Ans:

- As NAS uses IP network, bandwidth and latency issues associated with IP affect NAS performance.
- Other factors that affect NAS performance at different levels are:

1. Number of hops

- A large number of hops can increase latency because IP processing is required at each hop, adding to the delay caused at the router.

2. Authentication with a directory service such as LDAP, Active Directory, or NIS

- The authentication service must be available on the network, with adequate bandwidth, and must have enough resources to accommodate the authentication load.

3. Retransmission

- Link errors, buffer overflows, and flow control mechanisms can result in retransmission. This causes packets that have not reached the specified destination to be resent.

4. Overutilized routers and switches

- The amount of time that an overutilized device in a network takes to respond is always more than the response time of an optimally utilized or underutilized device.

5. File/directory lookup and metadata requests

- NAS clients access files on NAS devices.
- The processing required before reaching the appropriate file or directory can cause delays.

- Sometimes a delay is caused by deep directory structures and can be resolved by flattening the directory structure.

6. Overutilized NAS devices

- Clients accessing multiple files can cause high utilization levels on a NAS device which can be determined by viewing utilization statistics.
- High utilization levels can be caused by a poor file system structure or insufficient resources in a storage subsystem.

7. Overutilized clients

- The client accessing CIFS or NFS data may also be overutilized. An overutilized client requires longer time to process the responses received from the server, increasing latency.

Q.45 What is iSCSI? What are its components? Explain the iSCSI host connectivity.

Ans:

- **iSCSI**
  - iSCSI is an IP-based protocol that establishes and manages connections between storage, hosts, and bridging devices over IP.
  - iSCSI carries block-level data over IP-based networks, including Ethernet networks and the Internet.
- **Components of iSCSI**
  - Host (initiators), targets, and an IP-based network are the principal iSCSI components.
  - The simplest iSCSI implementation does not require any FC components.
  - If an iSCSI-capable storage array is deployed, a host itself can act as an iSCSI initiator, and directly communicate with the storage over an IP network.
  - However, in complex implementations that use an existing FC array for iSCSI connectivity, iSCSI gateways or routers are used to connect the existing FC SAN.
  - These devices perform protocol translation from IP packets to FC packets and vice-versa.
- **iSCSI host connectivity**
  - iSCSI host connectivity requires a hardware component, such as a NIC with a software component (iSCSI initiator) or an iSCSI HBA.

- In order to use the iSCSI protocol, a software initiator or a translator must be installed to route the SCSI commands to the TCP/IP stack.
- A standard NIC, a TCP/IP offload engine (TOE) NIC card, and an iSCSI HBA are the three physical iSCSI connectivity options.
- If a standard NIC is used in heavy I/O load situations, the host CPU may become a bottleneck.
- A TOE NIC offloads the TCP management functions from the host and leaves iSCSI functionality to the host processor.
- The host passes the iSCSI information to the TOE card and the TOE card sends the information to the destination using TCP/IP.
- An iSCSI HBA is capable of providing performance benefits, as it offloads the entire iSCSI and TCP/IP protocol stack from the host processor.

Q.46 Discuss the different topologies for iSCSI connectivity.

Ans: The topologies used to implement iSCSI can be categorized into two classes: native and bridged.

➢ **Native iSCSI Connectivity**

- Native topologies do not have any FC components; they perform all communication over IP.
- The initiators may be either directly attached to targets or connected using standard IP routers and switches.

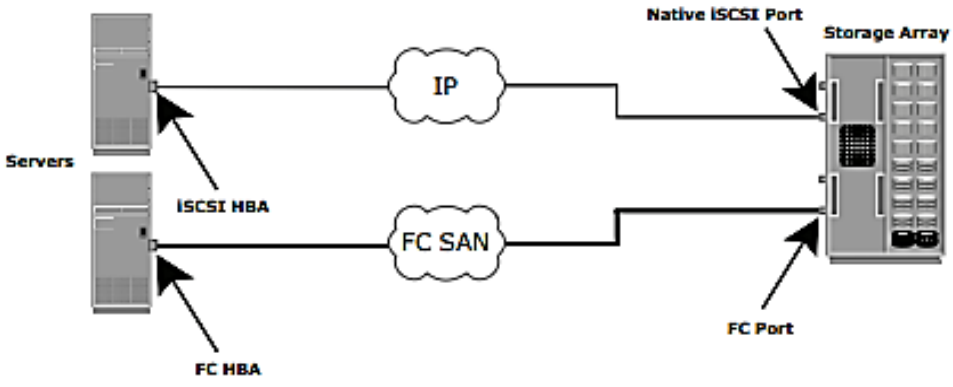

(a) Native iSCSI Connectivity

- As shown in Figure, the array has one or more Ethernet NICs that are connected to a standard Ethernet switch and configured with an IP address and listening port.

- Once a client/ initiator is configured with the appropriate target information, it connects to the array and requests a list of available LUNs.
- A single array port can service multiple hosts or initiators as long as the array can handle the amount of storage traffic that the hosts generate**.**

➢ **Bridged iSCSI Connectivity**

- A bridged iSCSI implementation includes FC components in its configuration.
- The array does not have any native iSCSI capabilities—that is, it does not have any Ethernet ports.
- Therefore, an external device, called a bridge, router, gateway, or a multi-protocol router, must be used to bridge the communication from the IP network to the FC SAN.
- These devices can be a stand-alone unit, or in many cases are integrated with an existing FC switch.
- In this configuration, the bridge device has Ethernet ports connected to the IP network, and FC ports connected to the storage.
- The iSCSI initiator/host is configured with the bridge's IP address as its target destination.
- The bridge is also configured with an FC initiator or multiple initiators. These are **called virtual initiators** because there is no physical device, such as an HBA, to generate the initiator record.

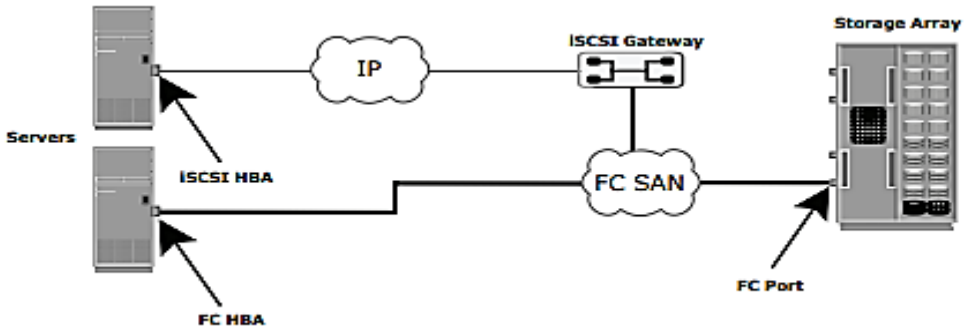

(b) Bridged iSCSI Connectivity

Q.47 Explain the iSCSI protocol stack

Ans:

- The architecture of iSCSI is based on the client/server model.

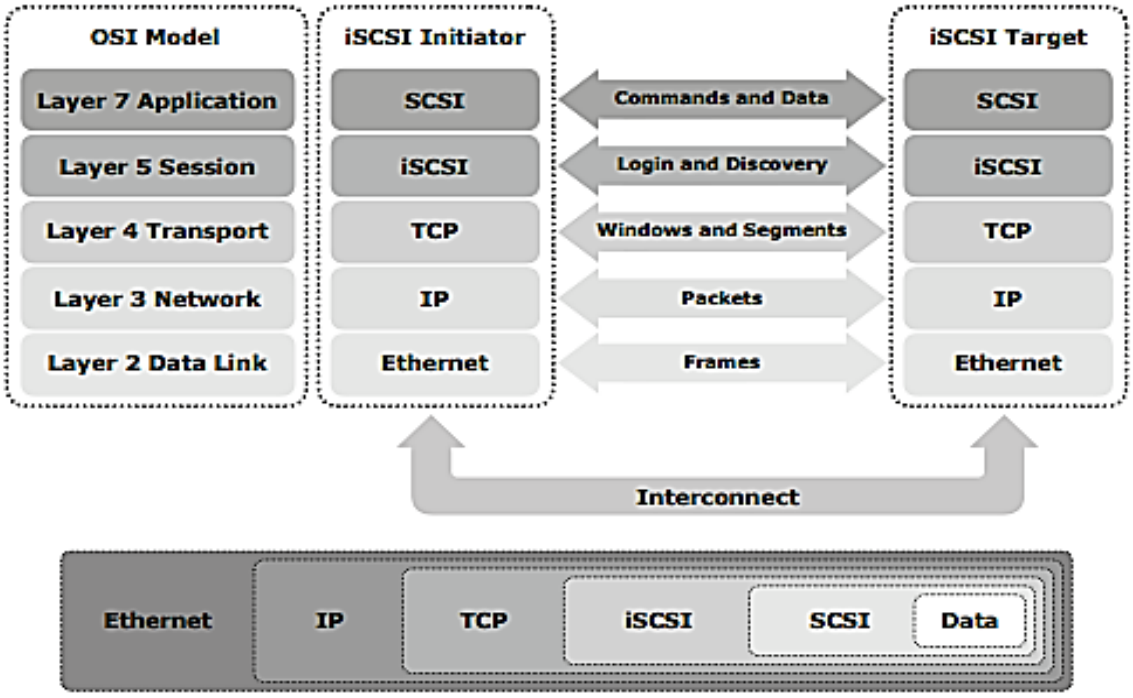

Figure: iSCSI protocol stack

- **SCSI** is the command protocol that works at the application layer of the OSI model. The initiators and targets use SCSI commands and responses to talk to each other.
- **iSCSI** is the session-layer protocol that initiates a reliable session between a device that recognizes SCSI commands and TCP/IP.
- The iSCSI session-layer interface is responsible for handling login, authentication, target discovery, and session management.
- **TCP** is used with iSCSI at the transport layer to provide reliable service.
- TCP is used to control message flow, windowing, error recovery, and retransmission.
- It relies upon the network layer of the OSI model to provide global addressing and connectivity.
- The layer-2 protocols at the data link layer of this model enable node-to-node communication for each hop through a separate physical network.

Q.48 What are the two ways in which an initiator discovers the location of the target network in iSCSI? Explain.

Ans:

- This discovery can take place in two ways: **SendTargets discovery** and **internet Storage Name Service (iSNS).**
- In SendTargets discovery, the initiator is manually configured with the target's network portal, which it uses to establish a discovery session with the iSCSI service on the target.
- The initiator issues the SendTargets command, and the target responds with the names and addresses of the targets available to the host.
- iSNS enables the automatic discovery of iSCSI devices on an IP network.
- The initiators and targets can be configured to automatically register themselves with the iSNS server.
- Whenever an initiator wants to know the targets that it can access, it can query the iSNS server for a list of available targets.

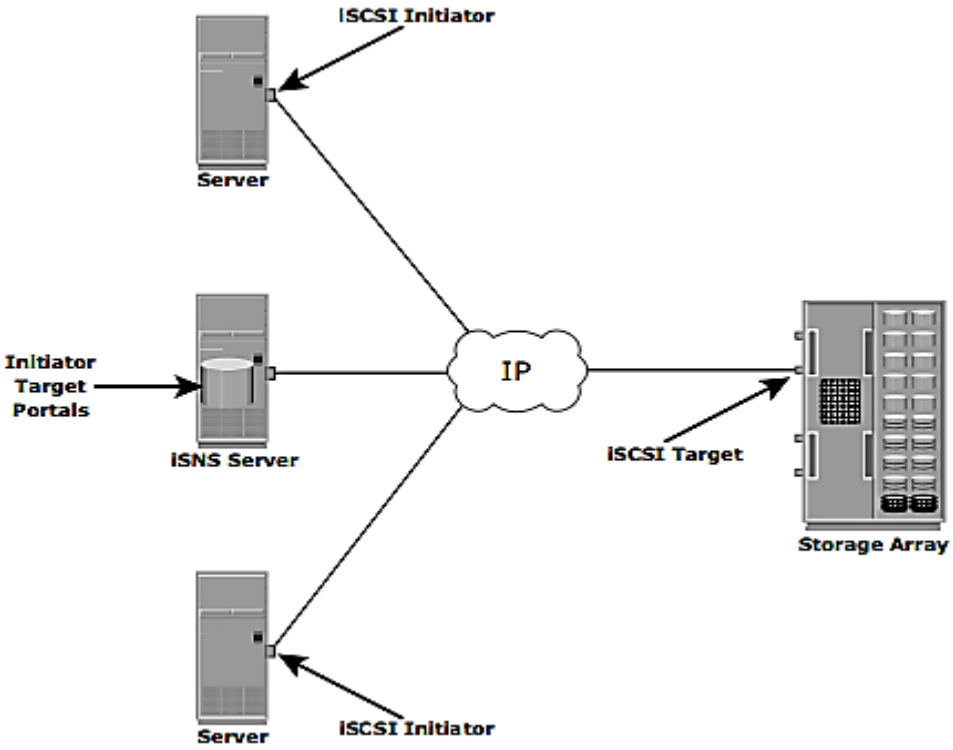

Figure: Discovery using iSNS

Q.49 What is iSCSI name? What are the two types of iSCSI names? Explain
Ans:

- A unique worldwide iSCSI identifier, known as an iSCSI name, is used to name the initiators and targets within an iSCSI network to facilitate communication.

- The unique identifier can be a combination of department, application, manufacturer name, serial number, asset number, or any tag that can be used to recognize and manage a storage resource.
- There are two types of iSCSI names:

➢ **iSCSI Qualified Name (IQN)**

- An organization must own a registered domain name in order to generate iSCSI Qualified Names.
- This domain name does not have to be active or resolve to an address.
- It just needs to be reserved to prevent other organizations from using the same domain name to generate iSCSI names.
- A date is included in the name to avoid potential conflicts caused by transfer of domain names.
- An example of an IQN is: **iqn.2008-02.com.example:optional_string**
- The optional_string provides a serial number, an asset number, or any of the storage device identifiers.

➢ **Extended Unique Identifier (EUI)**

- An EUI is a globally unique identifier based on the IEEE EUI-64 naming standard.
- An EUI comprises the eui prefix followed by a 16-character hexadecimal name, such as eui.0300732A32598D26.
- The 16-character part of the name includes 24 bits for the company name assigned by IEEE and 40 bits for a unique ID, such as a serial number.

Q.50 What is an iSCSI session? Why is it required?

Ans:

- An iSCSI session is established between an initiator and a target.
- A session ID (SSID), which includes an initiator ID (ISID) and a target ID (TSID), identifies a session.
- The session can be intended for one of the following:
  - Discovery of available targets to the initiator and the location of a specific target on a network.
  - Normal operation of iSCSI (transferring data between initiators and targets).

- TCP connections may be added and removed within a session. Each iSCSI connection within the session has a unique connection ID (CID).

Q.51 Explain the iSCSI protocol data units and its components in detail.
Ans:

- iSCSI initiators and targets communicate using iSCSI Protocol Data Units (PDUs).
- All iSCSI PDUs contain one or more header segments followed by zero or more data segments.
- The PDU is then encapsulated into an IP packet to facilitate the transport.
- A PDU includes the components shown in following figure:

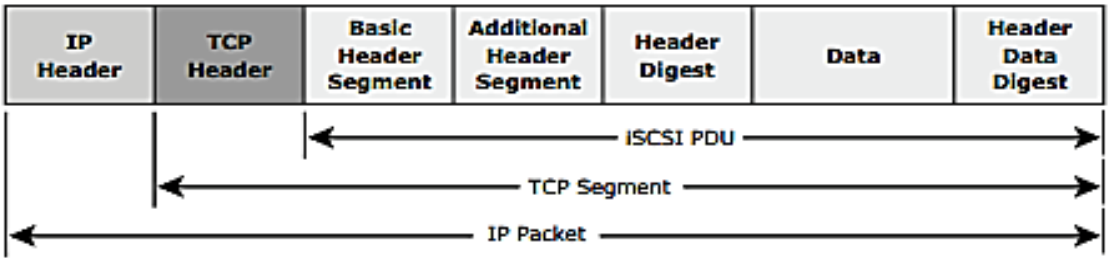

Figure: iSCSI PDU encapsulated in an IP packet

- The IP header provides packet-routing information that is used to move the packet across a network.
- The TCP header contains the information needed to guarantee the packet's delivery to the target.
- The iSCSI header describes how to extract SCSI commands and data for the target.
- iSCSI adds an optional CRC, known as the digest, beyond the TCP checksum and Ethernet CRC to ensure datagram integrity.
- The header and the data digests are optionally used in the PDU to validate integrity, data placement, and correct operation.

Q.52 Discuss ordering and numbering with reference to iSCSI.
Ans:

- iSCSI communication between initiators and targets is based on the requestresponse command sequences.
- A command sequence may generate multiple PDUs.

- A command sequence number (CmdSN) within an iSCSI session is used to number all initiator-to-target command PDUs belonging to the session.
- This number is used to ensure that every command is delivered in the same order in which it is transmitted, regardless of the TCP connection that carries the command in the session.
- Similar to command numbering, a status sequence number (StatSN) is used to sequentially number status responses.
- These unique numbers are established at the level of the TCP connection.

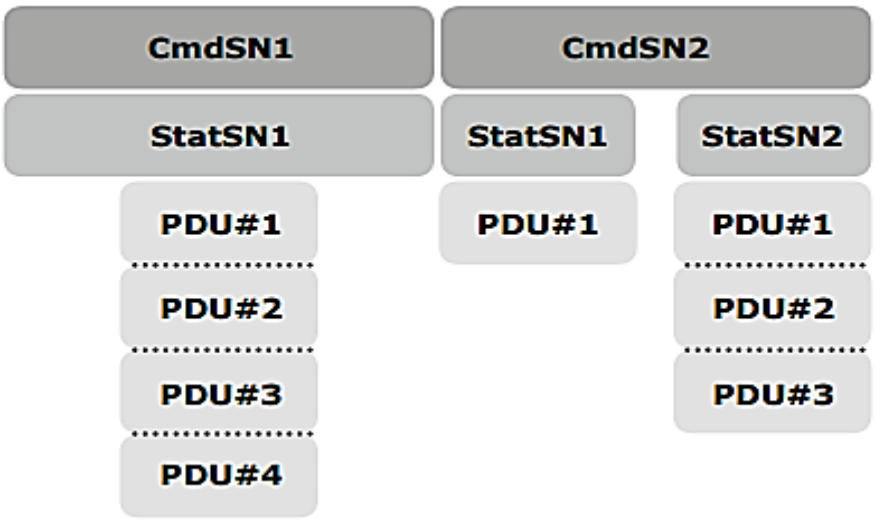

Figure: Command and status sequence number

- A target sends the request-to-transfer (R2T) PDUs to the initiator when it is ready to accept data.
- Data sequence number (DataSN) is used to ensure in-order delivery of data within the same command.
- The DataSN and R2T sequence numbers are used to sequence data PDUs and R2Ts, respectively.
- Each of these sequence numbers is stored locally as an unsigned 32-bit integer counter defined by iSCSI.
- These numbers are communicated between the initiator and target in the appropriate iSCSI PDU fields during command, status, and data exchanges.

Q.53 How is iSCSI error detection and recovery classified in iSCSI? Explain.

Ans:

- The iSCSI protocol addresses errors in IP data delivery.
- The error detection and recovery in iSCSI can be classified into three levels: Level 0 = Session Recovery, Level 1 = Digest Failure Recovery and Level 2 = Connection Recovery.

➢ **Level 0**
  - If an iSCSI session is damaged, all TCP connections need to be closed and all tasks and unfulfilled SCSI commands should be completed.
  - Then, the session should be restarted via the repeated login.

➢ **Level 1**
  - Each node should be able to selectively recover a lost or damaged PDU within a session for recovery of data transfer.
  - At this level, identification of an error and data recovery at the SCSI task level is performed, and an attempt to repeat the transfer of a lost or damaged PDU is made.

➢ **Level 2**
  - New TCP connections are opened to replace a failed connection. The new connection picks up where the old one failed.

- iSCSI may be exposed to the security vulnerabilities of an unprotected IP network.
- Some of the security methods that can be used are IPSec and authentication solutions such as Kerberos and CHAP (challenge-handshake authentication protocol).

Q.54 What is FCIP? Explain the encapsulation of FC frames into IP payload.
Ans:

- FCIP is a tunneling protocol that enables distributed FC SAN islands to be transparently interconnected over existing IP-based local, metropolitan, and wide-area networks.
- As a result, organizations now have a better way to protect, store, and move their data while leveraging investments in existing technology.
- FCIP uses TCP/IP as its underlying protocol.
- In FCIP, the FC frames are encapsulated onto the IP payload.

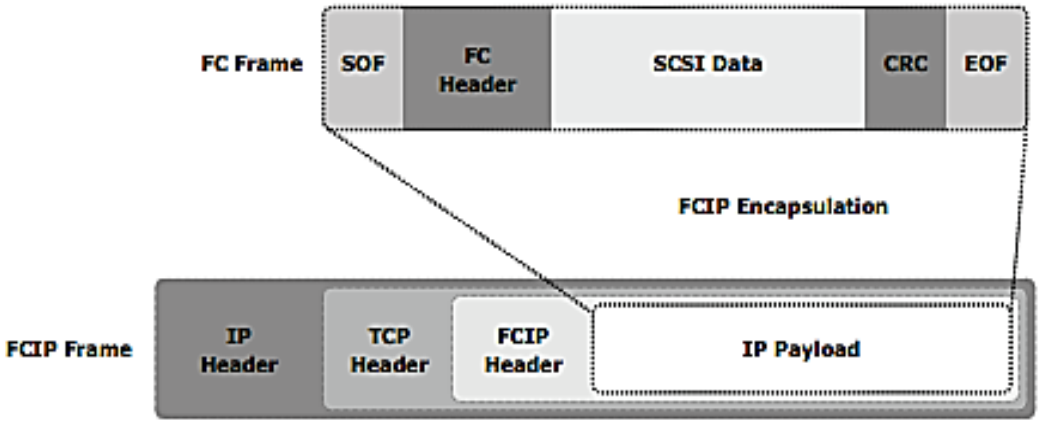

Figure: FCIP encapsulation

- When SAN islands are connected using FCIP, each interconnection is called an FCIP link.
- A successful FCIP link between two SAN islands results in a fully merged FC fabric.

Q.55 With the help of a diagram, explain the FCIP topology.

Ans:

- An FCIP gateway router is connected to each fabric via a standard FC connection.
- The fabric treats these routers like layer 2 fabric switches.
- The other port on the router is connected to an IP network and an IP address is assigned to that port.
- This is similar to the method of assigning an IP address to an iSCSI port on a gateway.
- Once IP connectivity is established, the two independent fabrics are merged into a single fabric.
- When merging the two fabrics, all the switches and routers must have unique domain IDs, and the fabrics must contain unique zone set names.
- Failure to ensure these requirements will result in a segmented fabric.
- The FC addresses on each side of the link are exposed to the other side, and zoning or masking can be done to any entity in the new environment.

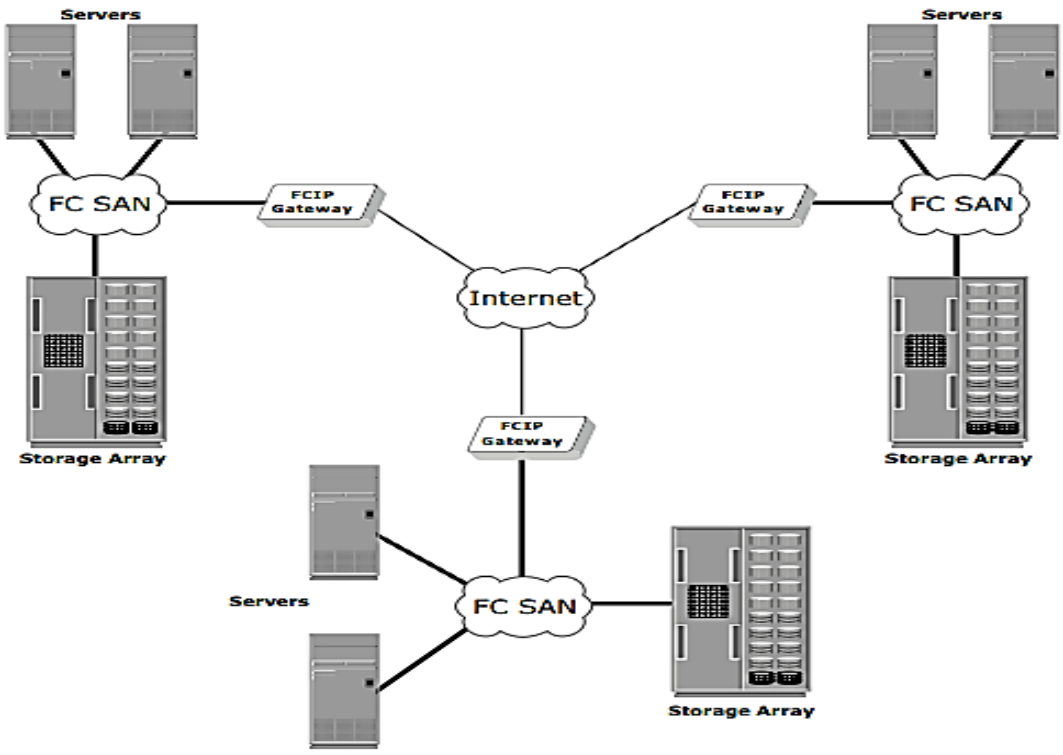

Figure: FCIP topology

Q.56 Explain FCIP performance and security.

Ans:

- From the perspective of performance, multiple paths to multiple FCIP gateways from different switches in the layer 2 fabric eliminates single points of failure and provides increased bandwidth.
- In addition, because FCIP creates a unified fabric, disruption in the underlying IP network can cause instabilities in the SAN environment.
- These include a segmented fabric, excessive RSCNs, and host timeouts.
- The vendors of FC switches have recognized some of the drawbacks related to FCIP and have implemented features to provide additional stability, such as the capability to segregate FCIP traffic into a separate virtual fabric.
- Security is also a consideration in an FCIP solution because the data is transmitted over public IP channels.
- Various security options are available to protect the data based on the router's support.
- IPSec is one such security measure that can be implemented in the FCIP environment.