CMR
INSTITUTE OF
TECHNOLOGY

USN | | | | | | | | | | |

Internal Assessment Test 2 – Nov. 2020

Sub:	Storage Area Networks							Sub Code:	18MCA554
Date:	06/11/2020	Duration:	90 min's	Max Marks:	50	Sem & Sec:	5 A & B	Branch:	MCA

Note : Answer any full FIVE Questions

	PART I	MARKS	OBE CO	OBE RBT
1	What is SAN? Explain the components of FC SAN (Fiber Channel Storage Area Networks) with a neat diagram?	[10]	CO2	L2
2	Explain the SCSI architecture and Parallel SCSI Addressing with a neat diagram	[10]	CO2	L2
3	What is Fabric? Explain SAN connectivity options with a neat diagram	[10]	CO2	L1
4	a) What is NAS? Explain its benefits? b) Explain the components of NAS with a neat diagram?	[4] [6]	CO2 CO3	L2
5	Explain the layers of FCP stack with a neat diagram. Also describe FCP addressing with a neat diagram	[10]	CO2	L2
6	Explain NAS I/O operation with a neat diagram	[10]	CO2	L2
7	What is zoning? Explain different types of zoning with a neat diagram	[10]	CO2	L3
8	Write a note on NAS file sharing protocol	[10]	CO3	L2

CMR
INSTITUTE OF
TECHNOLOGY

USN | | | | | | | | | | |

Internal Assessment Test 2 – Nov. 2020

Sub:	Storage Area Networks							Sub Code:	18MCA554
Date:	06/11/2020	Duration:	90 min's	Max Marks:	50	Sem & Sec:	5 A & B	Branch:	MCA

Note : Answer any full FIVE Questions

Q1	What is SAN? Explain the components of FC SAN (Fiber Channel Storage Area Networks) with a neat diagram?	[10]	CO1	L2

Sol:

Components of FC SAN

FC SAN is a network of servers and shared storage devices. Servers and storage are the endpoints or devices in the SAN (called nodes). **FC SAN infrastructure consists of node ports, cables, connectors, and interconnecting devices (such as FC switches or hubs), along with SAN management software**

Node Ports

In a Fibre Channel network, the end devices, such as hosts, storage arrays, and tape libraries, are all referred to as nodes. Each node is a source or destination of information. Each node requires one or more ports to provide a physical interface for communicating with other nodes. These ports are integral components of host adapters, such as HBA, and storage front-end controllers or adapters. In an FC environment a port operates in full-duplex data transmission mode with a transmit (Tx) link and a receive (Rx) link

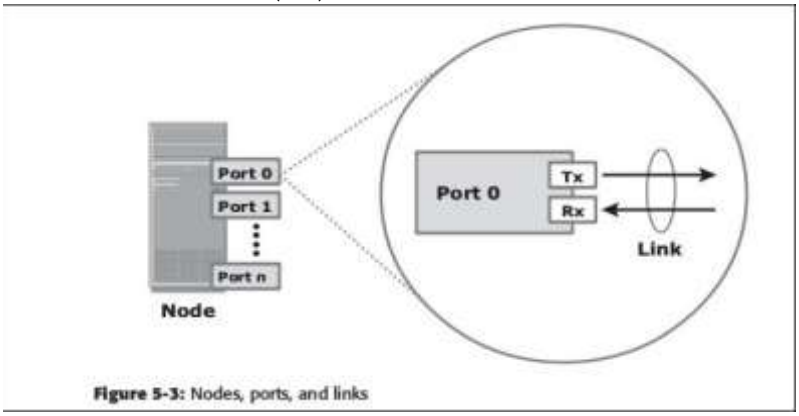

Figure 5-3: Nodes, ports, and links

Cables and Connectors

SAN implementations use optical fiber cabling. Copper can be used for shorter distances for back-end connectivity because it provides an acceptable 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: multimode and single-mode. Multimode fiber (MMF) cable carries multiple beams of light projected at different angles simultaneously onto the core of the cable (see Figure 5-4 [a]). Based on the bandwidth, multimode fibers are classified as OM1 (62.5µm core), OM2 (50µm core), and laser-optimized OM3 (50µm core). In an MMF transmission, multiple light beams traveling inside the

cable tend to disperse and collide. This collision weakens the signal strength after it travels a certain distance — a process known as modal dispersion. An MMF cable is typically used for short distances because of signal degradation (attenuation) due to modal dispersion. Single-mode fiber (SMF) carries a single ray of light projected at the center of the core (see Figure 5-4 [b]). These cables are available in core diameters of 7 to 11 microns; the most common size is 9 microns. In an SMF transmission, a single light beam travels in a straight line through the core of the fiber. The small core and the single light wave help to limit modal dispersion. Among all types of fiber cables, single-mode provides minimum signal attenuation over maximum distance (up to 10 km). A single-mode cable is used for long-distance cable runs, and distance usually depends on the power of the laser at the transmitter and sensitivity of the receiver.

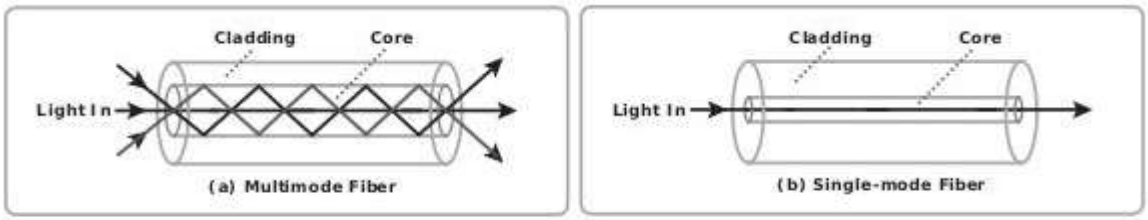

Figure 5-4: Multimode fiber and single-mode fiber

MMFs are generally used within data centers for shorter distance runs, whereas SMFs are used for longer distances. A connector is attached at the end of a cable to enable swift connection and disconnection of the cable to and from a port. A Standard connector (SC) (see Figure 5-5 [a]) and a Lucent connector (LC) (see Figure 5-5 [b]) are two commonly used connectors for fiber optic cables. Straight Tip (ST) is another fiber-optic connector, which is often used with fiber patch panels (see Figure 5.5 [c]).

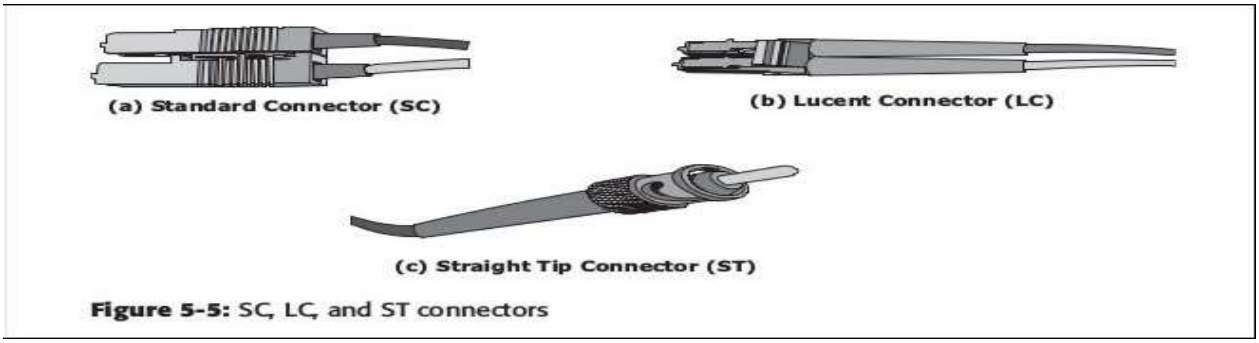

Figure 5-5: SC, LC, and ST connectors

Interconnect Devices

FC hubs, switches, and directors are the interconnect devices commonly used in FC 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 loop because

data travels through all the connection points. Because of the availability of low-cost and high-performance switches, hubs are no longer used in FC SANs.

Switches are more intelligent than hubs and directly route data from one physical port to another. Therefore, nodes do not share the bandwidth. Instead, each node has a dedicated communication path.

Directors are high-end switches with a higher port count and better fault-tolerance capabilities. Switches are available with a fixed port count or with modular design. In a modular switch, the port count is increased by installing additional port cards to open slots. The architecture of a director is always modular, and its port count is increased by inserting additional line cards or blades to the director's chassis. High-end switches and directors contain redundant components to provide high availability. Both switches and directors have management ports (Ethernet or serial) for connectivity to SAN management servers.

A port card or blade has multiple ports for connecting nodes and other FC switches. Typically, a Fibre Channel transceiver is installed at each port slot that houses the transmit (Tx) and receive (Rx) link. In a transceiver, Tx and Rx links share common circuitry. Transceivers inside a port card are connected to an application specific integrated circuit, also called port ASIC. Blades in a director usually have more than one ASIC for higher throughput.

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 zoning.

2	Explain the SCSI architecture and Parallel SCSI Addressing with a neat diagram	[10]	CO2	L2

SCSI-3 Architecture The SCSI-3 architecture defines and categorizes various SCSI-3 standards and requirements for SCSI-3 implementations. The SCSI-3 architecture was approved and published as the standard X.3.270-1996 by the ANSI. This architecture helps developers, hardware designers, and users to understand and effectively utilize SCSI. The three major components of a SCSI architectural model are as follows: **1. SCSI-3 command protocol:** This consists of primary commands that are common to all devices as well as device-specific commands that are unique to a given class of devices. **2. Transport layer protocols:** These are a standard set of rules by which devices communicate and share information. **3. Physical layer interconnects:** These are interface details such as electrical signaling methods and data transfer modes. ***Common access methods*** are the ANSI software interfaces for SCSI devices.

Figure 5-3 shows the SCSI-3 standards architecture with interrelated groups of other standards within SCSI-3.

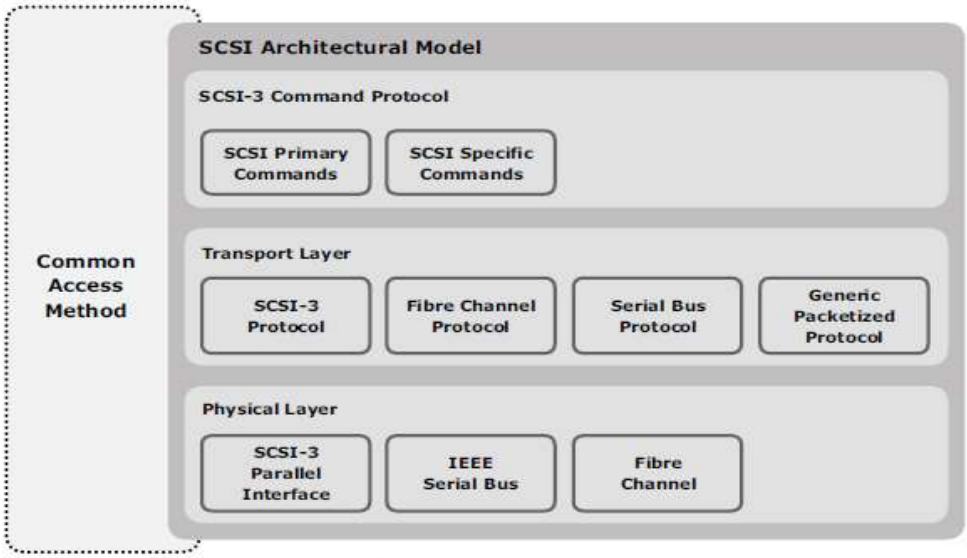

Figure 5-3: SCSI-3 standards architecture

SCSI-3 Client-Server Model SCSI-3 architecture derives its base from the client-server relationship, in which a client directs a service request to a server, which then fulfills the client's request. In a SCSI environment, an initiator-target concept represents the client server model. In a SCSI-3 client-server model, a particular SCSI device acts as a SCSI target device, a SCSI initiator device, or a SCSI target/initiator device. Each device performs the following functions: **1. SCSI initiator device:** Issues a command to the SCSI target device, to perform a task. A SCSI host adaptor is an example of an initiator. **2. SCSI target device:** Executes commands to perform the task received from a SCSI initiator. Typically a SCSI peripheral device acts as a target device. However, in certain implementations, the host adaptor can also be a target device. Figure 5-4 displays the SCSI-3 client-server model, in which a SCSI initiator, or a client, sends a request to a SCSI target, or a server. The target performs the tasks requested and sends the output to the initiator, using the protocol service interface. A SCSI target device contains one or more logical units. A logical unit is an object that implements one of the device functional models as described in the SCSI command standards. A logical unit has two components, a *device server* and a *task manager*, as shown in Figure 5-4. The logical unit processes the commands sent by a SCSI initiator. The device server addresses client requests, and the task manager performs management functions. The SCSI initiator device is comprised of an application client and task management function, which initiates device service and task management requests.

Each device service request contains a *Command Descriptor Block (CDB)*. The CDB defines the command to be executed and lists command-specific inputs and other parameters specifying how to process the command. The SCSI devices are identified by a specific number called a SCSI ID. In narrow SCSI (bus width=8), the devices are numbered 0 through 7; in wide (bus width=16) SCSI, the devices are numbered 0 through 15. These ID numbers set the device priorities on the SCSI bus. In narrow SCSI, 7 has the highest priority and 0 has the lowest priority. In wide SCSI, the device IDs from 8 to 15 have the highest priority, but the entire sequence of wide SCSI IDs has lower priority than narrow SCSI IDs. Therefore, the overall priority sequence for a wide SCSI is 7, 6, 5, 4, 3, 2, 1, 0, 15, 14, 13, 12, 11, 10, 9, and 8.

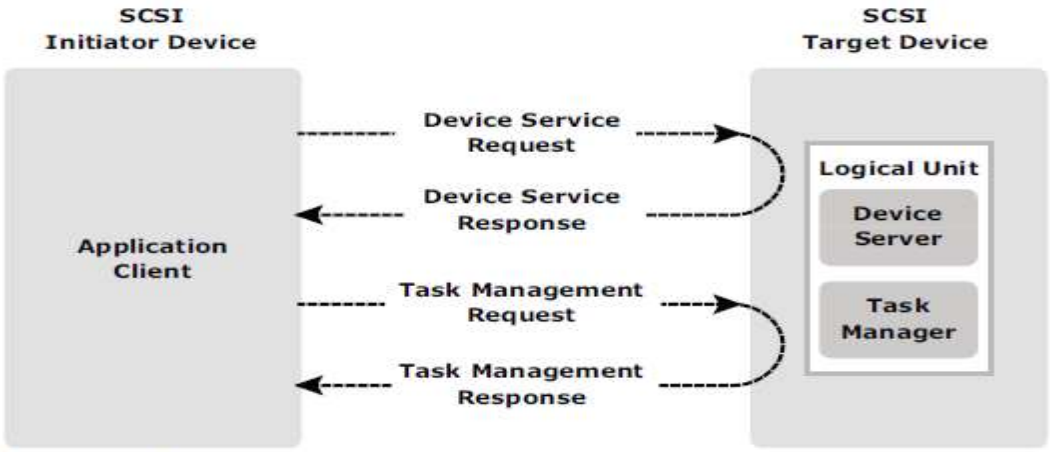

Figure 5-4: SCSI-3 client-server model

SCSI Communication Model A SCSI communication model (see Figure 5-6) is comprised of three interconnecting layers as defined in the
SAM-3 and is similar to the OSI seven-layer model. Lower-level layers render their services to the upper-level layers. A highl evel layer communicates with a low-level layer by invoking the services that the low-level layer provides. The protocol at each layer defines the communication between peer layer entities.

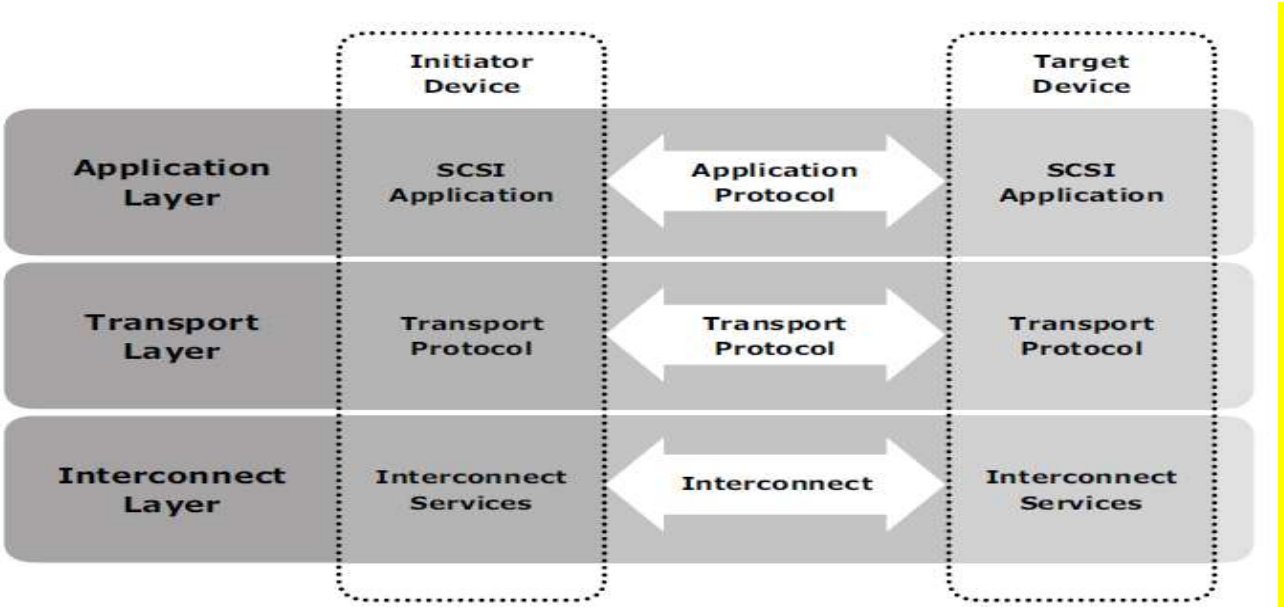

Figure 5-6: SCSI communication model

There are three layers in the SCSI communication model: 1. **SCSI application layer (SAL):** This layer contains both client and server applications that initiate and process SCSI I/O operations using a SCSI application protocol. 2. **SCSI transport protocol layer (STPL):** This layer contains the services and protocols that allow communication between an initiator and targets. **3. Interconnect layer:** This layer facilitates data transfer between the initiator and targets. The interconnect layer is also known as the *service delivery subsystem* and comprises the services, signalling mechanisms, and interconnects for data transfer.

Parallel SCSI Addressing In the Parallel SCSI Initiator-Target communication (see Figure 5-7), an initiator ID uniquely identifies the
initiator and is used as an originating address. This ID is in the range of 0 to 15, with the range 0 to 7 being the most common. A target ID uniquely identifies a target and is used as the address for

exchanging commands and status information with initiators. The target ID is in the range of 0 to 15.

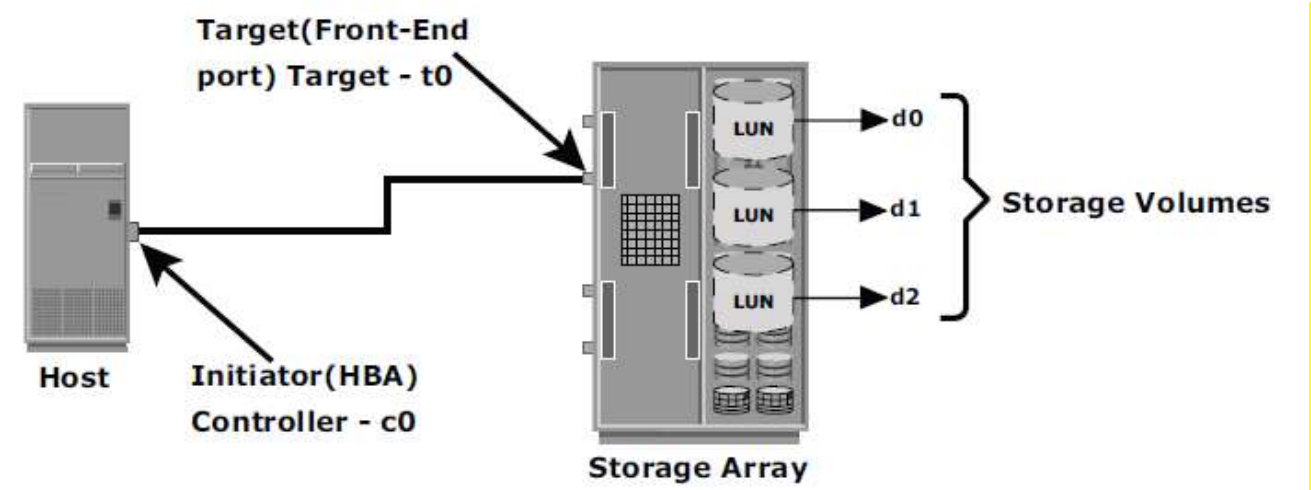

Host Addressing :
Storage Volume 1 - c0 t0 d0
Storage Volume 2 - c0 t0 d1
Storage Volume 3 - c0 t0 d2

Figure 5-7: SCSI Initiator-Target communication

SCSI addressing is used to identify hosts and devices. In this addressing, the UNIX naming convention is used to identify a disk and the three identifiers — initiator ID, target ID, and a LUN — in the **c**n|**t**n|**d**n format, which is also referred as *ctd addressing*. Here, cn is the initiator ID, commonly referred to as the controller ID; tn is the target ID of the device, such as t0, t1, t2, and so on; and dn is the device number reflecting the actual address of the device unit, such as d0, d1, and d2. A LUN identifies a specific logical unit in a target. The implementation of SCSI addressing may differ from one vendor to another. Figure 5-7 shows ctd addressing in the SCSI architecture.

3	What is Fabric? Explain SAN connectivity options with a neat diagram	[10]	CO2	L1

Sol :
FC Connectivity
The FC architecture supports three basic interconnectivity options:

i) Point-to-Point
ii) Arbitrated loop (FC-AL)
iii) Fibre Channel Switched Fabric

i) 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. Moreover, it cannot be scaled to accommodate a large number of network devices. Standard DAS uses point- to- point connectivity.

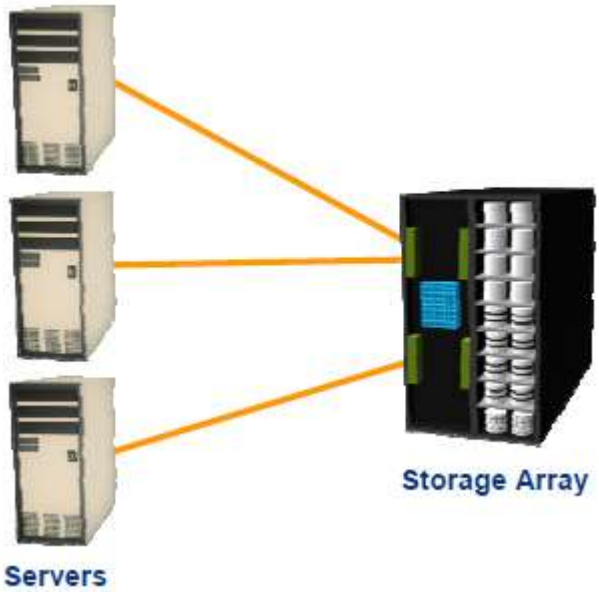

ii) Fibre Channel Arbitrated Loop

In the FC-AL configuration, devices are attached to a shared loop, as shown in Figure 6-7. 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. As a loop configuration, FC-AL can be implemented without any interconnecting devices by directly connecting one device to another in a ring through cables.

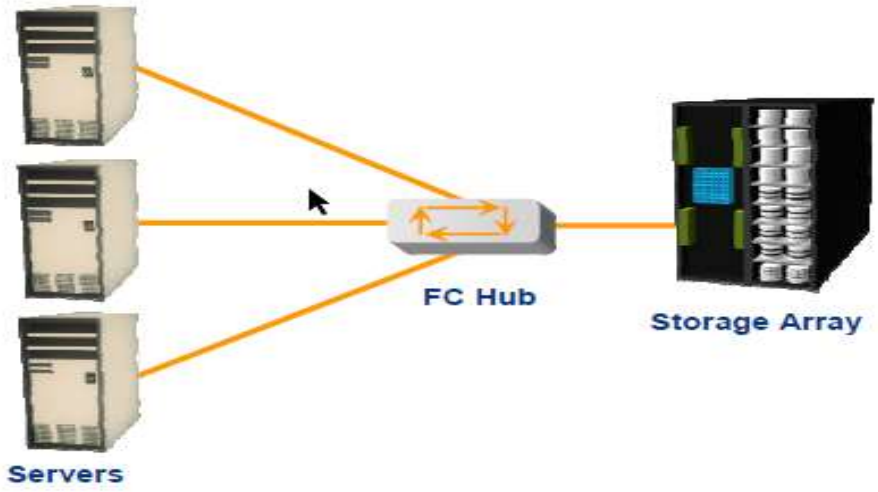

The FC-AL configuration has the following limitations in terms of scalability: 1. FC-AL shares the bandwidth in the loop. Only one device can perform I/O operations at a time. Because each device in a loop has to wait for its turn to process I/O request, the speed of data transmission is low in an FC-AL topology.

2. FC-AL uses 8-bit addressing. It can support up to 127 devices on a loop. 3. Adding or removing a device results in loop re-initialization, which can cause a momentary pause in loop traffic. ***FC-AL Transmission*** When a node in the FC-AL topology attempts to transmit data, the node sends an *arbitration (ARB)* frame to

each node on the loop. If two nodes simultaneously attempt to gain control of the loop, the node with the highest priority is allowed to communicate with another node. This priority is determined on the basis of Arbitrated Loop Physical Address (AL-PA) and Loop ID. When the initiator node

receives the ARB request it sent, it gains control of the loop. The initiator then transmits data to the node with which it has established a virtual connection. Figure 6-8 illustrates the process of data transmission in an FC-AL configuration.

FC-AL Data Transmission

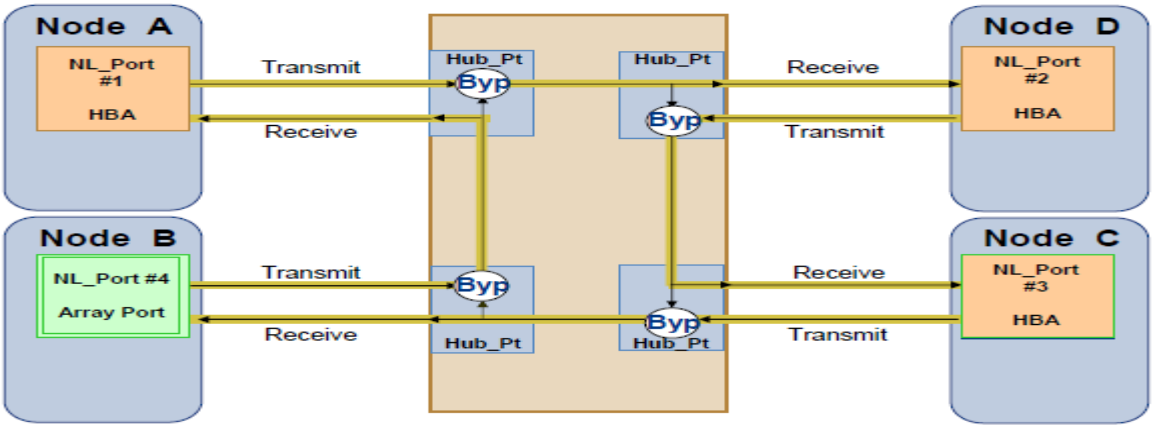

Arbitration is done in FC-AL. Consider Node A want to communicate with Node B, the steps are as follows:

1. High priority initiator, Node A inserts the ARB frame in the loop.
2. ARB frame is passed to the next node (Node D) in the loop.
3. Node D receives high priority ARB, therefore remains idle.
4. ARB is forwarded to next node (Node C) in the loop.
5. Node C receives high priority ARB, therefore remains idle.
6. ARB is forwarded to next node (Node B) in the loop.
7. Node B receives high priority ARB, therefore remains idle and
8. ARB is forwarded to next node (Node A) in the loop.
9. Node A receives ARB back; now it gains control of the loop and can start communicating with target Node B.

iii) **Fibre Channel Switched Fabric**

Unlike a loop configuration, a Fibre Channel switched fabric (FC-SW) network provides interconnected devices, dedicated bandwidth, and scalability. The addition or removal of a device in a switched fabric is minimally disruptive; it does not affect the ongoing traffic between other devices. 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. Each port in a fabric has a unique 24-bit fibre channel address for communication.

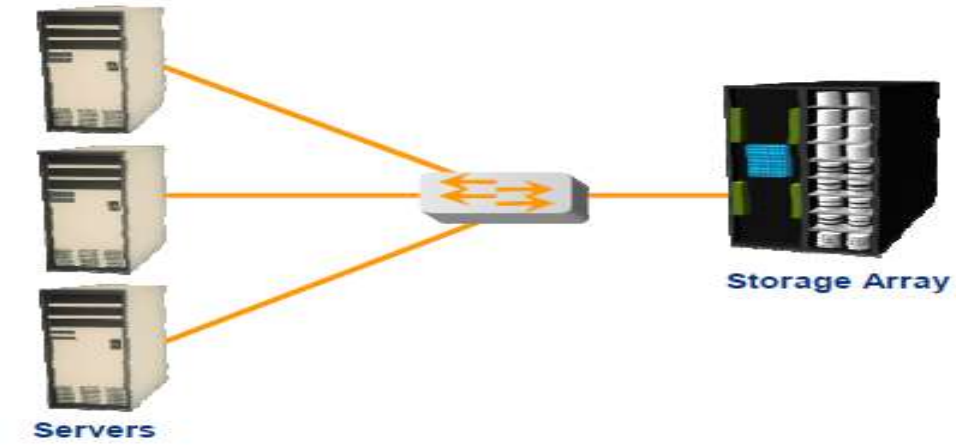

When the number of tiers in a fabric increases, the distance that a fabric management message must travel to reach each switch in the fabric also increases. The increase in the distance also increases the time taken to propagate and complete a fabric reconfiguration event, such as the addition of a new switch, or a zone set propagation event (detailed later in this chapter). Figure 6-10 illustrates two-tier and three-tier fabric architecture.

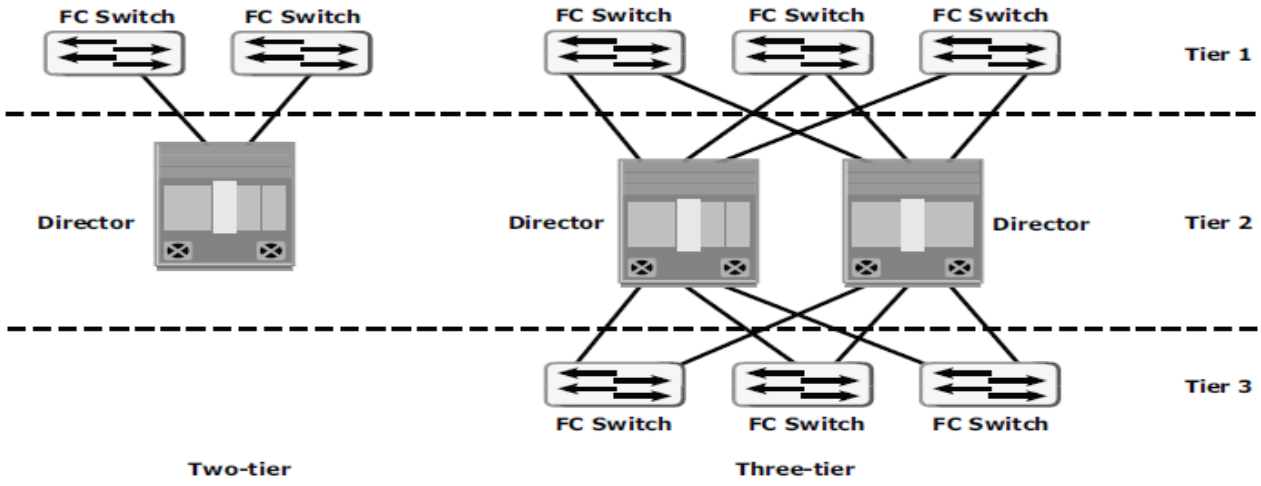

Figure 6-10: Tiered structure of FC-SW topology

FC-SW Transmission FC-SW uses switches that are intelligent devices. They can switch data traffic from an initiator node to a target
node directly through switch ports. Frames are routed between source and destination by the fabric. As shown in Figure 6-11, if node B wants to communicate with node D, Nodes should individually login first and then transmit data via the FC-SW. This link is considered a dedicated connection between the initiator and the target.

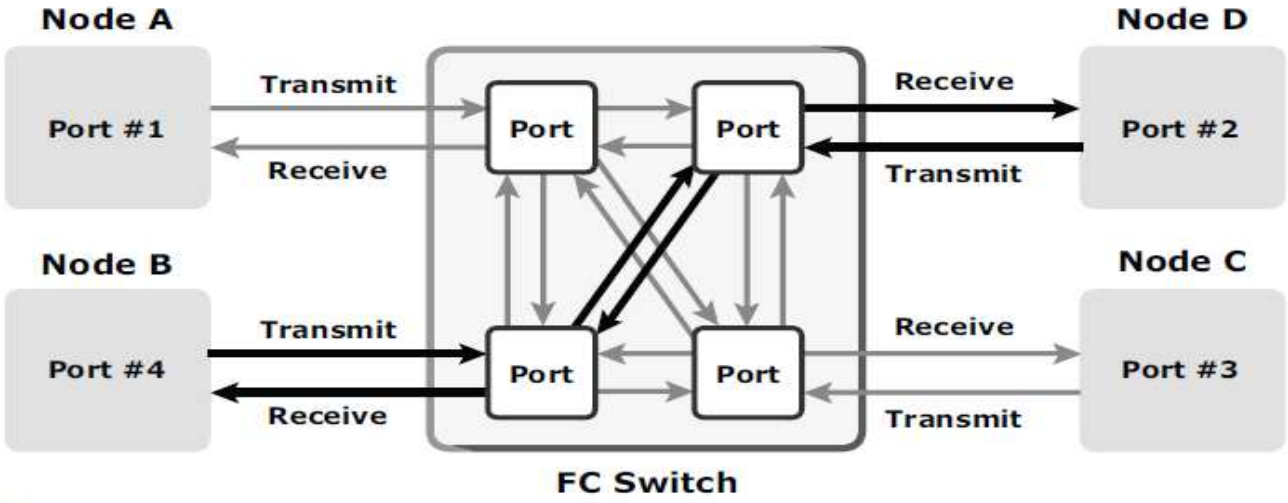

Figure 6-11: Data tansmission in FC-SW topology

4	a) What is NAS? Explain its benefits? b) Explain the components of NAS with a neat diagram?	[4] [6]	CO2 CO3	L2

a) N*etwork-attached storage (NAS)* is an IP-based file-sharing device attached to a local area network. A NAS device is a dedicated, high-performance, high-speed, single-purpose file serving and storage
system.

i) NAS provides the advantages of server consolidation by eliminating the need for multiple file servers.

ii) It provides storage consolidation through file-level data access and sharing.

iii) NAS helps to eliminate bottlenecks that users face when accessing files from a general purpose server.

iv) 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.

v) NAS uses NFS for UNIX, CIFS for Windows, and File Transfer Protocol (FTP) and other protocols for both environments.

Benefits of NAS

NAS offers the following benefits:

■ **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.

■ **Improved efficiency:** NAS uses an operating system specialized for file serving. It improves the utilization of general-purpose servers by relieving them of file-server operations.

■ **Improved flexibility:** NAS is flexible and can serve requests from different types of clients from the same source.

■ **Centralized storage:** Centralizes data storage to minimize data duplication on client workstations, simplify data management, and ensures greater data protection.

■ **Simplified management:** Provides a centralized console that makes it possible to manage file systems efficiently.

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

■**High availability:** Offers efficient replication and recovery options, enabling high data availability. NAS uses redundant networking components that provide maximum connectivity options. A NAS device can use clustering technology for failover.

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

b) Components of NAS

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

The NAS environment includes clients accessing a NAS device over an IP network using standard protocols.

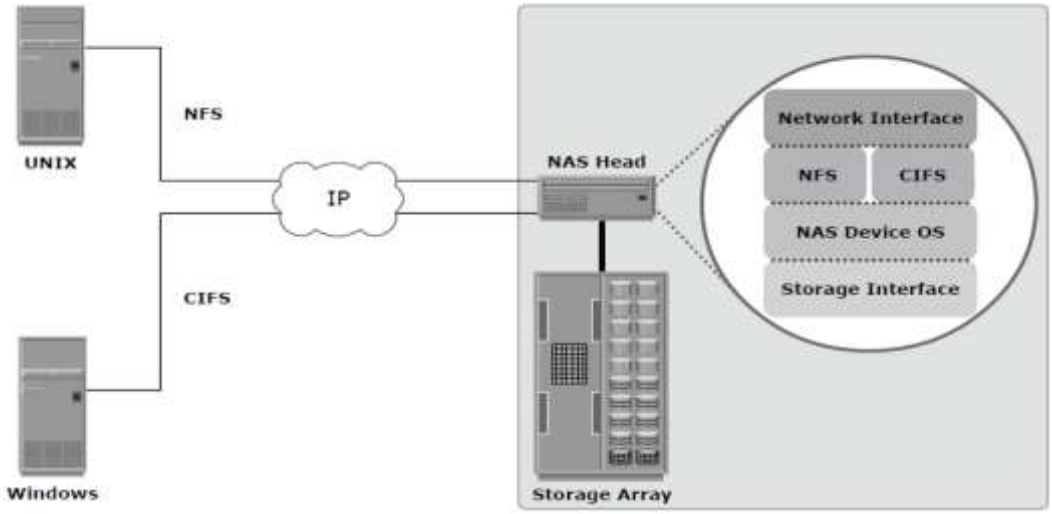

5	Explain the layers of FCP stack with a neat diagram. Also describe FCP addressing with a neat diagram	[10]	CO2	L2

Sol:

Fibre Channel Protocol Stack It is easier to understand a communication protocol by viewing it as a structure of independent layers. FCP defines the communication protocol in five layers: FC-0 through FC-4 (except FC-3 layer, which is not implemented). In a layered communication model, the peer layers on each node talk to each other through defined protocols. Figure 6-13 illustrates the fibre channel protocol stack.

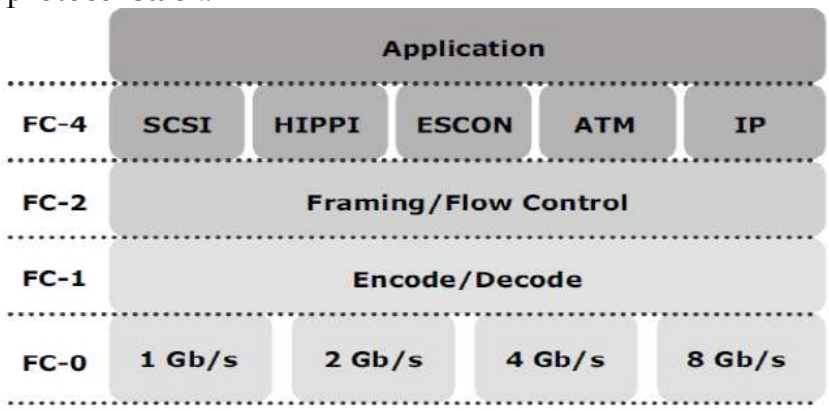

Figure 6-13: Fibre channel protocol stack

Fibre Channel Protocol Stack

FC layer	Function	SAN relevant features specified by FC layer
FC-4	Mapping interface	Mapping upper layer protocol (e.g. SCSI-3 to FC transport
FC-3	Common services	Not implemented
FC-2	Routing, flow control	Frame structure, ports, FC addressing, buffer credits
FC-1	Encode/decode	8b/10b encoding, bit and frame synchronization
FC-0	Physical layer	Media, cables, connector

i) 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 (see Figure 6-7). Some of the protocols include SCSI, HIPPI Framing Protocol, Enterprise Storage Connectivity (ESCON), ATM, and IP.

ii) ***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.

iii) ***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.

iv) ***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. The FC transmission can use both electrical and optical media.

6	Explain NAS I/O operation with a neat diagram	[10]	CO2	L2

Sol.

NAS I/O Operation

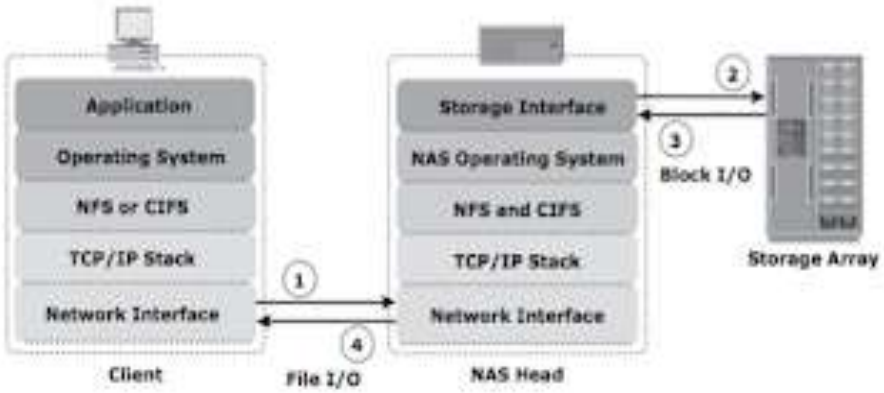

Figure 7-4: NAS I/O operation

TCP/IP I/O request -> Block-level IO request -> Get block data -> File IO data

1. The NAS device receives client **I/O request** from the **TCP/IP** network.
2. The NAS device converts the I/O request into an appropriate physical **storage request,** which is a **block-level I/O**, and then performs the operation on the physical storage.
3. When the **NAS device receives data from the storage**, it processes and repackages the data into an appropriate **file protocol response**.
4. The NAS device packages this response into **TCP/IP** again and forwards it to the client through the network.

4.5.1 Hosting and Accessing Files on NAS

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.

4.6 Factors Affecting NAS Performance

IP network bandwidth and latency issues(Network congestion is the most significant sources of latency).

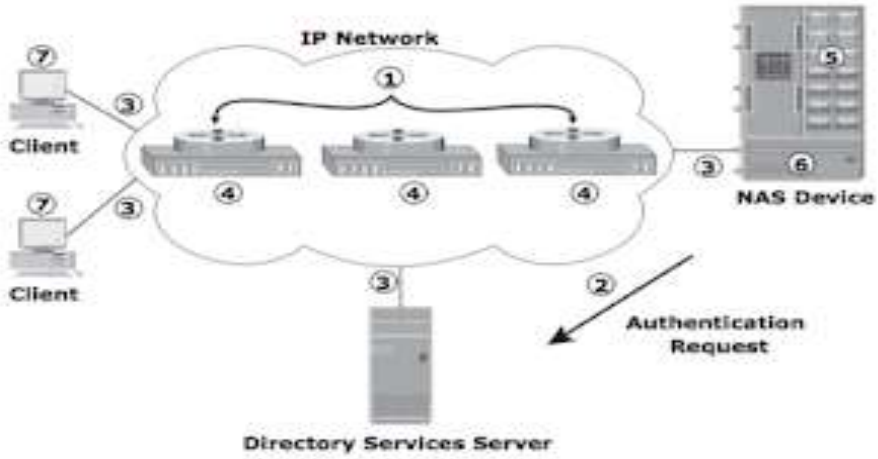

Figure 7-8: Causes of latency

1. **Number of hops (**router).
2. **Authentication with a directory service such as ActiveDirectory or NIS:** a large number of authentication requests can increase latency.
3. **Retransmission:** Link errors and buffer overflows can result in retransmission. This causes **packets that have not reached the specified destination** to be re-sent. Care must be taken to match both **speed and duplex settings** on the network devices and the NAS heads. **Improper configuration** might result in errors and retransmission, adding to latency.
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. Additional devices should be added if the current devices are overutilized.
5. **File system lookup and metadata requests:** NAS clients access files on NAS devices. The **processing required to reach 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. **Poor file system layout and an overutilized** disk system can also degrade performance.
6. **Over utilized NAS devices:** Clients accessing multiple files can cause **high utilization levels on a NAS device.** High memory, CPU, or disk subsystem 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 might also be over utilized. An overutilized client requires a longer time to process the requests and responses.

Q7	What is zoning? Explain different types of zoning with a neat diagram	[10]	CO2	L3

Sol :

Zoning is an FC switch function that enables nodes within the fabric to be logically segmented into groups that can communicate with each other. The zoning function controls this process by allowing only the members in the same zone to establish these link-level services.

When a device (host or storage array) logs onto a fabric, it is registered with the name server. When a port logs onto the fabric, it goes through a device discovery process with other devices registered in the name server.

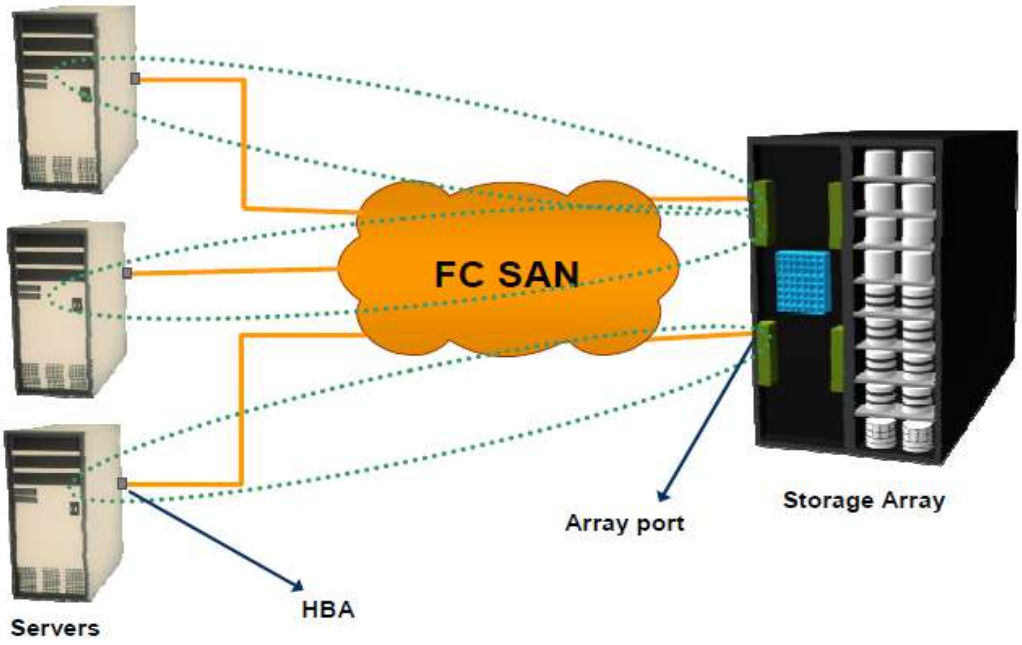

Types of Zoning

Zoning can be categorized into three types:

i) **Port zoning**: Uses the physical address of switch ports to define zones. 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 (switch domain ID and port number) to which HBA and its targets (storage devices) are connected. If a node is moved to another switch port in the fabric, then zoning must be modified to allow the node, in its new port, to participate in its original zone. However, if an HBA or storage device port fails, an administrator just has to replace the failed device without changing the zoning configuration.

ii) **WWN zoning:** Uses World Wide Names to define zones. The zone members are the unique WWN addresses of the HBA and its targets (storage devices). A major advantage of WWN zoning is its flexibility. WWN zoning allows nodes to be moved to another switch port in the fabric and maintain connectivity to its zone partners without having to modify the zone configuration. This is possible because the WWN is static to the node port.

iii) **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.

Zoning is used with LUN masking to control server access to storage. However, these are two different activities. Zoning takes place at the fabric level and LUN masking is performed at the array level.

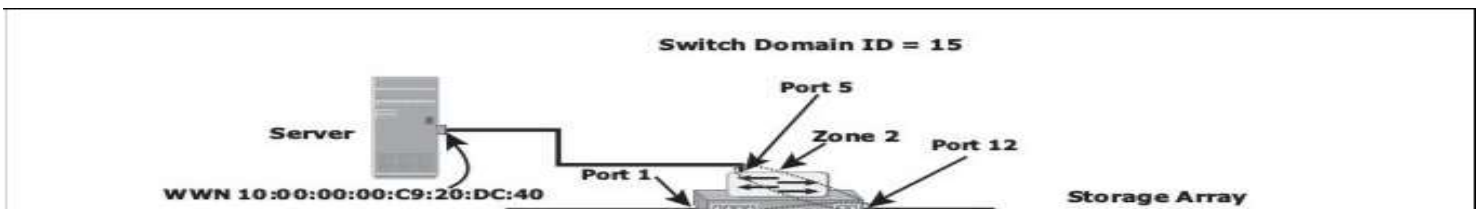

8. Write a note on NAS file sharing protocol

NAS File I/O

NAS uses file-level access for all of its I/O operations. File I/O is a high-level request that specifies the file to be accessed, but does not specify its logical block address. For example, a file I/O request from a client may specify reading 256 bytes from byte number 1152 onward in a specific file. Unlike block I/O, there is no disk volume or disk sector information in a file I/O request. The NAS operating system keeps track of the location of files on the disk volume and

converts client file I/O into block-level I/O to retrieve data. The NAS operating system issues a block I/O request to fulfill the file read and write requests that it receives. The retrieved data is again converted to file level I/O for applications and clients.

i) NFS

NFS is a client/server protocol for file sharing that is most commonly used on UNIX systems.

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

Currently, three versions of NFS are in use:

■■NFS version 2 (NFSv2): Uses UDP to provide a stateless network connection between a client and a server. Features such as locking are handled outside the protocol.

■■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.

■■NFS version 4 (NFSv4): This version uses TCP and is based on a stateful protocol design. It offers enhanced security.

ii) 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. It is a public, or open, variation of Server Message Block (SMB) protocol.

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.

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. The fault tolerance features of CIFS depend on whether an application is written to take advantage of these features.