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

Internal Assessment Test II – APRIL 2024-Solution

Sub:	Storage Area Network					Sub Code:	18CS822	Branch:	ISE
Date:	13-04-2024	Duration:	90 min's	Max Marks:	50	Sem/Sec:	VIII C		OBE

Answer any FIVE FULL Questions

		MARKS	CO	RBT
1	a) Mention the different types of iSCSI connectivity topologies. Explain those clearly using diagrams.	5+5	CO1	L2

Solution:

Two topologies of iSCSI implementations are native and bridged. **Native topology** does not have FC components. The initiators may be either directly attached to targets or connected through the IP network. **Bridged topology** enables the coexistence of FC with IP by providing iSCSI-to-FC bridging functionality.

Native iSCSI Connectivity

FC components are not required for iSCSI connectivity if an iSCSI-enabled array is deployed. In Figure 6-2 (a), the array has one or more iSCSI ports configured with an IP address and is connected to a standard Ethernet switch. After an initiator is logged on to the network, it can access the available LUNs on the storage array. A single array port can service multiple hosts or initiators as long as the array port can handle the amount of storage traffic that the hosts generate.

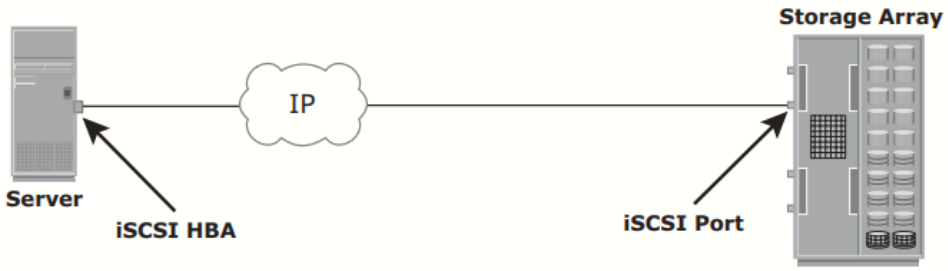

(a) Native iSCSI Connectivity

Bridged iSCSI Connectivity

A bridged iSCSI implementation includes FC components in its confi guration. Figure 6-2 (b) illustrates iSCSI host connectivity to an FC storage array. In this case, the array does not have any iSCSI ports. Therefore, an external device, called a gateway or a multiprotocol router, must be used to facilitate the communication between the iSCSI host and FC storage. The gateway converts IP packets to FC frames and vice versa. The bridge devices contain both FC and Ethernet ports to facilitate the communication between the FC and IP environments. In a bridged iSCSI implementation, the iSCSI initiator is configured with the gateway's IP address as its target destination. On the other side, the gateway is configured as an FC initiator to the storage array.

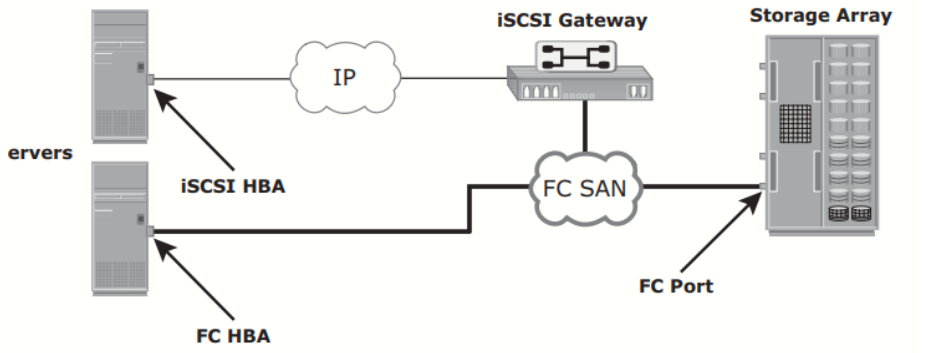

(b) Bridged iSCSI Connectivity

Combining FC and Native iSCSI Connectivity

The most common topology is a combination of FC and native iSCSI. Typically, a storage array comes with both FC and iSCSI ports that enable iSCSI and FC connectivity in the same environment, as shown in Figure 6-2 (c).

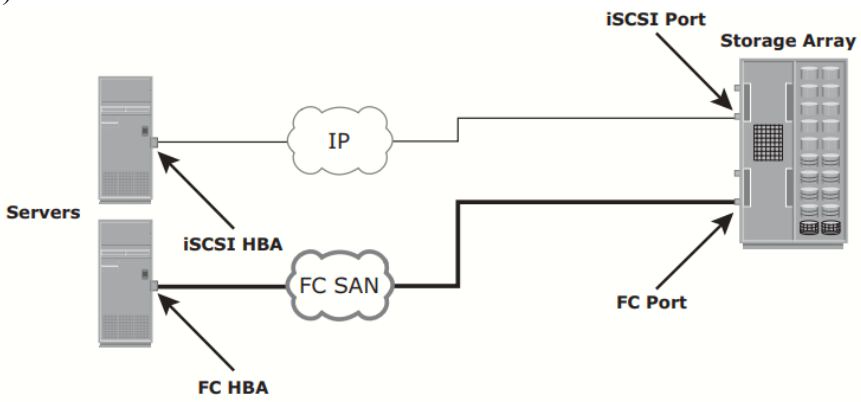

(c) Combining FC and Native iSCSI Connectivity

b) Draw a picture to illustrate the iSCSI protocol stack

Solution:

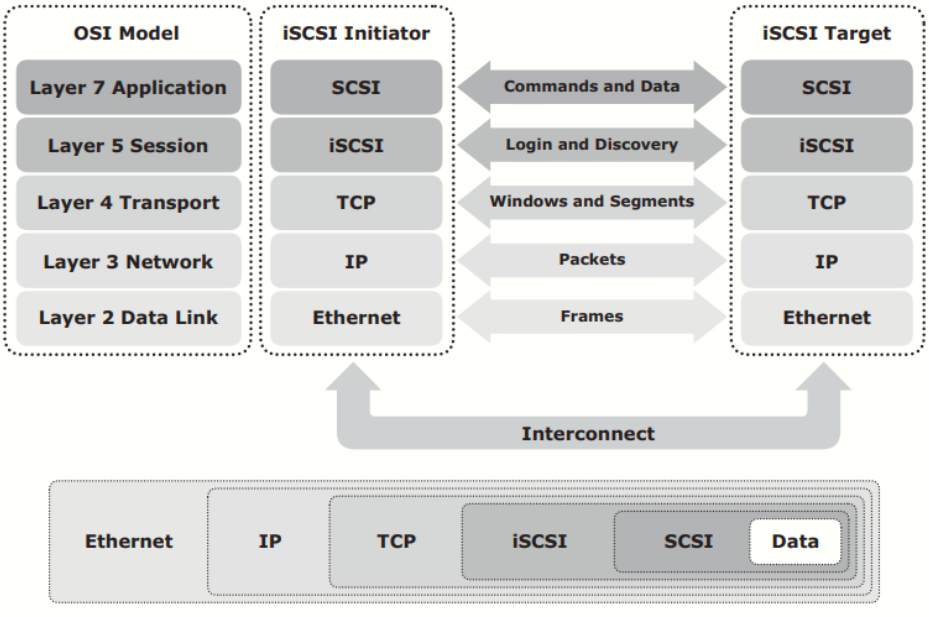

Figure 6-3: iSCSI protocol stack

2	a) Draw a diagram to describe the components of NAS. **Solution:**	10	CO2	L2

7.4 Components of NAS

A NAS device has two key components: NAS head and storage (see Figure 7-3). In some NAS implementations, the storage could be external to the NAS device and shared with other hosts. The NAS head includes the following components:

- CPU and memory
- One or more network interface cards (NICs), which provide connectivity to the client network. Examples of network protocols supported by NIC include Gigabit Ethernet, Fast Ethernet, ATM, and Fiber Distributed Data Interface (FDDI).
- An optimized operating system for managing the NAS functionality. It translates file-level requests into block-storage requests and further converts the data supplied at the block level to file data.
- NFS, CIFS, and other protocols for file sharing
- Industry-standard storage protocols and ports to connect and manage physical disk resources

The NAS environment includes clients accessing a NAS device over an IP network using file-sharing protocols.

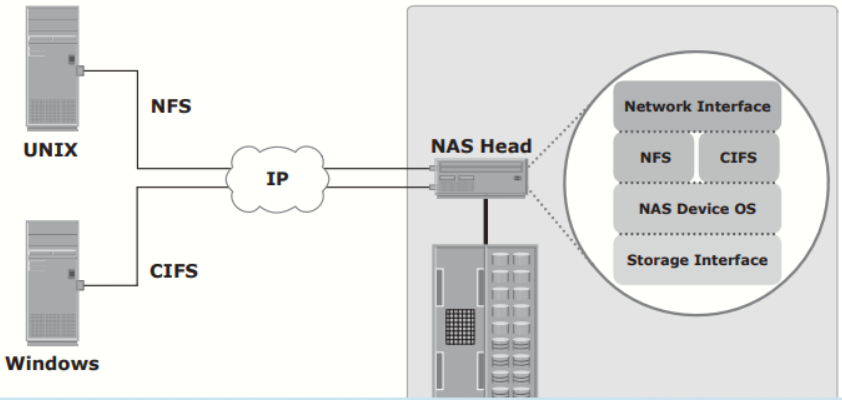

b) Describe and list the benefits of NAS.

Solution:

NAS offers the following benefits:

- **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.
- **Improved efficiency:** NAS delivers better performance compared to a general-purpose file server because NAS uses an operating system specialized for file serving.
- **Improved flexibility:** Compatible with 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.
- **Centralized storage:** Centralizes data storage to minimize data duplication on client workstations, and ensure greater data protection
- **Simplified management:** Provides a centralized console that makes it possible to manage file systems efficiently

	■ **Scalability:** Scales well 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 components that provide maximum connectivity options. A NAS device supports clustering technology for failover. ■ **Security:** Ensures security, user authentication, and file locking with industry-standard security schemas ■ **Low cost:** NAS uses commonly available and inexpensive Ethernet components. ■ **Ease of deployment:** Configuration at the client is minimal, because the clients have required NAS connection software built in.			
3	Discuss how to use NAS file-sharing protocols and how to use NAS I/O. **Solution**: Common Internet File System (CIFS) • Client-server application protocol that enables clients to access files that are on a server over TCP/IP ▸ An open variation of the Server Message Block (SMB) protocol • Stateful Protocol ▸ Maintains connection information regarding every connected client ▸ If a network failure or CIPS server failure occurs, client receives a disconnection notification ▸ Can automatically restore connections and reopen files that were open prior to interruption • Operates at the Application/Presentation layer of the OSI model • Most commonly used with Microsoft operating systems, but is platform-independent (available to Unix/Linux through Samba) Network File System (NFS) • Client-server application protocol that enables clients to access files that are on a server • Uses Remote Procedure Call (RPC) mechanism to provide access to remote file system ▸ Searching files and directories ▸ Opening, reading, writing to, and closing a file ▸ Changing file attributes ▸ Modifying file links and directories • Currently, 3 versions of NFS are in use: ▸ NFS v2 is stateless and uses UDP as transport layer protocol ▸ NFS v3 is stateless and uses UDP or optionally TCP as transport layer protocol ▸ NFS v4 is stateful and uses TCP as transport layer protocol	10	CO2	L2

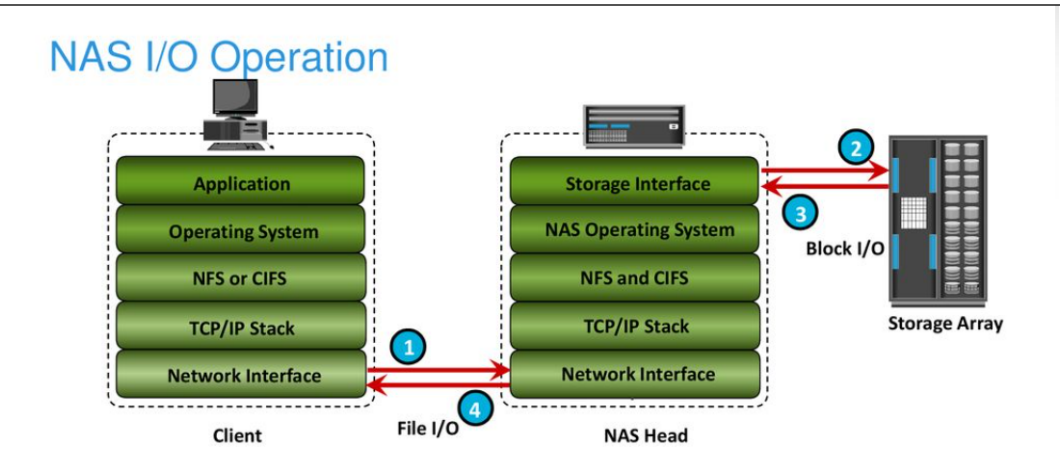

1. The requestor packages an **I/O request** into TCP/IP and forwards it through the network stack. The NAS device **receives** this request from the **network**.
2. The NAS device **converts** the I/O request into an **appropriate** physical storage **request**, which is a block-level I/O, and then performs the operation against the physical storage pool.
3. When the data is **returned** from the physical storage pool, the NAS device **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	Describe in detail about FCIP protocol along with its protocol stack, topology and performance, and security considerations.
	
	Solution:
	
	• FCIP gateway is connected to each fabric via a standard FC connection
	  • FCIP gateway at one end of the IP network **encapsulates** the FC frames into IP packets
	  • FCIP gateway at the other end **removes** the IP wrapper and **sends** the FC data to the layer 2 fabric
	• Fabric treats these gateways as layer 2 fabric switches
	• **IP address** is assigned to the port on the gateway, which is connected to an IP network.
	• After the **IP connectivity** is **established**, the nodes in the two independent fabrics can communicate with other
		10	CO1	L2

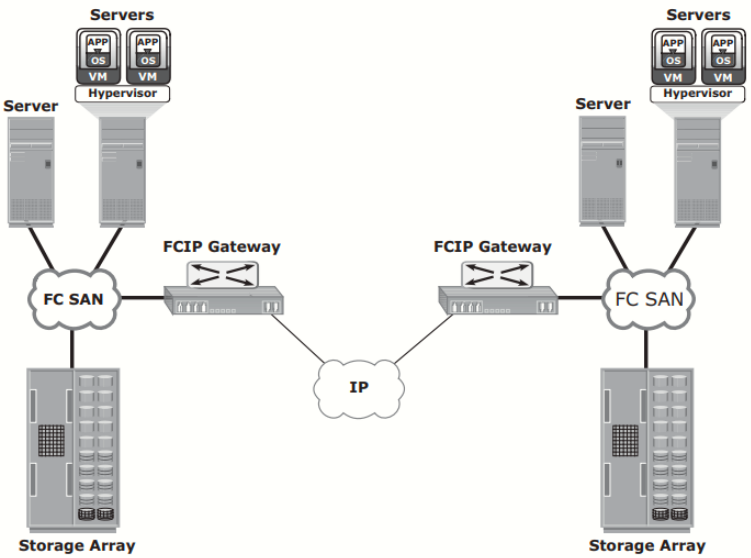

• **Figure 6-11:** FCIP topology

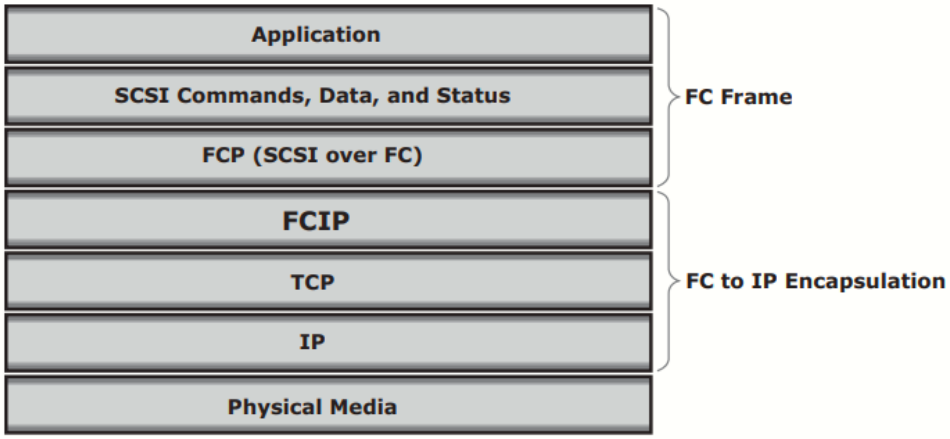

Figure 6-9: FCIP protocol stack

- **Performance, reliability, and security** should always be taken into consideration when implementing storage solutions.
- From the perspective of performance, configuring multiple paths between FCIP gateways eliminates single points of failure and provides increased bandwidth.
- 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

5	a. Explain how the performance of NAS can be affected if the TCP window size at the sender and receiver are not synchronized. **Solution:** This will affect the NAS performance as this may lead to **re-transmission** of data, **lower bandwidth utilization**, and performance degradation of the network, intermittent **connectivity, and data link errors**. b. A NAS implementation configured jumbo frames on the NAS head, with 9,000 as its MTU. However, the implementers did not see any performance improvement and actually experienced performance degradation. What could be the cause? Research the end-to-end jumbo frame support requirements in a network. **Solution:** ○ Jumbo frames are used at the end point (NAS Head) with MTU of 9000 ○ Check if the **intermediate network** uses a **different MTU size (e.g. 1500).** ○ This can cause the router to drop packets which then have to **re-transmit** at the TCP layer.	5+5	CO2	L3

	○ Packets are then **fragmented** and have to reassemble to accommodate the different MTU sizes. This degrades network performance.			
6	Describe in detail three common NAS implementations and justify which implementation is most suited for FC **Solution:** NAS Implementation • 3 types of NAS implementations: **Unified NAS** • Has **all** of its components and storage system in a **single** enclosure or frame **Gateway NAS** • NAS **head shares** its storage with SAN environment. **Scale-out NAS** • Ideal for enterprise data centers • Consolidating both virtualized and non-virtualized file storage into **one storage pool** with a single point of management NAS Implementation – Unified NAS • Consolidates **NAS-based** (file-level) and **SAN-based** (block-level) access on a **single** storage platform • Supports both CIFS and NFS protocols for file access and iSCSI and FC protocols for block level access • Provides **unified** management for both NAS head and storage 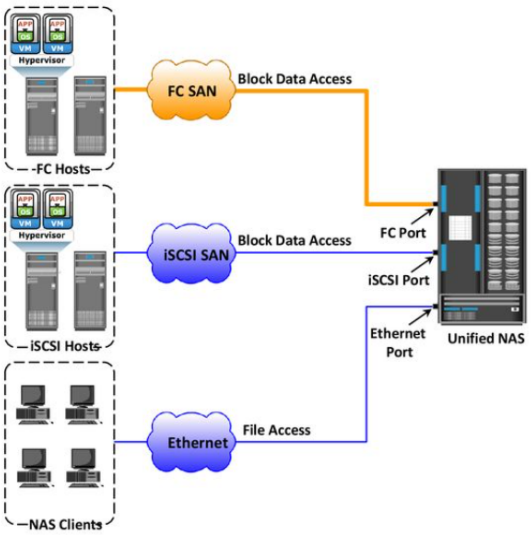 	10	CO2	L2

NAS Implementation – Gateway NAS

- Uses external and independently-managed storage
  - NAS heads access SAN-attached or direct-attached storage arrays
- NAS heads **share storage** with other application servers that perform block I/O
- Requires **separate management** of NAS head and storage
- The gateway NAS is the **most scalable** because NAS heads and storage arrays can be **independently scaled up when required**.
- Gateway NAS enables **high utilization** of storage capacity by sharing it with SAN environment.

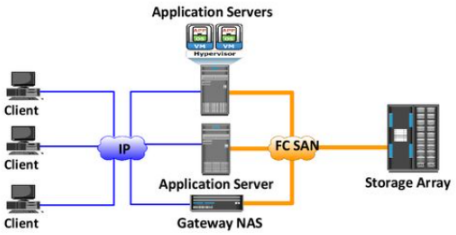

NAS Implementation – Scale-out NAS

- Pools **multiple nodes** together in a **cluster** that works as a single NAS device
  - Pool is managed centrally
- Scales performance and/or capacity with addition of nodes to the pool non-disruptively
- Creates a **single file system** that runs on **all nodes** in the cluster
  - Clients, connected to any node, can access entire file system
  - File system grows dynamically as nodes are added
- **Stripes** data across all nodes in a pool along with **mirror** or **parity** protection

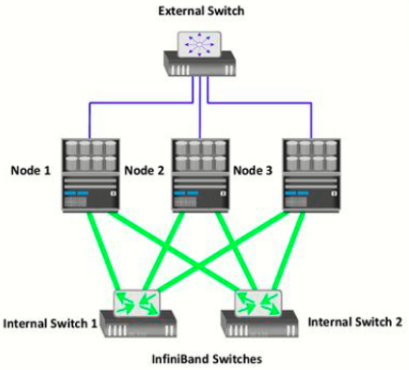

InfiniBand is a networking technology that provides a low-latency, high-bandwidth communication link between hosts and peripherals.