VTU QUESTION PAPER SOLUTION

18CS822-STORAGE AREA NETWORKS

1. a. Define Data center. List and explain key elements of the data center. Explain the characteristics of data center elements with a diagram.(10)

Solution:

Data Center

Data-center provides centralized data processing capabilities to companies.

Core Elements of a Data Center

Five core elements of a data-center:

1) Application

An application is a program that provides the logic for computing-operations.

- For example: Order-processing-application.

2) Database

DBMS is a structured way to store data in logically organized tables that are interrelated.

3 Server and OS

A computing-platform that runs 1) applications and 2) databases.

4. **Network**

A data-path that facilitates communication

1) between clients and servers or
2) between servers and storage.

5. Storage Array

A device that stores data permanently for future-use.

Key Characteristics for Data Center

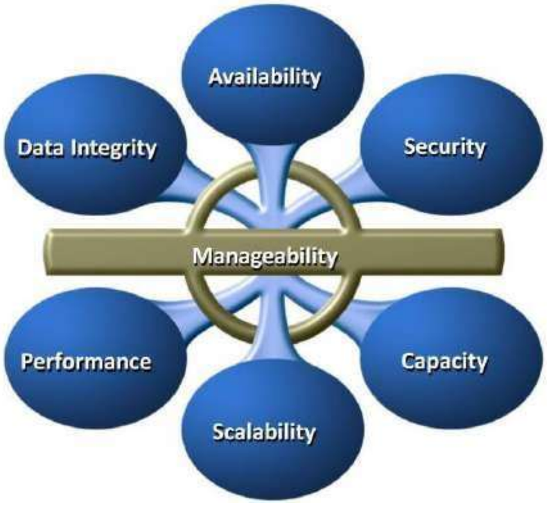

1) Availability

In data-center, all core-elements must be designed to ensure availability.

If the users cannot access the data in time, then it will have negative impact on the company.

2) Security

To prevent unauthorized-access to data,

Good polices & procedures must be used.

Proper integration of core-elements must be established.

Security-mechanisms must enable servers to access only their allocated-resources on the storage.

3) Scalability

It must be possible to allocate additional resources on-demand w/o interrupting normal-operations.

The additional resources includes CPU-power and storage.

The data-center must be able to support performance-requirements.

4) Data Integrity

Data integrity ensures that data is written to disk exactly as it was received.

For example: Parity-bit or ECC (error correction code).

Data-corruption may affect the operations of the company.

5) Storage Capacity

The data-center must have sufficient resources to store and process large amount of data efficiently.

When capacity-requirement increases, the data-center must be able

Capacity must be managed by reallocation of existing-resources rather than by adding new resources

6) Manageability

A data-center must perform all operations and activities in the most efficient manner.

Manageability is achieved through automation i.e. reduction of human-intervention in common tasks.

1.b Explain with neat diagram the evolution of storage architecture.(6)

<u>Solution:</u>

Evolution of Storage Architecture

Historically, organizations had centralized computers (mainframes) and information storage devices (tape reels and disk packs) in their data center. The evolution of open systems, their affordability, and ease of deployment made it possible for business units/departments to have their own servers and storage. In earlier implementations of open systems, the storage was typically internal to the server. These storage devices could not be shared with any other servers. This approach is referred to as server-centric storage architecture (see Figure 1-4)

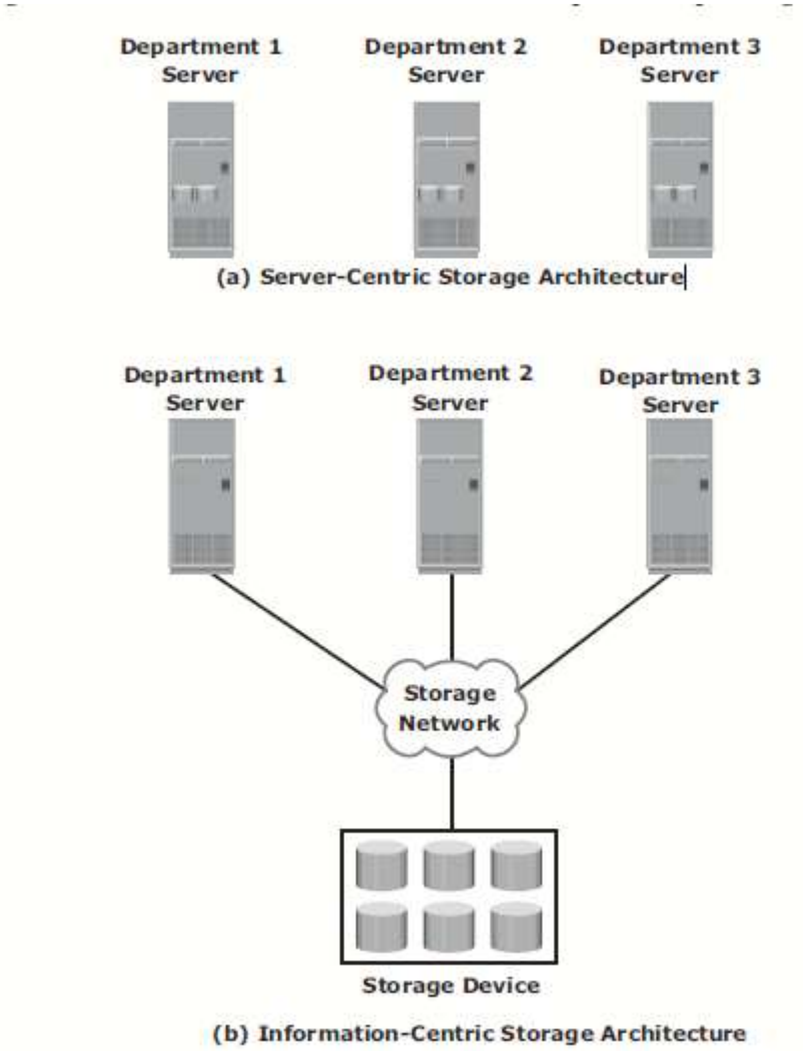

(a) Server-Centric Storage Architecture

(b) Information-Centric Storage Architecture

Figure 1-4: Evolution of storage architecture

To overcome these challenges, storage evolved from server-centric to *information-centric architecture* (see Figure 1-4 [b]). In this architecture, storage devices are managed centrally and independent of servers. These centrally managed storage devices are shared with multiple servers. When a new server is deployed in the environment, storage is assigned from the same shared storage devices to that server. The capacity of shared storage can be increased dynamically by adding more storage devices without impacting information availability. In this architecture, information management is easier and cost-effective. Storage technology and architecture continue to evolve, which enables organizations to consolidate, protect, optimize, and leverage their data to achieve the highest return on information assets.

1 c. With the diagram explain compute virtualization (4)

Compute virtualization is a technique for masking or abstracting the physical hardware from the operating system. It enables multiple operating systems to run concurrently on single or clustered physical machines. This technique enables creating portable virtual compute systems called *virtual machines* (VMs). Each VM runs an operating system and application instance in an isolated manner. Compute virtualization is achieved by a virtualization layer that resides between the hardware and virtual machines. This layer is also called the *hypervisor*. The hypervisor provides hardware resources, such as CPU, memory, and network to all the virtual machines. Within a physical server, a large number of virtual machines can be created depending on the hardware capabilities of the physical server. les, and so on.

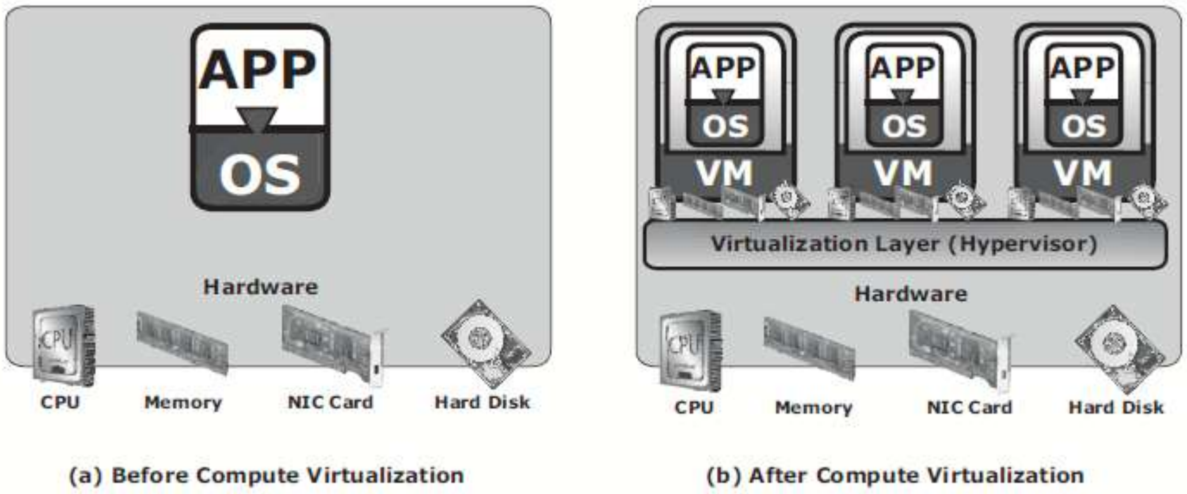

(a) Before Compute Virtualization

(b) After Compute Virtualization

Figure 2-3: Server virtualization

2. a. Explain disk drive components with diagram (10)

A disk drive uses a rapidly moving arm to read and write data across a flat platter coated with magnetic particles.

- Data is transferred from the magnetic platter through the R/W head to the computer.
- Several platters are assembled together with the R/W head and controller, most commonly referred to as a hard disk drive (HDD).
- Data can be recorded and erased on a magnetic disk any number of times.
- This section details the different components of the disk, the mechanism for organizing and storing data on disks, and the factors that affect disk performance.
- Key components of a disk drive area platter, spindle, read/write head, actuator arm assembly, andcontroller.

  1. Platter

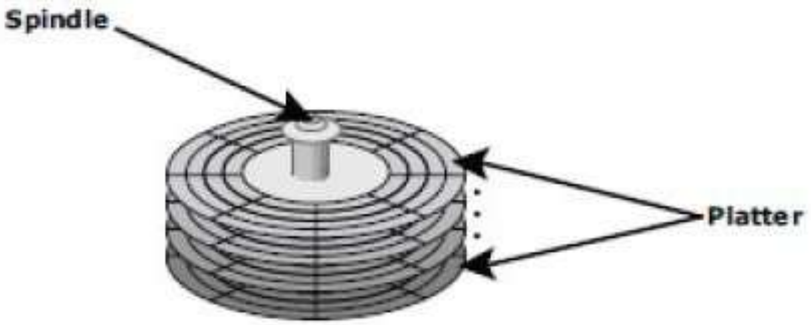

Figure: Spindle and platter

2. Spindle

A spindle connects all the platters, as shown in above figure and is connected to a motor.

The motor of the spindle rotates with a constant speed.

3. Read/Write Head

Read/Write (R/W) heads, shown in Figure, read and write data from or to a platter.

- Drives have two R/W heads per platter, one for each surface of the platter.

4. Actuator Arm Assembly

- The R/W heads are mounted on the *actuator arm assembly* which positions the R/W head at the location on the platter where the data needs to be written or read.

5. Controller

- The *controller* is a printed circuit board, mounted at the bottom of a disk drive.
- It consists of a microprocessor, internal memory, circuitry and firmware.

6. Physical Disk Structure

Data on the disk is recorded on the spindle.

2.b Discuss disk service time and disk Disk I/O controller utilization (10)

Disk Service Time
Disk service time is the time taken by a disk to complete an I/O request. Components that contribute to the service time on a disk drive are *seek time*, *rotational latency*, and *data transfer rate*.

Seek Time
The *seek time* (also called *access time*) describes the time taken to position the R/W heads across the platter with a radial movement (moving along the radius of the platter). In other words, it is the time taken to position and settle the arm and the head over the correct track. Therefore, the lower the seek time, the faster the I/O operation.

Full Stroke: The time taken by the R/W head to move across the entire width of the disk, from the innermost track to the outermost track.

Average: The average time taken by the R/W head to move from one random track to another, normally listed as the time for one-third of a full stroke.

Track-to-Track: The time taken by the R/W head to move between adjacent tracks.Each of these specifications is measured in milliseconds. The seek time of a disk is typically specified by the drive manufacturer. The average seek time on a modern disk is typically in the range of 3 to 15 milliseconds. Seek time has more impact on the read operation of random tracks rather than adjacent tracks

Rotational Latency
To access data, the actuator arm moves the R/W head over the platter to a particular track while the platter spins to position the requested sector under the R/W head. The time taken by the platter to rotate and position the data under the R/W head is called *rotational latency*.

Data Transfer Rate

The *data transfer rate* (also called *transfer rate*) refers to the average amount of data per unit of time that the drive can deliver to the HBA. It is important to first understand the process of read/write operations to calculate data transfer rates. In a *read operation*, the data first moves from disk platters to R/W heads; then it moves to the drive's internal *buffer*. Finally, data moves from the buffer through the interface to the host HBA. In a *write operation*, the data moves from the HBA to the internal buffer of the disk drive through the drive's interface.

Module II

3.a. List the different RAID levels, where parity techniques has been adopted. Explain any 3.(10)

Level	**Description**	**Fault isolation**
RAID 3	Byte-level striping with dedicated parity	Drive Firmware and Parity
RAID 4	Block-level striping with dedicated parity	Drive Firmware and Parity
RAID 5	Block-level striping with distributed parity	Drive Firmware and Parity

RAID-3

- RAID-3 uses both striping & parity techniques.
  1) Striping is used to improve performance of a storage-device.
  2) Parity is used to provide data-protection in case of disk-failure.
- Parity-information is stored on separate, dedicated-disk.

- Data is striped across all disks except the parity-disk in the array.
- In case of disk-failure, parity can be used for reconstruction of the missing-data.

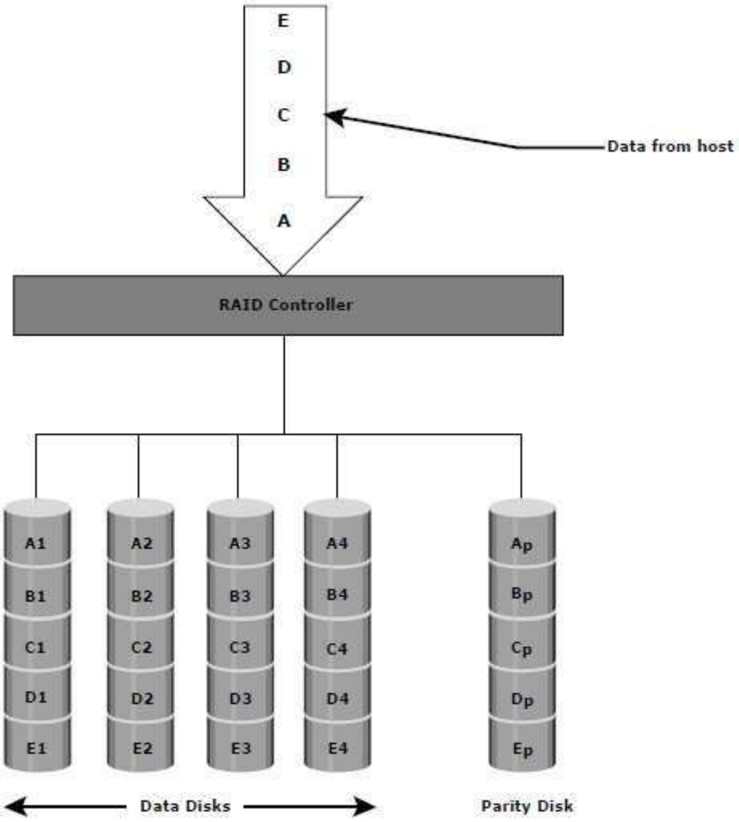

RAID-4

- Similar to RAID-3, RAID-4 uses both striping & parity techniques.
  1) Striping is used to improve performance of a storage-device.
  2) Parity is used to provide data-protection in case of disk-failure.
- Parity-information is stored on a separate dedicated-disk.
- Data is striped across all disks except the parity-disk.
- In case of disk-failure, parity can be used for reconstruction of the missing-data.

RAID-5

- Problem:

  In RAID-3 and RAID-4, parity is written to a dedicated-disk.

  If parity-disk fails, we will lose our entire backup.

  Solution: To overcome this problem, RAID-5 is proposed.

  In RAID-5, we distribute the parity-information evenly among all the disks.
- RAID-5 similar to RAID-4 because

  it uses striping and the drives (strips) are independently accessible.

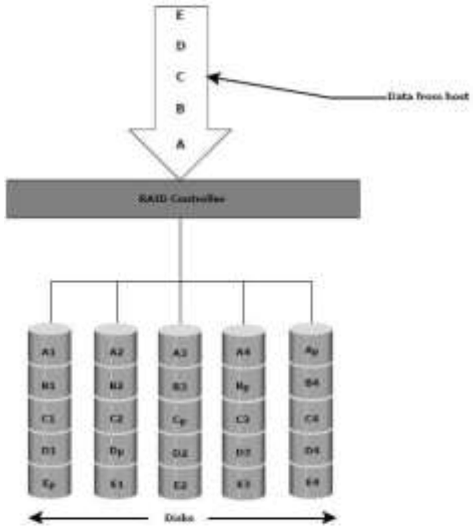

Figure 1-18: RAID-5

3.b Differentiate between hardware and software RAID (4)

S. No	Hardware RAID	Software RAID
1.	Hardware RAID is a customized processing system, using different controllers or RAID cards to manage the RAID design independently from the OS.	Software RAID uses the processing power of that computer's operating system in which the RAID disks are installed.
2.	Hardware RAID is more reliable and expensive.	The cost is low because no additional hardware RAID controller is required.
3.	Inconsistent performance for certain hardware RAID setups that use flash storage (SSD), HDD arrays.	In Software RAID the processors can easily handle RAID 0 & 1 processing with no noticeable performance hit.
4.	Replacing failed disk is simple – Just take it out and put in a new one	Replacing a failed disk in the software RAID is a bit more complex. We have to first tell our system to stop using the disk and then replace the disk.
5.	When the RAID controller goes down, it should get replaced with an identical model to avoid malfunction.	We can implement Software RAID configuration on one operating system (e.g. Ubuntu) and use it across other systems.

3.c With diagram explain different RAID techniques (6)

RAID Techniques

RAID-levels are defined based on following 3 techniques:

Striping (used to improve performance of storage)

Mirroring (used to improve data-availability) and

Parity (used to provide data-protection)

Striping

Striping is used to improve performance of a storage-device.

It is a technique of splitting and distribution of data across multiple disks.

Main purpose: To use the disks in parallel.

It can be bitwise, byte-wise or block wise.

A **RAID-set** is a group of disks.

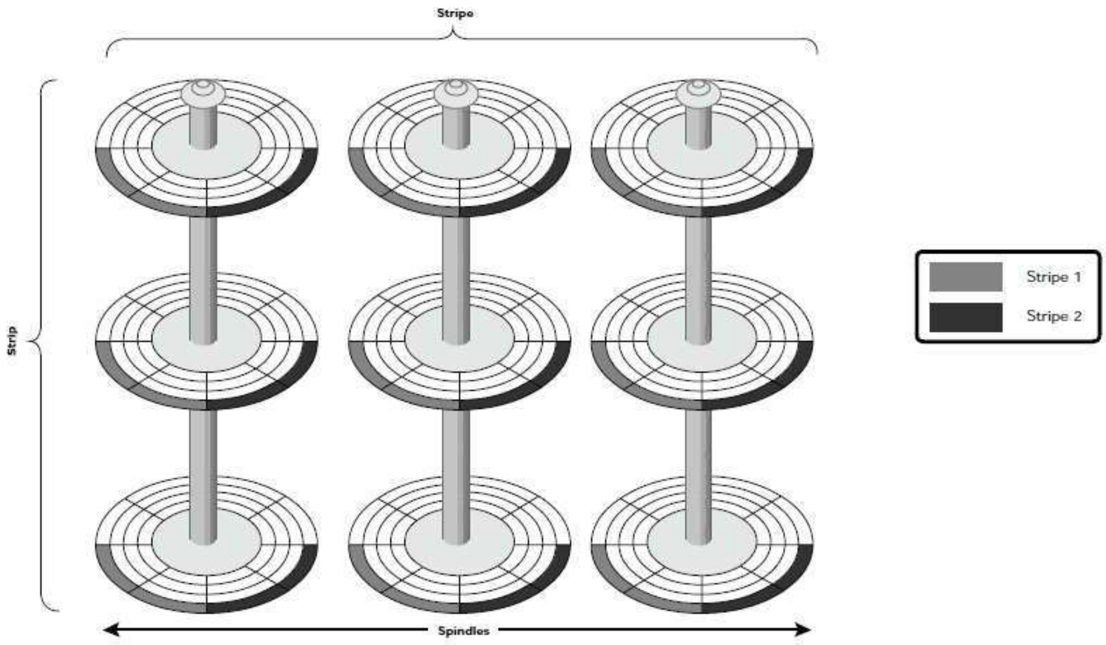

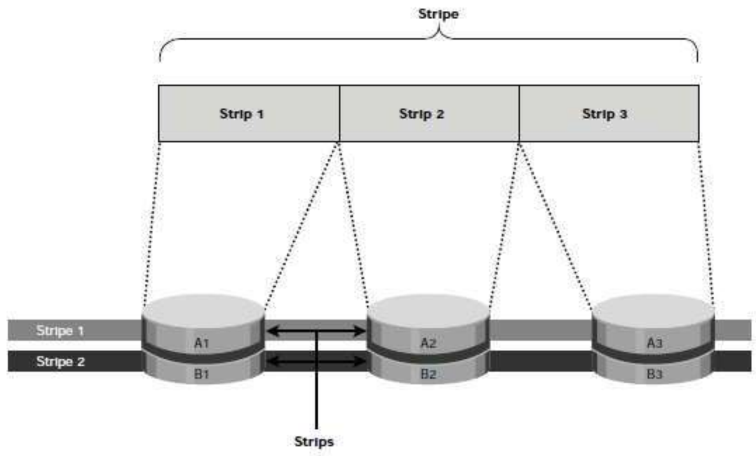

Mirroring

Mirroring is used to improve data-availability (or data-redundancy).

All the data is written to 2 disks simultaneously. Hence, we have 2 copies of the data.

In case of failure of one disk, the data can be accessed on the surviving-disk (Figure 1-12).

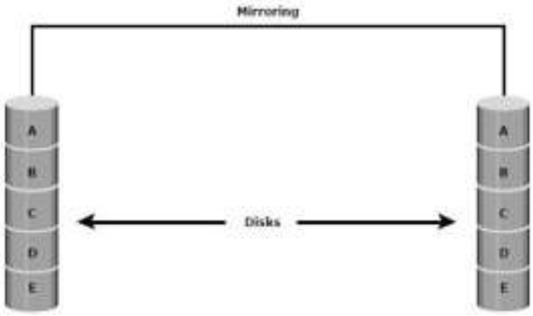

Figure 1-12: Mirrored disks in an array

Parity

Parity is used to provide data-protection in case of a disk-failure.

An additional disk is added to the stripe-width to hold parity.

In case of disk-failure, parity can be used for reconstruction of the missing-data.

Parity is a technique that ensures protection of data without maintaining a duplicate-data

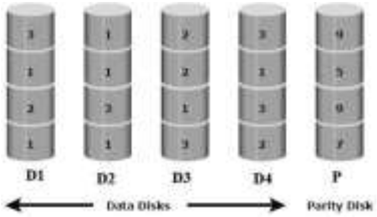

Figure 1-13: Parity RAID

4.a. Explain types of intelligent storage systems (10)

Types of Intelligent Storage System

- Storage-devices can be classified into 2 types:
  1) High-end storage-system
  2) Midrange storage-system

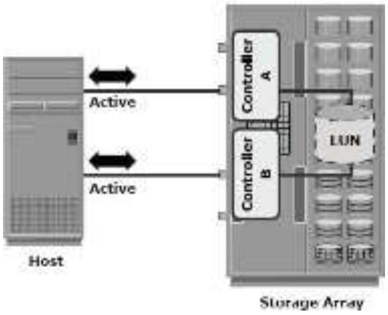

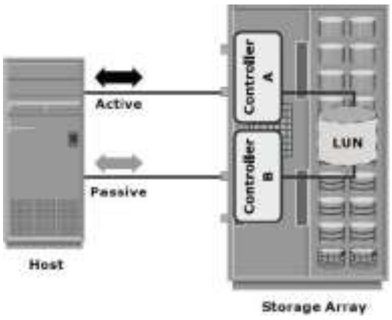

Figure 1-28: Active-active configuration

Figure 1-29: Active-passive configuration

High End Storage System

High-end storage-device is also known as **active-active array**.

It is suitable for large companies for centralizing corporate-data. (e.g. Big Bazaar)

An active-active array implies that

→ the host can transfer the data to its logical-units using any of the available paths (Figure 1-28)
It provides the following capabilities:

1) Large number of controllers and cache.
2) Multiple front-end ports to serve a large number of hosts.
3) Multiple back-end controllers to perform disk-operations optimally.
4) Large storage-capacity.
5) Large cache-capacity to service host's requests optimally.
6) Mirroring technique to improve data-availability.
7) Interoperability: Connectivity to mainframe-computers and open-systems hosts.

8) Scalability to support following requirements:
9) -requests from a no. of servers and applications.

Midrange Storage System

Midrange storage-device is also referred to as **active-passive array**.

It is suitable for small- and medium-sized companies for centralizing corporate-data.

It is designed with 2 controllers.

Each controller contains

→ host-interfaces
→ cache

→ RAID-controllers, and
→ interface to disks.

4.b Explain with diagram the components of fiber channel components (10).

Describe the Fibre Channel SAN components.

Solution:

Components of FC SAN

1. Node (server and storage) ports
2. Cables
3. Connectors
4. Interconnecting devices such as FC switches and hubs
5. SAN management software

Node (server and storage) ports

In FC 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.

Exist on

1. Host bus adapter (HBA) in server
2. Front-end adapters in storage

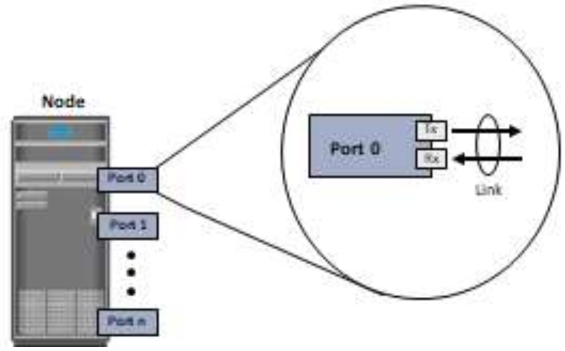

3. **Cables**
    1. Copper cables for short distance (back-end)(acceptable signal-to-ratio for distance up to 30 meters).

2. Optical fiber cables for long distance. carry data in the form of light.
3. Two types of optical cables: single-mode and multimode

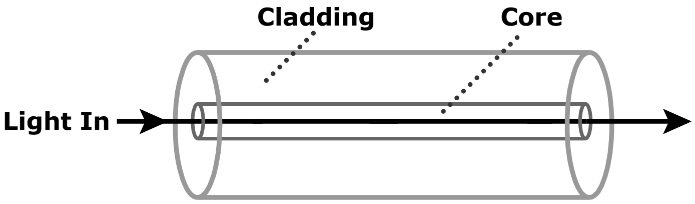

Single-mode Fiber

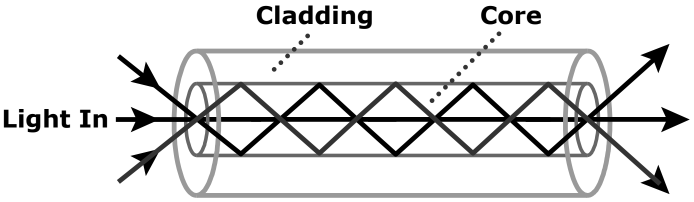

Multimode Fiber

3**. Connectors**

Attached at the end of a cable to enable swift (rapid) connection and disconnection of the cable to and from a port.
Commonly used connectors for fiber optic cables are:

- Standard Connector (SC)
  - Duplex connectors
- Lucent Connector (LC)
  - Duplex connectors
- Straight Tip (ST)

4. Patch panel connectors
5. Simplex connectors

4. Interconnecting Devices

Hubs

- **Physically** connect nodes in a logical loop or a physical star topology
- Provide **limited** connectivity and scalability
- All 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

- More **intelligent** than hubs and directly route data from one physical port to another
- Switches are available with fixed port count or modular design
- Nodes **do not share** the bandwidth.
- Instead, each node has a dedicated communication path

Directors

- **High-end switches** with a **higher port count** and better fault tolerance capabilities.
- Always modular, and its port count can be increased by inserting additional 'line cards' or 'blades'
- High-end switches and directors contain redundant components

5. **SAN Management Software**

- A suite of tools used in a SAN to manage interfaces between host and storage arrays
- Management of various resources from one central console.
- Provides integrated management of SAN environment
- Key management functions:
  1. Mapping of storage devices, switches and servers
  2. Monitoring and generating alerts for discovered devices
  3. Logical partitioning of SAN (zoning)

Management of SAN components (HBAs, storage components and interconnecting devices)

Module 3

5.a Explain different iSCSI topologies with diagram.(6)

(1) **Native iSCSI**

- iSCSI initiators are either **directly attached to storage array** or **connected through IP network**
  - No FC component
  - Storage array has iSCSI port
- Each iSCSI port on the array is configured with an IP address and port number.

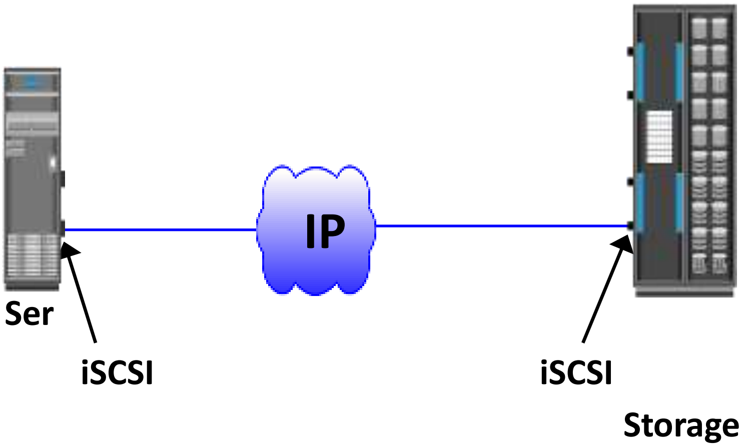

(2) **Bridged iSCSI**

- iSCSI gateway is used to enable communication between iSCSI host and FC storage
- iSCSI gateway works as **bridge** between FC and IP network
  - Converts IP packets to FC frames and vice versa
- iSCSI initiator is configured with gateway's IP address as its target
- iSCSI gateway is configured as FC initiator to storage array

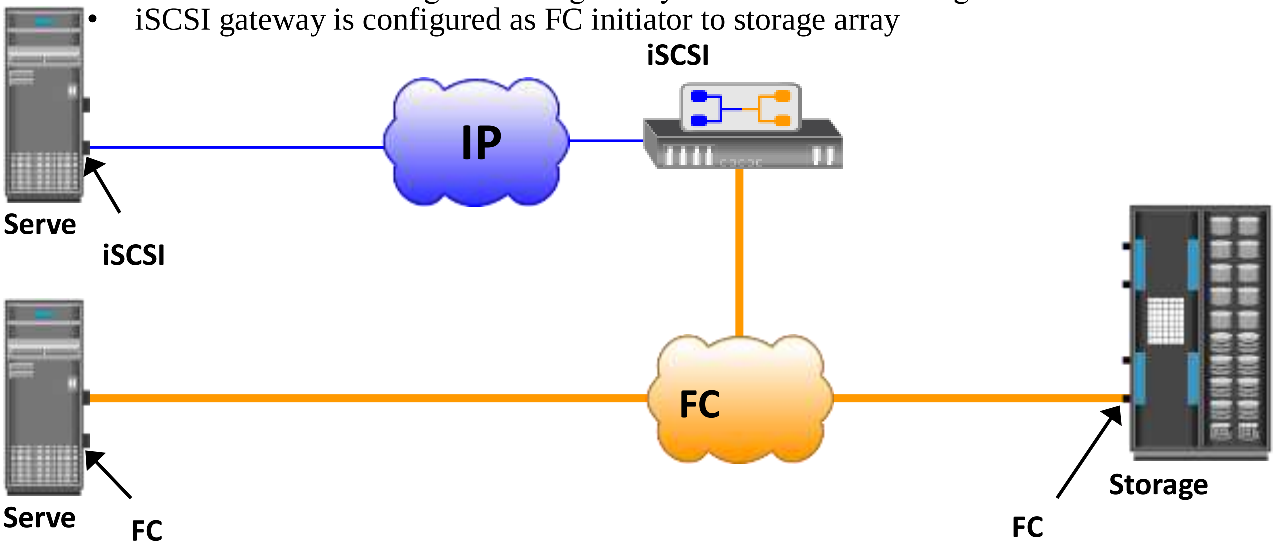

5.b Briefly explain about iSCSI protocol stack (4)

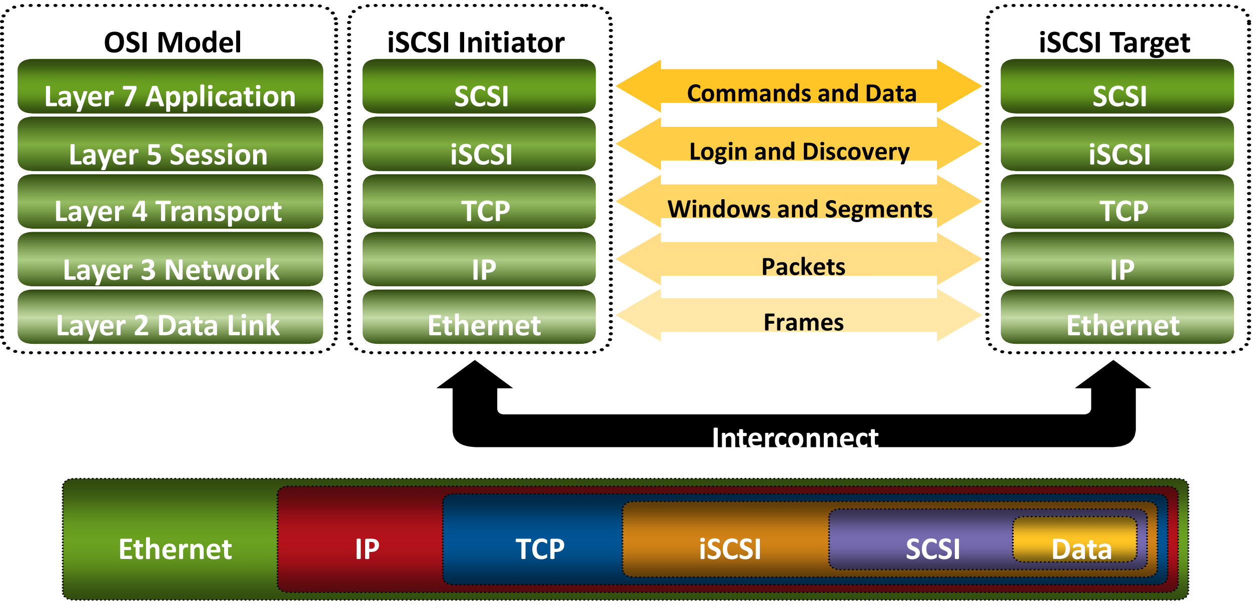

5.c Explain the components of NAS with neat sketch. Briefly explain the benefits of NAS(10)

Components of NAS

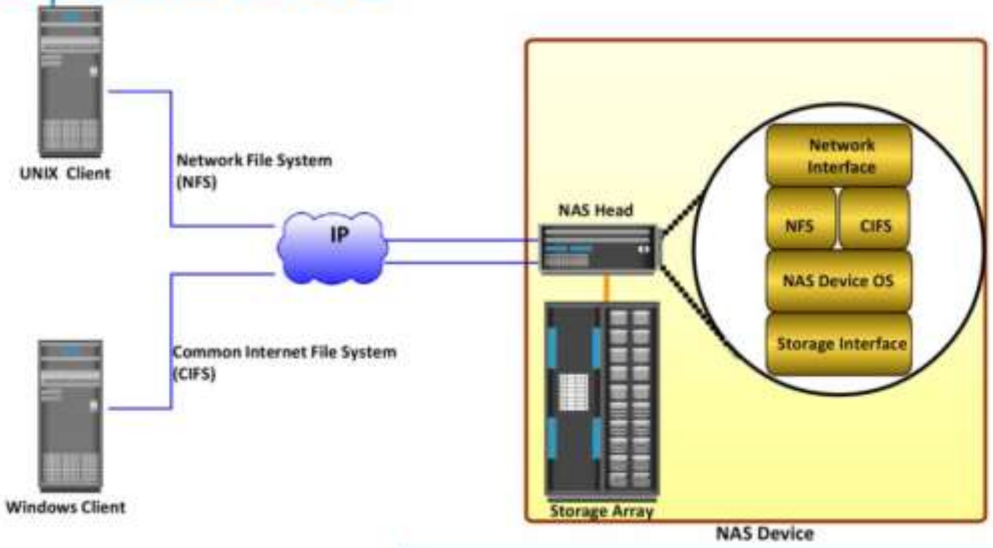

1. NAS head (CPU and Memory)
2. One or more NICs - provide connectivity to the network
3. An optimized operating system for managing NAS functionality
4. NFS and CIFS protocols for file sharing
5. Industry-standard storage protocols to connect and manage physical disk resources, such as ATA, SCSI, or FC

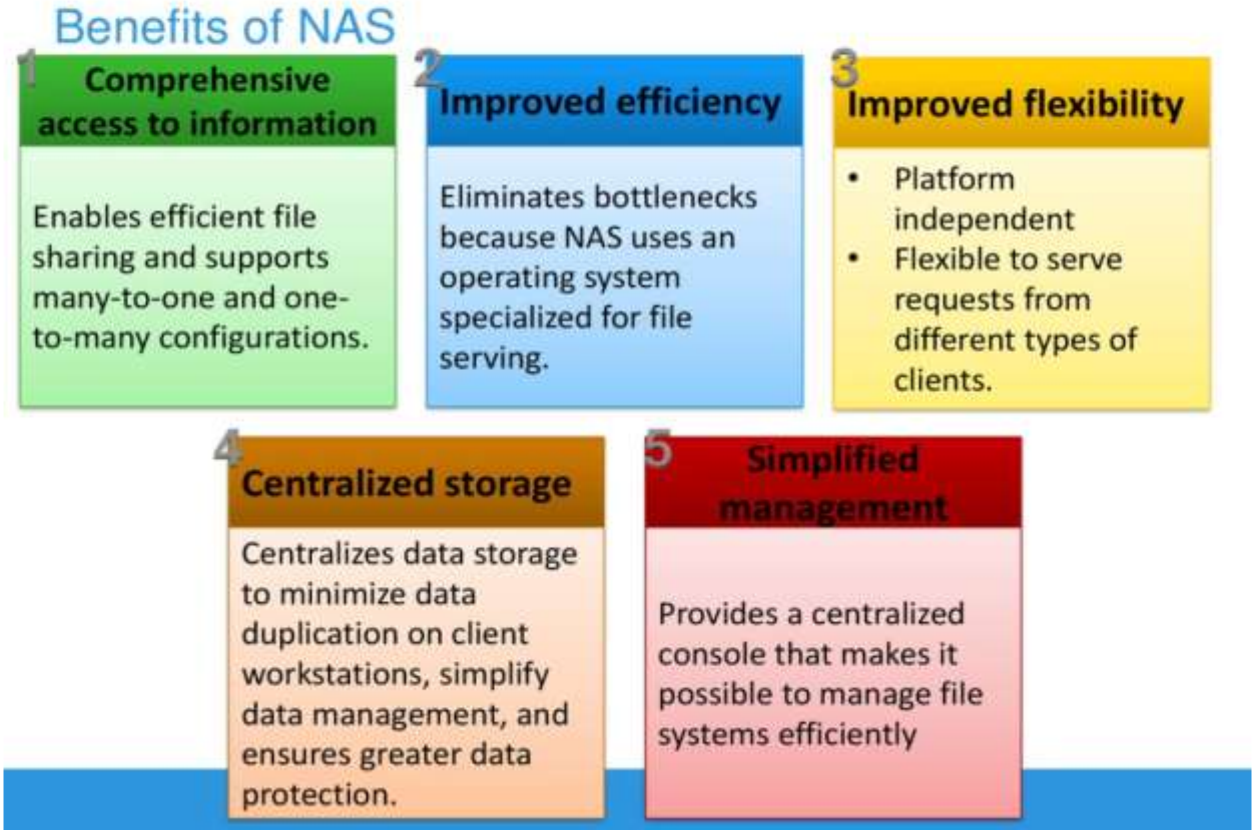

Benefits of NAS

1 **Comprehensive access to information**

Enables efficient file sharing and supports many-to-one and one-to-many configurations.

2 **Improved efficiency**

Eliminates bottlenecks because NAS uses an operating system specialized for file serving.

3 **Improved flexibility**

- Platform independent
- Flexible to serve requests from different types of clients.

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

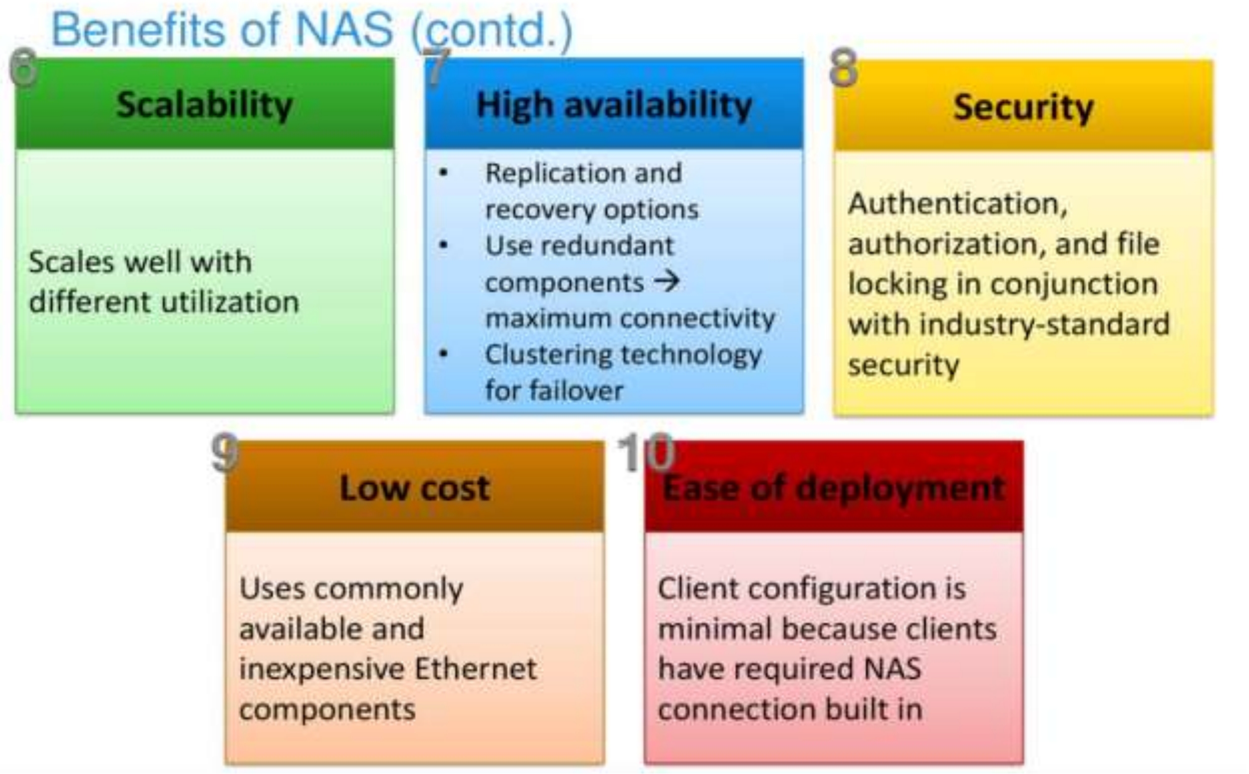

Benefits of NAS (contd.)

6 **Scalability**

Scales well with different utilization

7 **High availability**

- Replication and recovery options
- Use redundant components → maximum connectivity
- Clustering technology for failover

8 **Security**

Authentication, authorization, and file locking in conjunction with industry-standard security

9 **Low cost**

Uses commonly available and inexpensive Ethernet components

10 **Ease of deployment**

Client configuration is minimal because clients have required NAS connection built in

6.a Explain NAS implementation in detail (6)

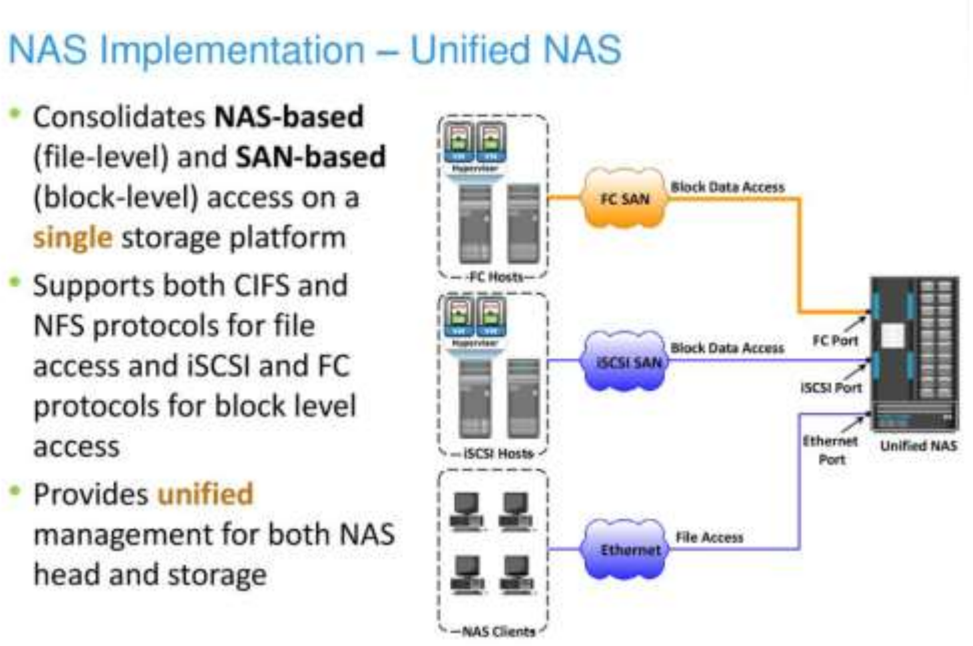

NAS Implementation – Unified NAS

- 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

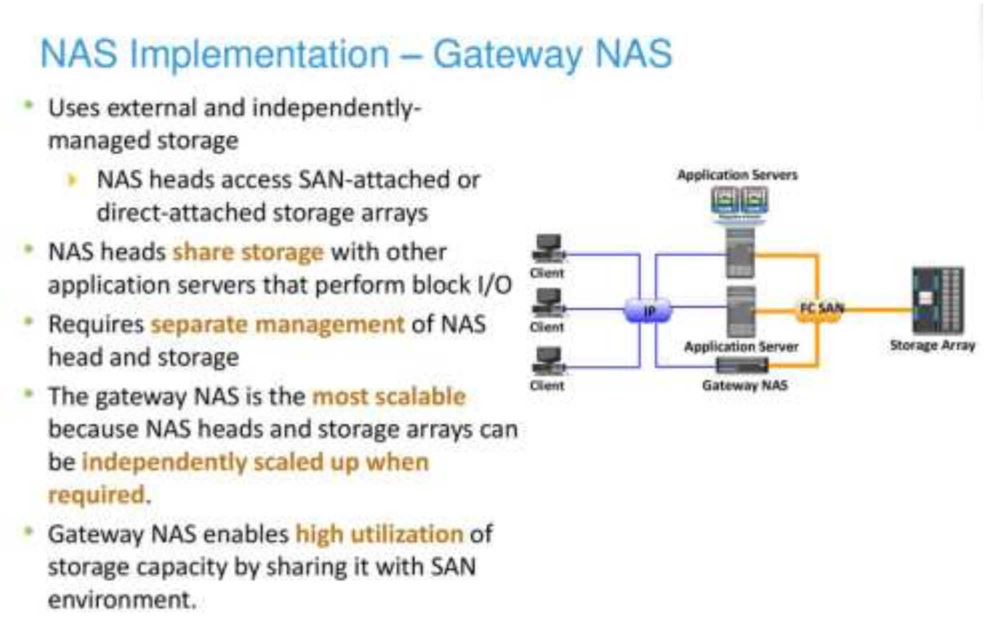

NAS Implementation – Gateway NAS

- 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.

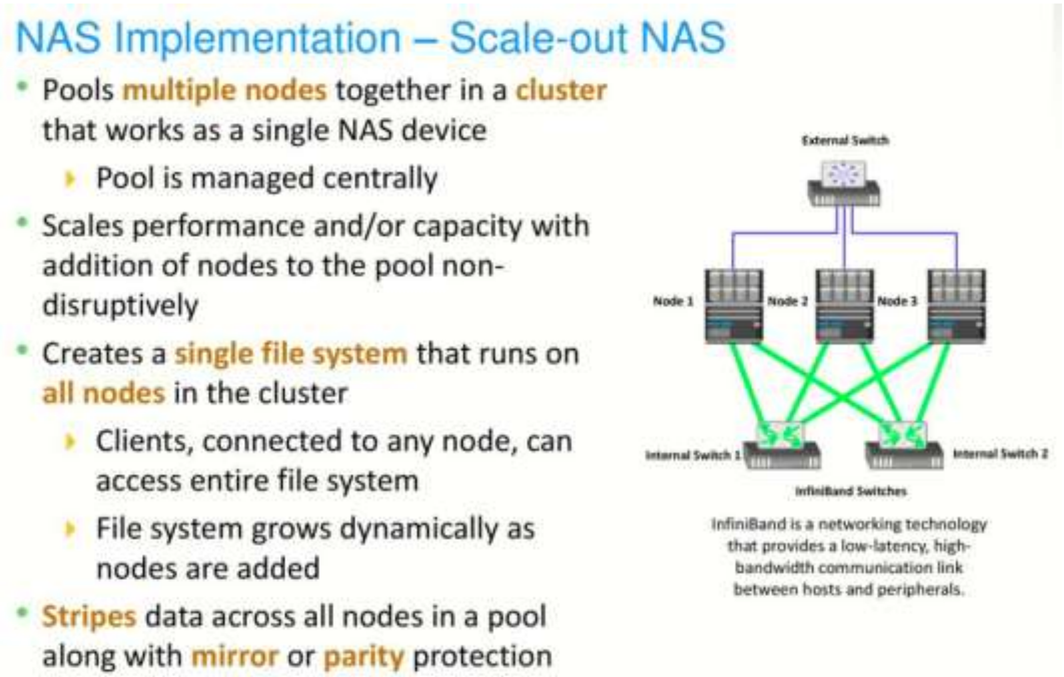

NAS Implementation – Scale-out NAS

- 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.

6.b Discuss about NAS file sharing protocols.(4)

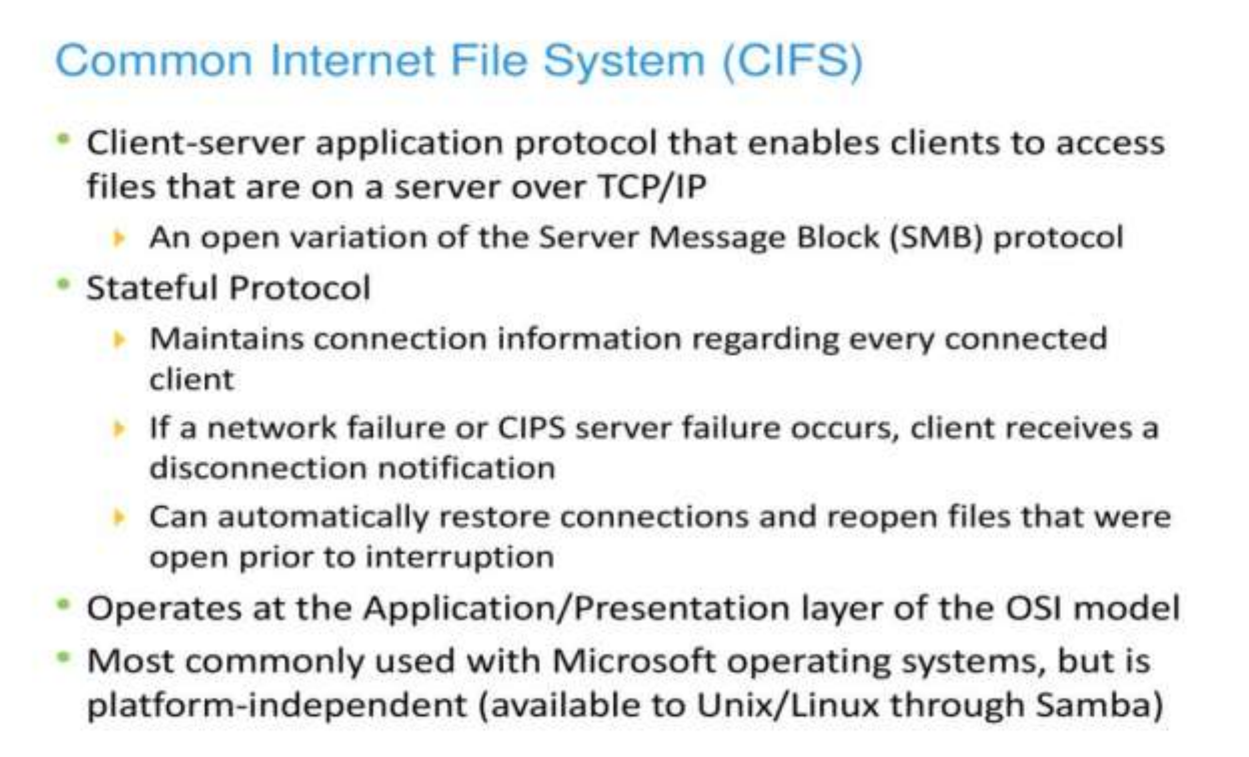

Common Internet File System (CIFS)

- 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)

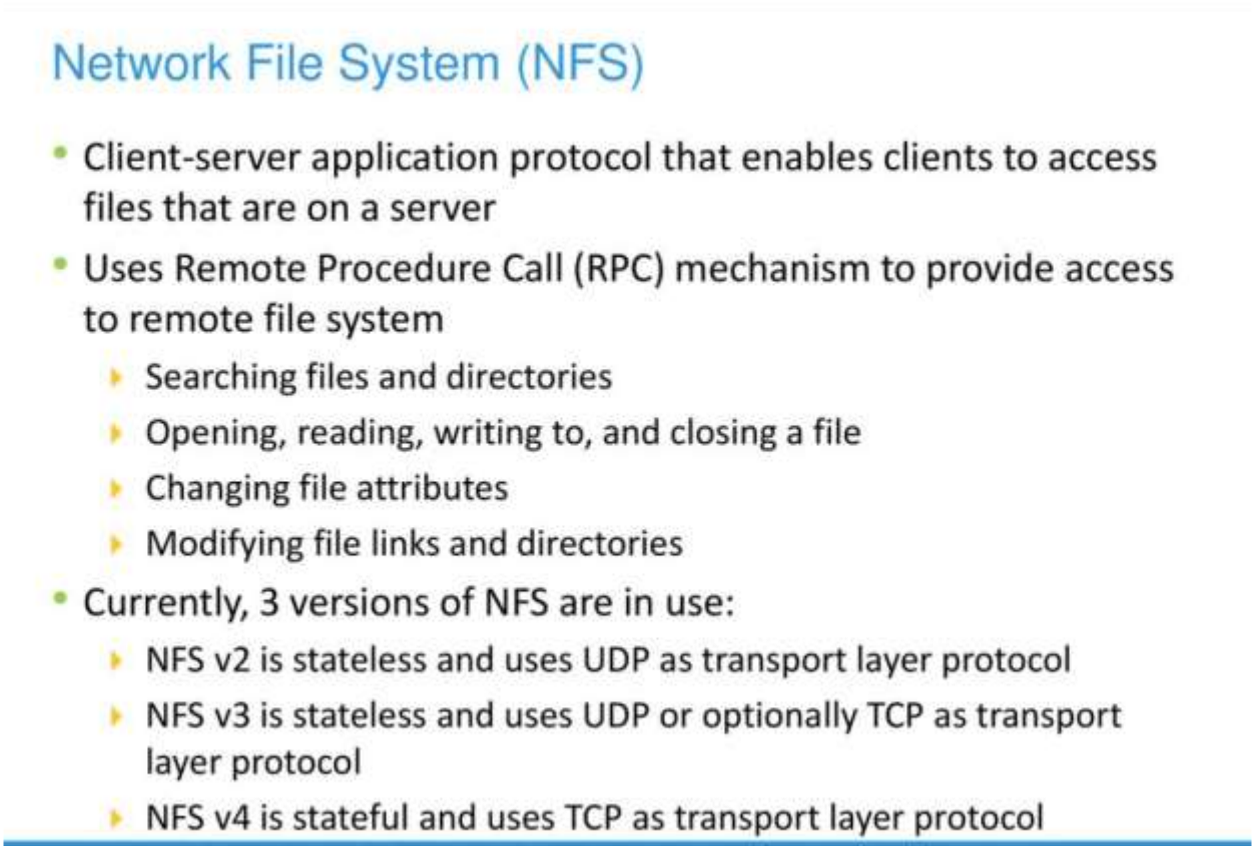

Network File System (NFS)

- 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

6.c With diagram explain about FCIP protocol stack and FCIP topology(10)

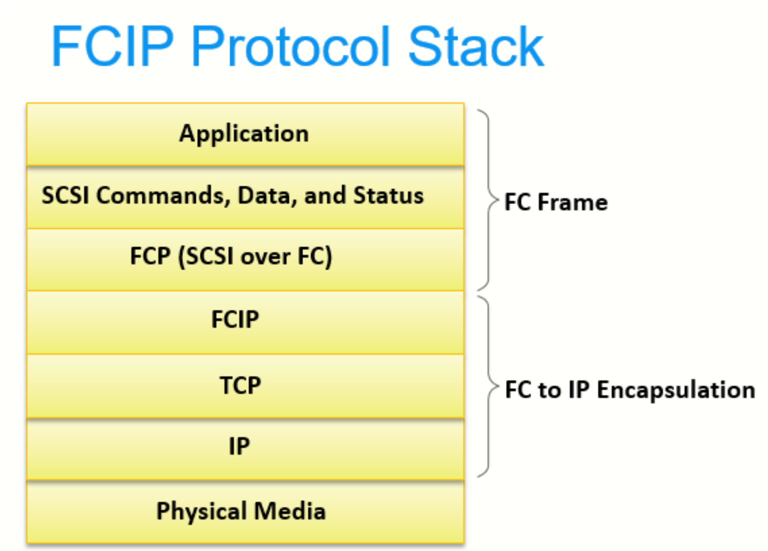
FCIP Protocol Stack

- Applications generate SCSI commands and data
- Upper layer protocol SCSI includes the SCSI driver program that executes the read-and-write commands
- FibreChannel Protocol (FCP) layer enables the FC frames to run natively within a SAN fabric environment
- FCIP layer encapsulates the FibreChannel frames onto the IP payload and passes them to the TCP layer
- TCP and IP are used for transporting the encapsulated information across Ethernet, wireless, or other media that support the TCP/IP traffic

FCIP Topology:

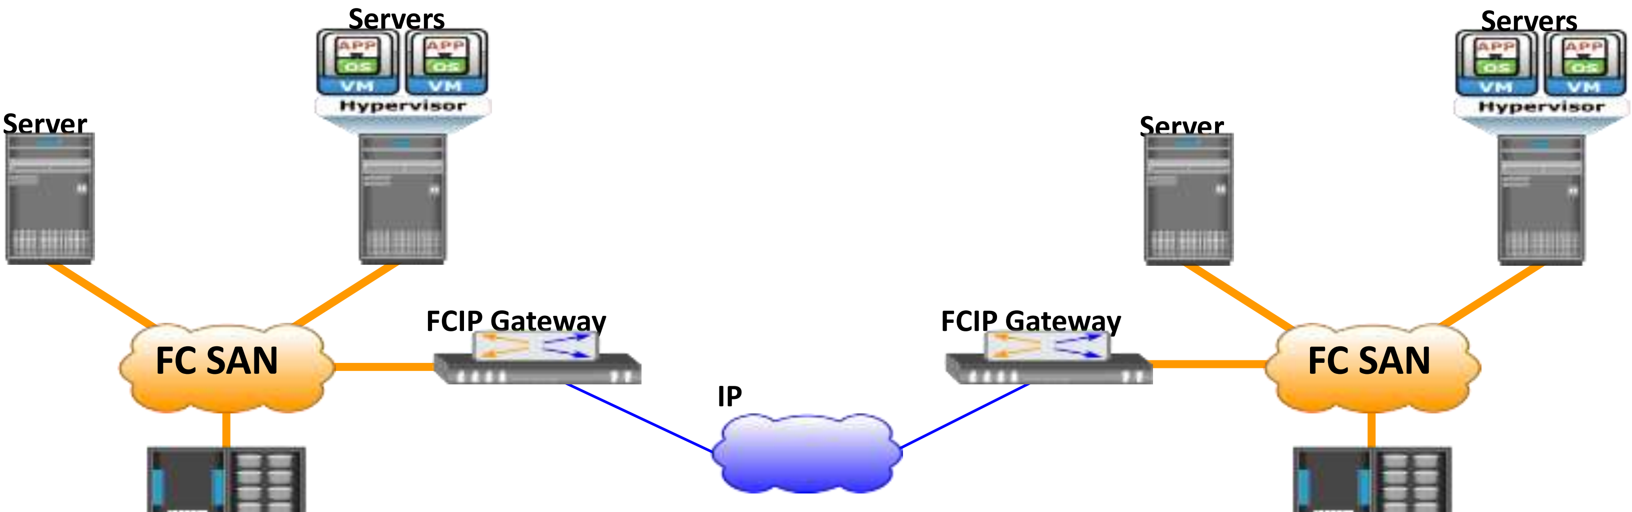

Module 4

7. a Describe failure analysis in business continuity. Mention important BC technologies solutions.(10)

Solution:

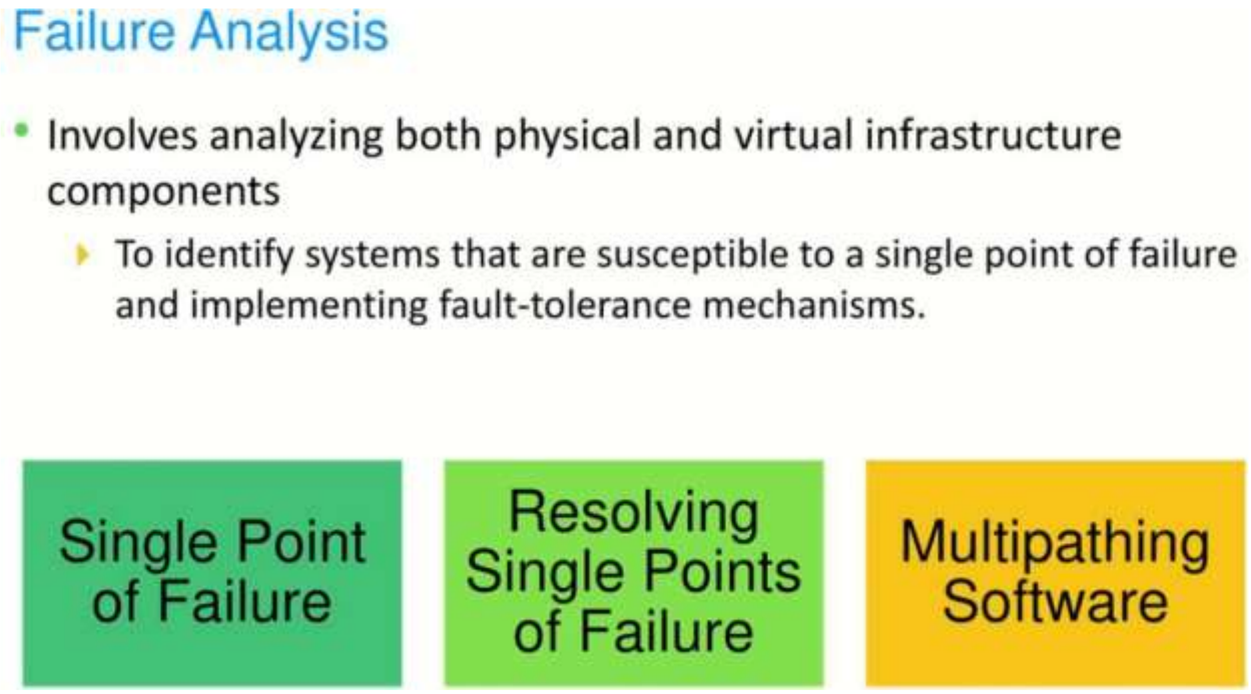

Failure Analysis

- Involves analyzing both physical and virtual infrastructure components
  - To identify systems that are susceptible to a single point of failure and implementing fault-tolerance mechanisms.

Single Point of Failure

Resolving Single Points of Failure

Multipathing Software

Failure Analysis: (1) Single Points of Failure

Single Points of Failure

It refers to the failure of a component of a system that can terminate the availability of the entire system or IT service.

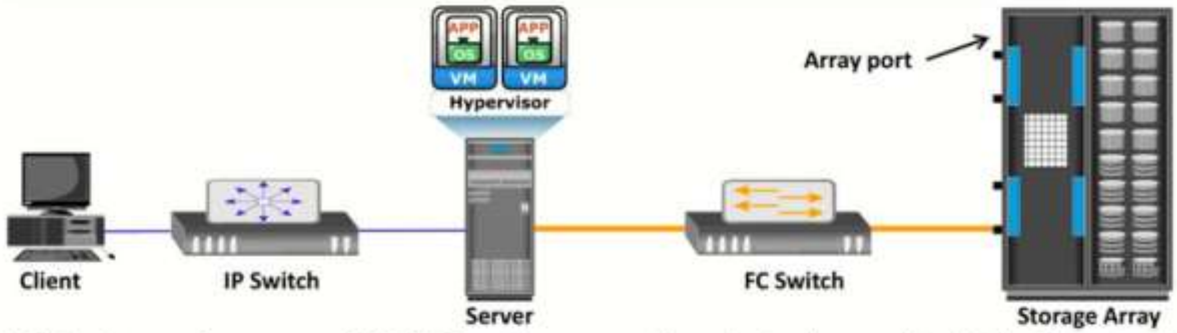

A VM, a hypervisor, or an HBA/NIC on the server, the physical server itself, the IP network, the FC switch, the storage array port, or even the storage array could be a potential single point of failure

E.g.: For example, failure of a hypervisor can affect all the running VMs and virtual network, which are hosted on it

Failure Analysis: (2) Resolving Single Points of Failure / Fault tolerant

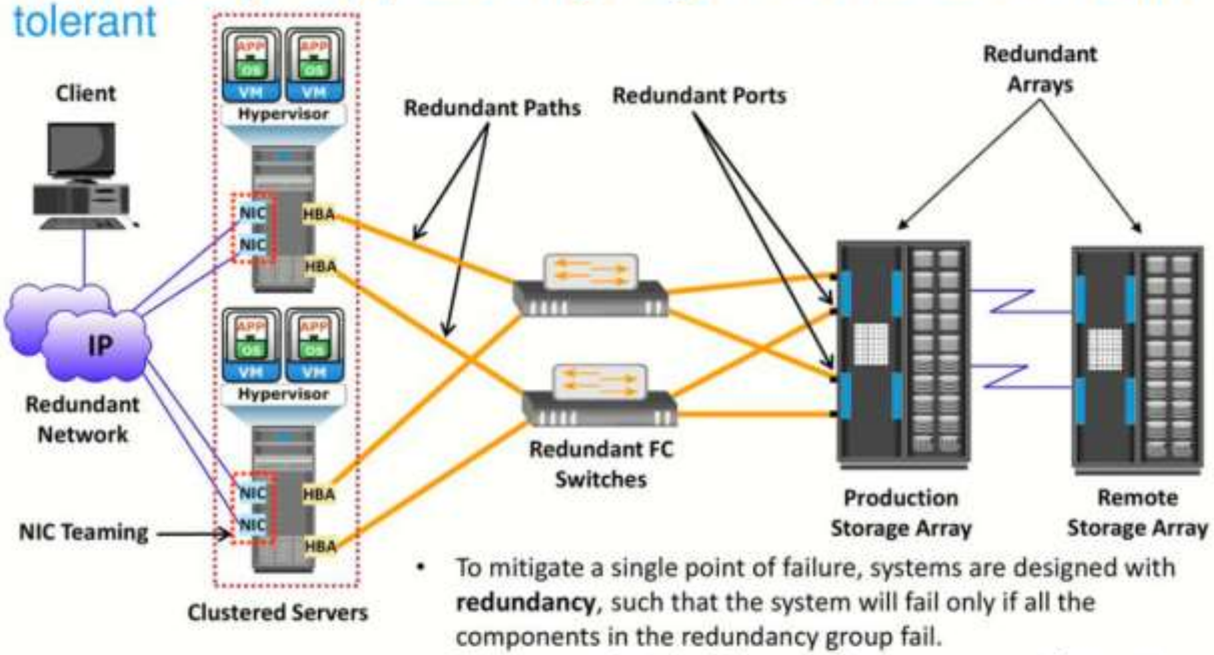

- To mitigate a single point of failure, systems are designed with **redundancy**, such that the system will fail only if all the components in the redundancy group fail.
- This ensures that the failure of a single component **does not affect data availability**.
- **Careful analysis** is performed to eliminate every single point of failure

7. b With neat diagram explain the steps involved in backup and restore operation.(10)

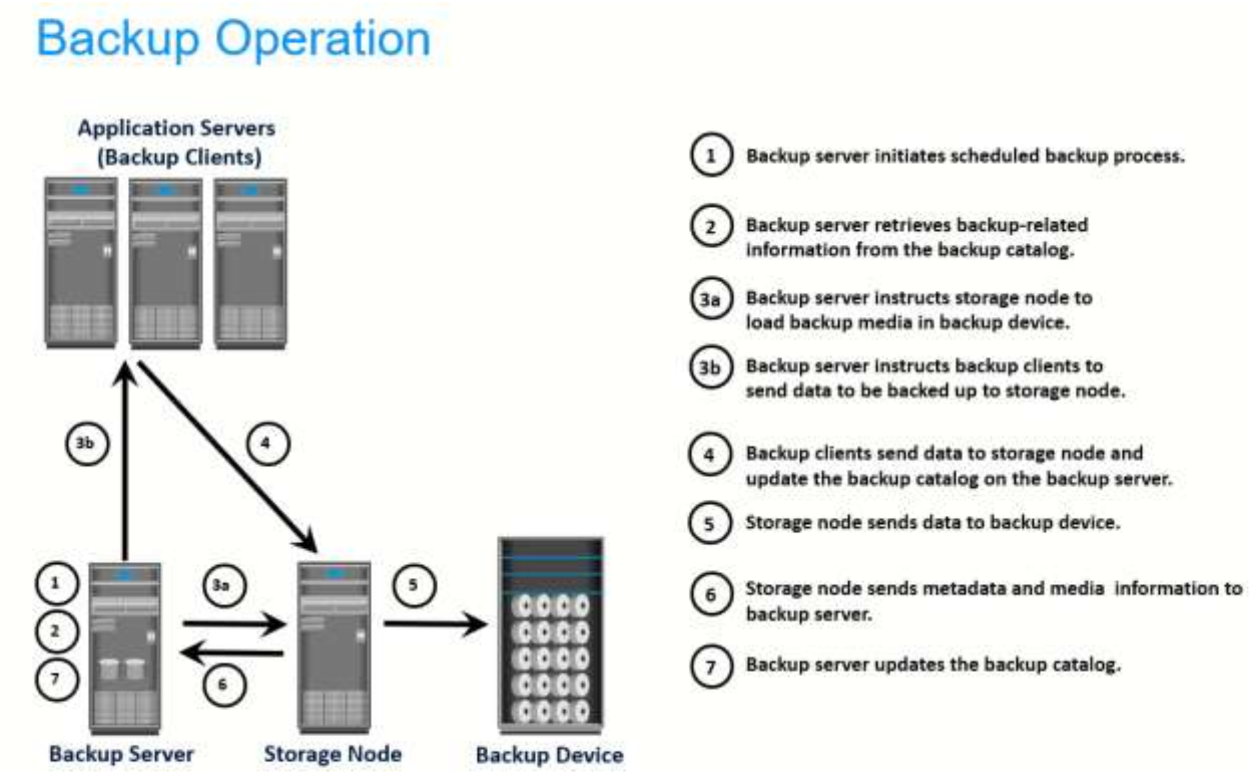

Backup Operation

1. Backup server initiates scheduled backup process.
2. Backup server retrieves backup-related information from the backup catalog.
3a. Backup server instructs storage node to load backup media in backup device.
3b. Backup server instructs backup clients to send data to be backed up to storage node.
4. Backup clients send data to storage node and update the backup catalog on the backup server.
5. Storage node sends data to backup device.
6. Storage node sends metadata and media information to backup server.
7. Backup server updates the backup catalog.

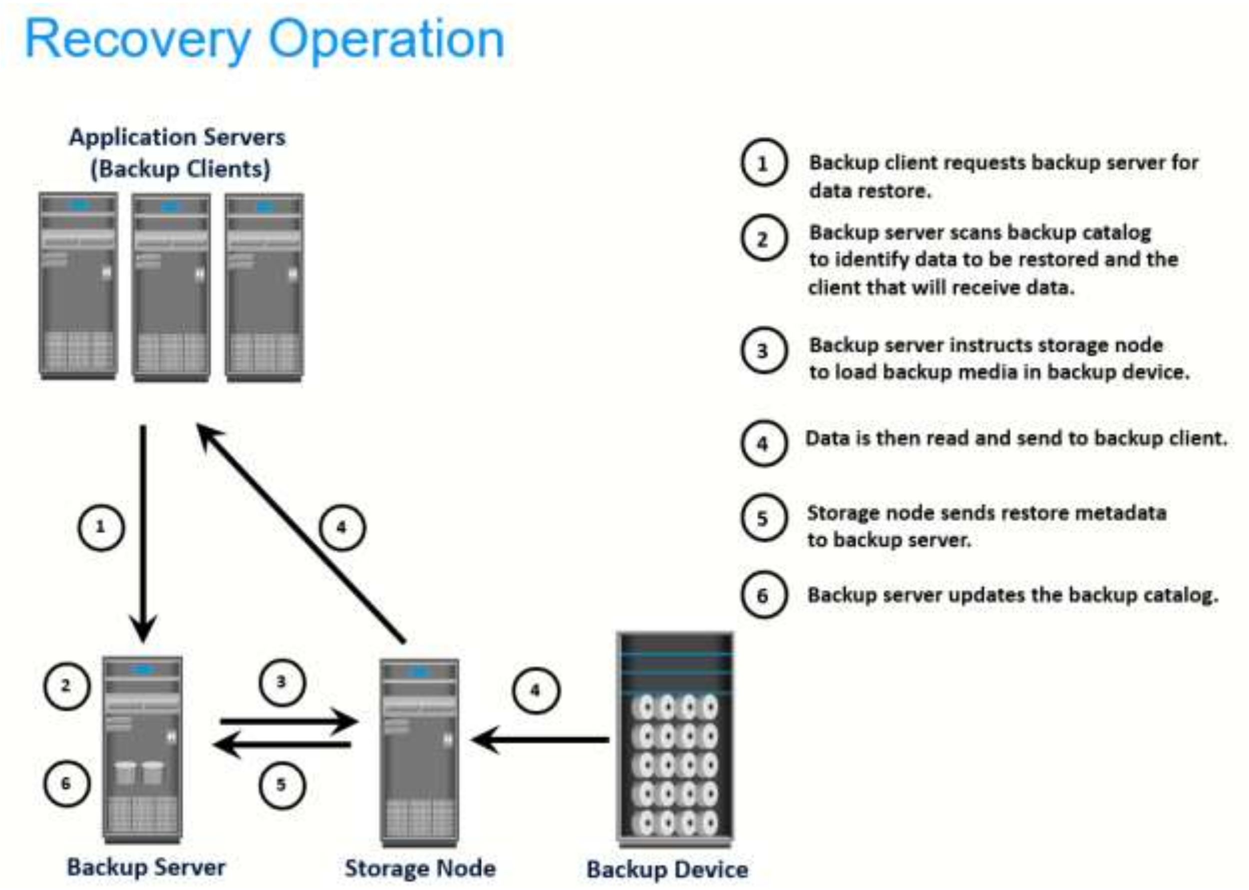

Recovery Operation

1. Backup client requests backup server for data restore.
2. Backup server scans backup catalog to identify data to be restored and the client that will receive data.
3. Backup server instructs storage node to load backup media in backup device.
4. Data is then read and send to backup client.
5. Storage node sends restore metadata to backup server.
6. Backup server updates the backup catalog.

8.a What is business continuity? Explain the BC terminology in detail.(10)

Solution:

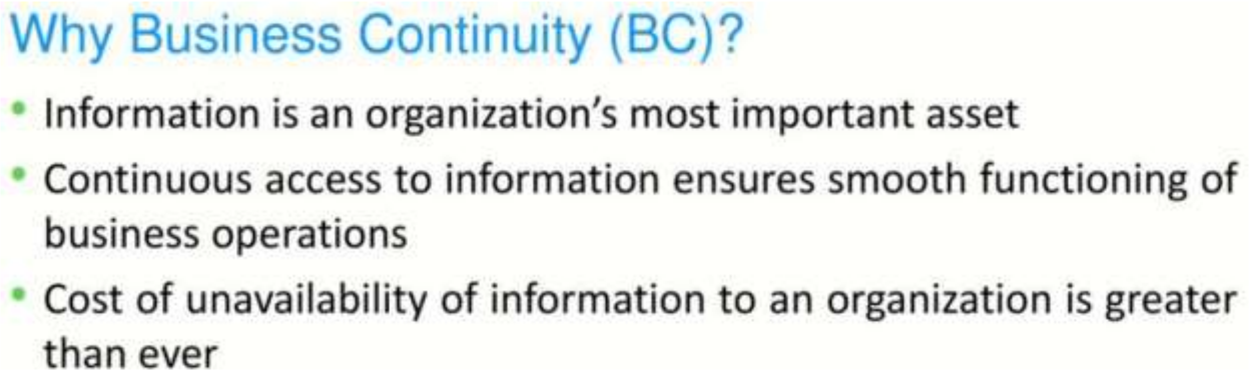

Why Business Continuity (BC)?

- Information is an organization's most important asset
- Continuous access to information ensures smooth functioning of business operations
- Cost of unavailability of information to an organization is greater than ever

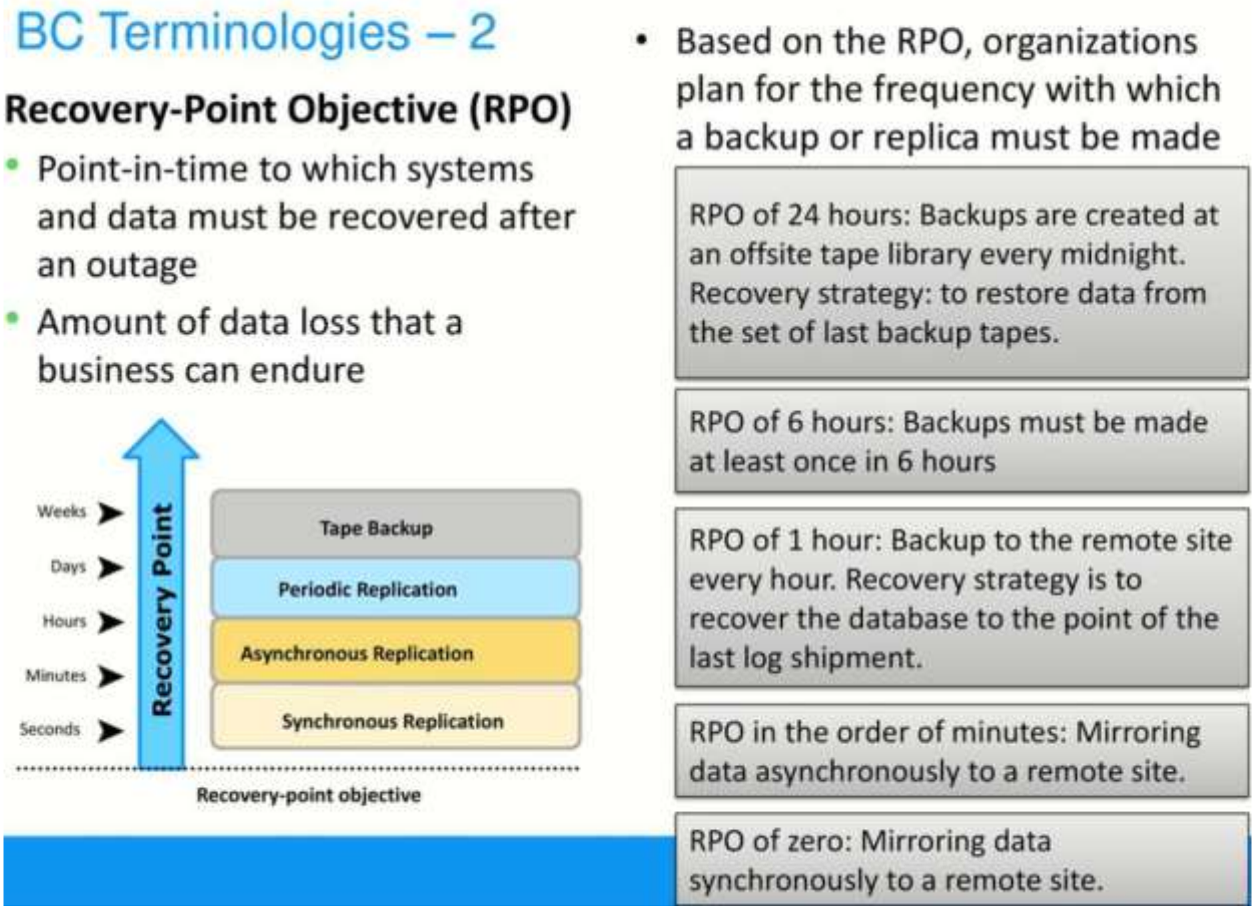

BC Terminologies – 2

Recovery-Point Objective (RPO)

- Point-in-time to which systems and data must be recovered after an outage
- Amount of data loss that a business can endure

- Based on the RPO, organizations plan for the frequency with which a backup or replica must be made

RPO of 24 hours: Backups are created at an offsite tape library every midnight. Recovery strategy: to restore data from the set of last backup tapes.

RPO of 6 hours: Backups must be made at least once in 6 hours

RPO of 1 hour: Backup to the remote site every hour. Recovery strategy is to recover the database to the point of the last log shipment.

RPO in the order of minutes: Mirroring data asynchronously to a remote site.

RPO of zero: Mirroring data synchronously to a remote site.

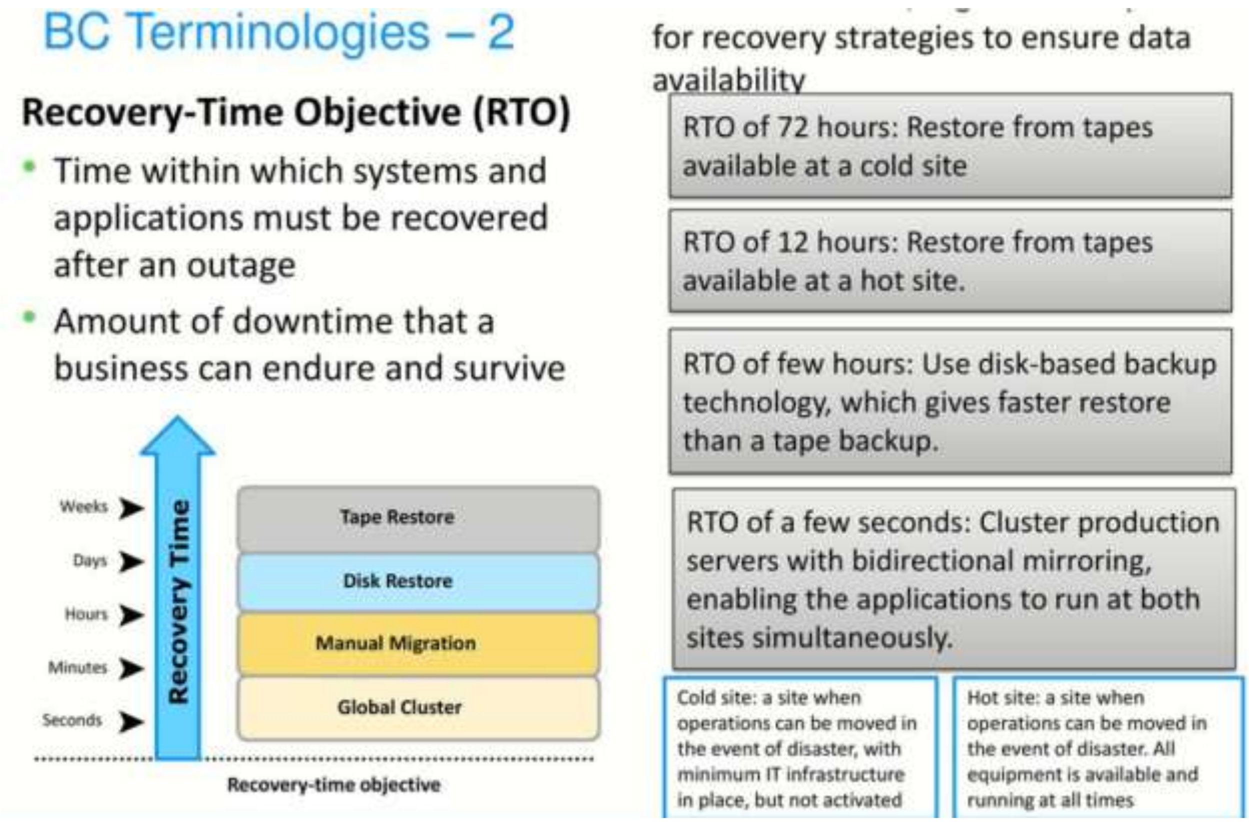

BC Terminologies – 2

Recovery-Time Objective (RTO)

- Time within which systems and applications must be recovered after an outage
- Amount of downtime that a business can endure and survive

for recovery strategies to ensure data availability

RTO of 72 hours: Restore from tapes available at a cold site

RTO of 12 hours: Restore from tapes available at a hot site.

RTO of few hours: Use disk-based backup technology, which gives faster restore than a tape backup.

RTO of a few seconds: Cluster production servers with bidirectional mirroring, enabling the applications to run at both sites simultaneously.

Cold site: a site when operations can be moved in the event of disaster, with minimum IT infrastructure in place, but not activated

Hot site: a site when operations can be moved in the event of disaster. All equipment is available and running at all times

8.b Explain different backup topologies(10)

Direct-Attached Backup

Back device is not shared

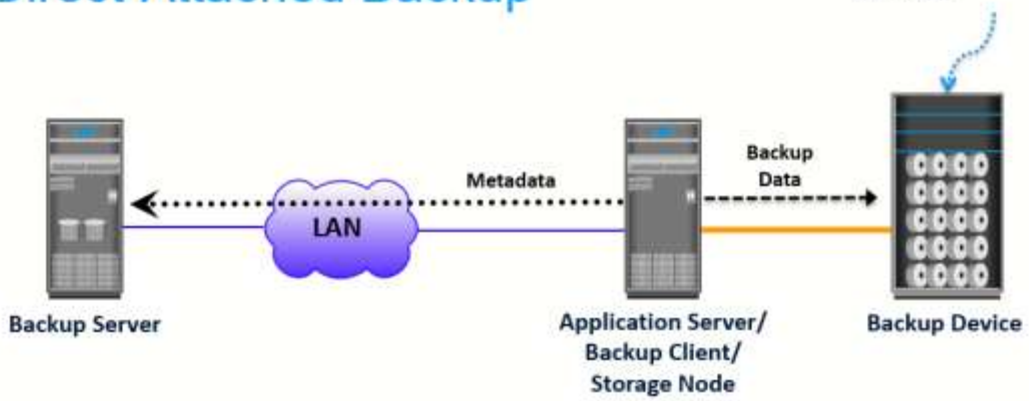

- This configuration frees the LAN from backup traffic.
- As the environment grows, however, there will be a need for central management of all backup devices and to share the resources to optimize costs.
- An appropriate solution is to share the backup devices among multiple servers.
- In this example, the client also acts as a storage node that writes data on the backup device.

LAN-based Backup

- The data to be backed up is transferred from the backup client (source), to the backup device (destination) over the LAN, which may affect network performance.
- Streaming across the LAN also affects network performance of all systems connected to the same segment as the backup server.
- This impact can be minimized by adopting a number of measures, such as configuring separate networks for backup and installing dedicated storage nodes for some application servers.

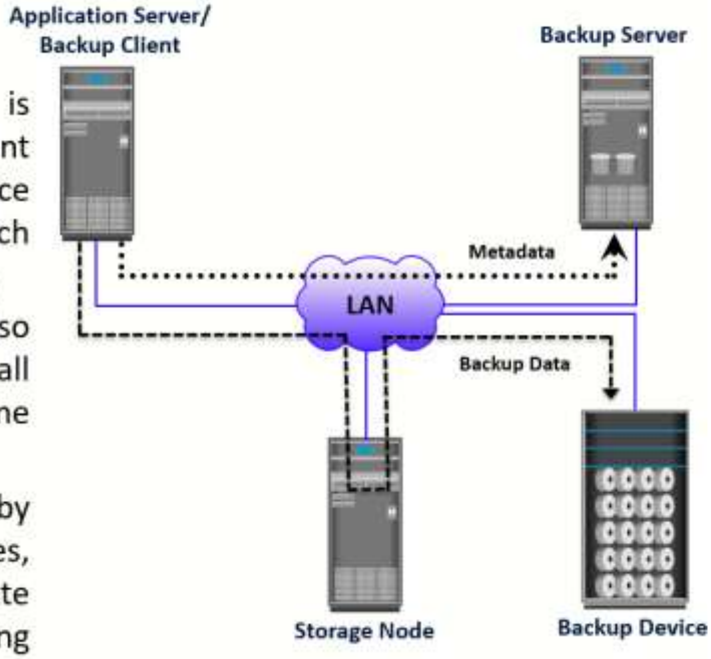

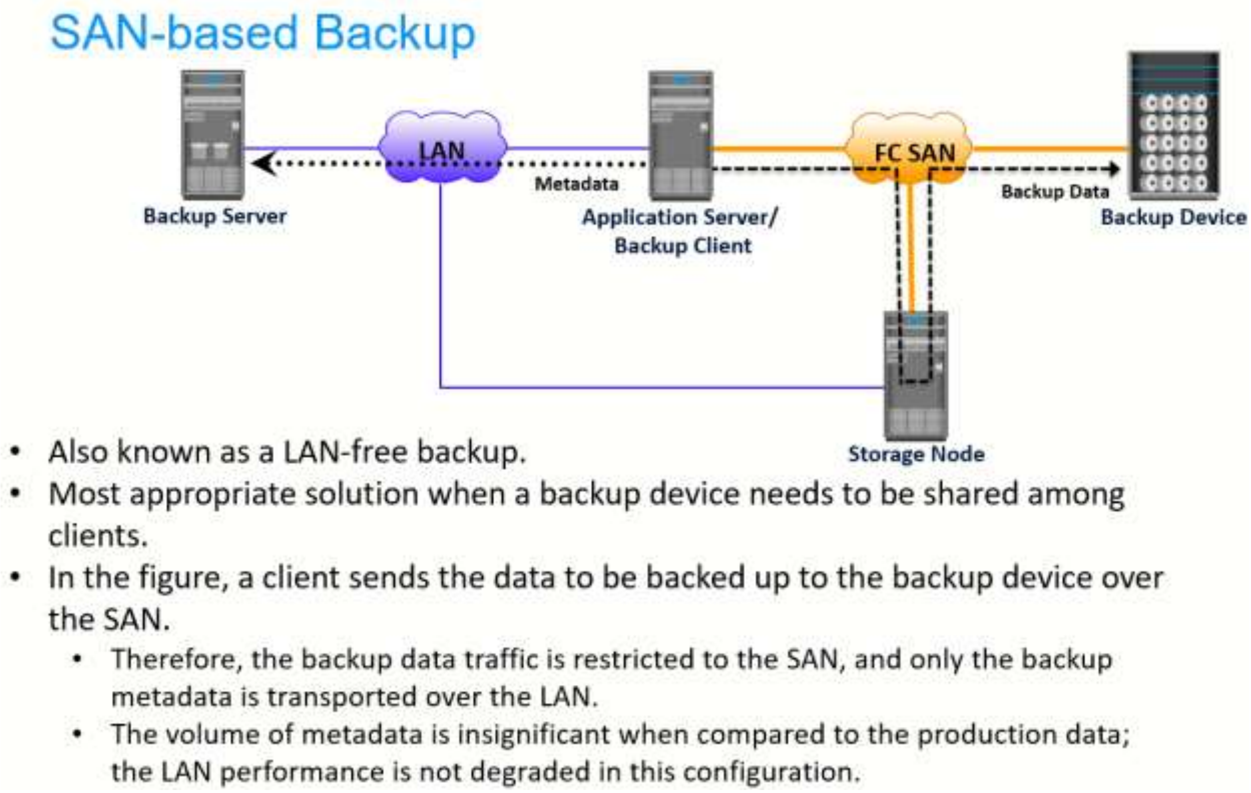

SAN-based Backup

- Also known as a LAN-free backup.
- Most appropriate solution when a backup device needs to be shared among clients.
- In the figure, a client sends the data to be backed up to the backup device over the SAN.
  - Therefore, the backup data traffic is restricted to the SAN, and only the backup metadata is transported over the LAN.
  - The volume of metadata is insignificant when compared to the production data; the LAN performance is not degraded in this configuration.

Module 5

9.a Explain FC-SAN security architecture and IP SAN security implementations with diagram (10)

FC SAN Security Architecture

Figure 14-5 illustrates various levels (zones) of a storage networking environment that must be secured and the security measures that can be deployed.

FC SANs not only suffer from certain risks and vulnerabilities that are unique, but also share common security problems associated with physical security and remote administrative access. In addition to implementing SAN-specific security measures, organizations must simultaneously leverage other security implementations in the enterprise. Table 14-1 provides a comprehensive list of protection strategies that must be implemented in various security zones.

Basic SAN Security Mechanisms

LUN masking and zoning, switch-wide and fabric-wide access control, RBAC, and logical partitioning of a fabric (Virtual SAN) are the most commonly used SAN security methods.

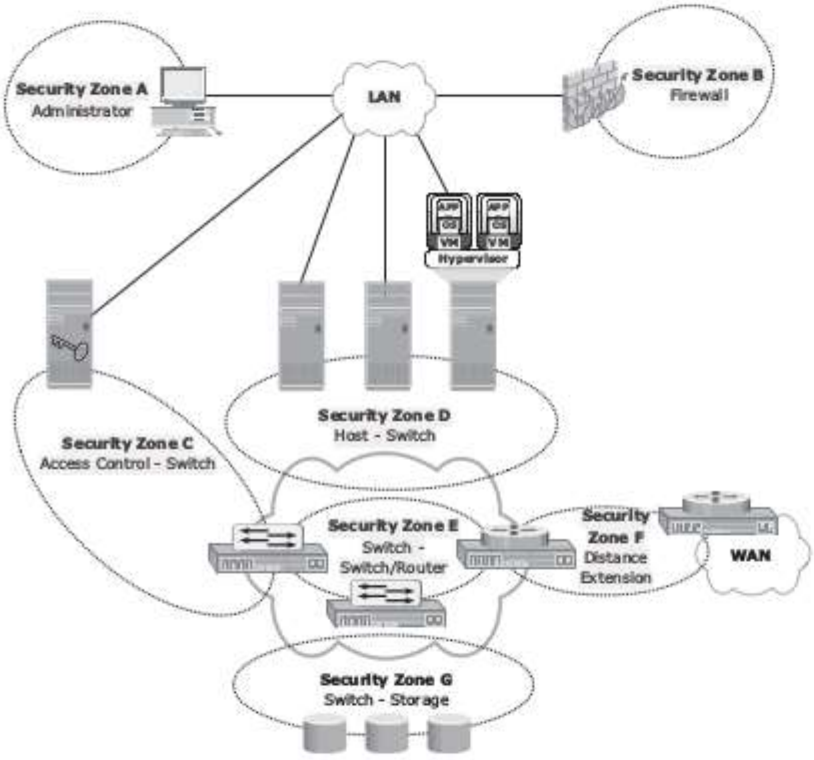

Figure 14-5: FC SAN security architecture

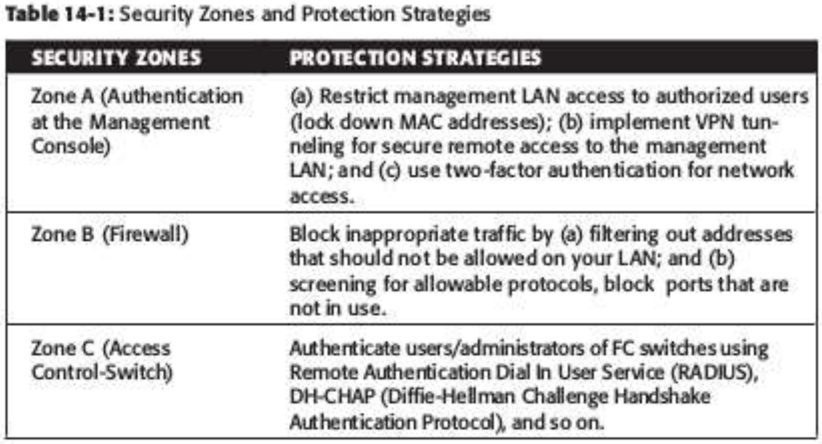

Table 14-1: Security Zones and Protection Strategies

SECURITY ZONES	PROTECTION STRATEGIES
Zone A (Authentication at the Management Console)	(a) Restrict management LAN access to authorized users (lock down MAC addresses); (b) implement VPN tunneling for secure remote access to the management LAN; and (c) use two-factor authentication for network access.
Zone B (Firewall)	Block inappropriate traffic by (a) filtering out addresses that should not be allowed on your LAN; and (b) screening for allowable protocols, block ports that are not in use.
Zone C (Access Control-Switch)	Authenticate users/administrators of FC switches using Remote Authentication Dial In User Service (RADIUS), DH-CHAP (Diffie-Hellman Challenge Handshake Authentication Protocol), and so on.

IP SAN

This section describes some of the basic security mechanisms used in IP SAN environments. The Challenge-Handshake Authentication Protocol (CHAP) is a basic authentication mechanism that has been widely adopted by network devices and hosts. CHAP provides a method for initiators and targets to authenticate each other by utilizing a secret code or password. CHAP secrets are usually random secrets of 12 to 128 characters. Figure 14-10 depicts the CHAP authentication process.

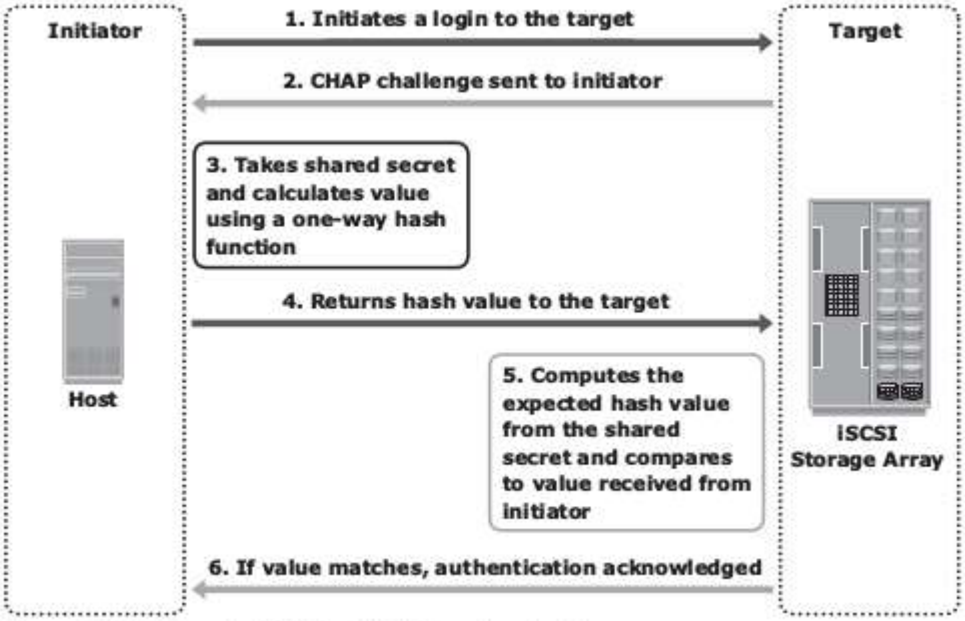
Initiator
Target
1. Initiates a login to the target
2. CHAP challenge sent to initiator
3. Takes shared secret and calculates value using a one-way hash function
Host
4. Returns hash value to the target
iSCSI Storage Array
5. Computes the expected hash value from the shared secret and compares to value received from initiator
6. If value matches, authentication acknowledged

Figure 14-10: Securing IPSAN with CHAP authentication

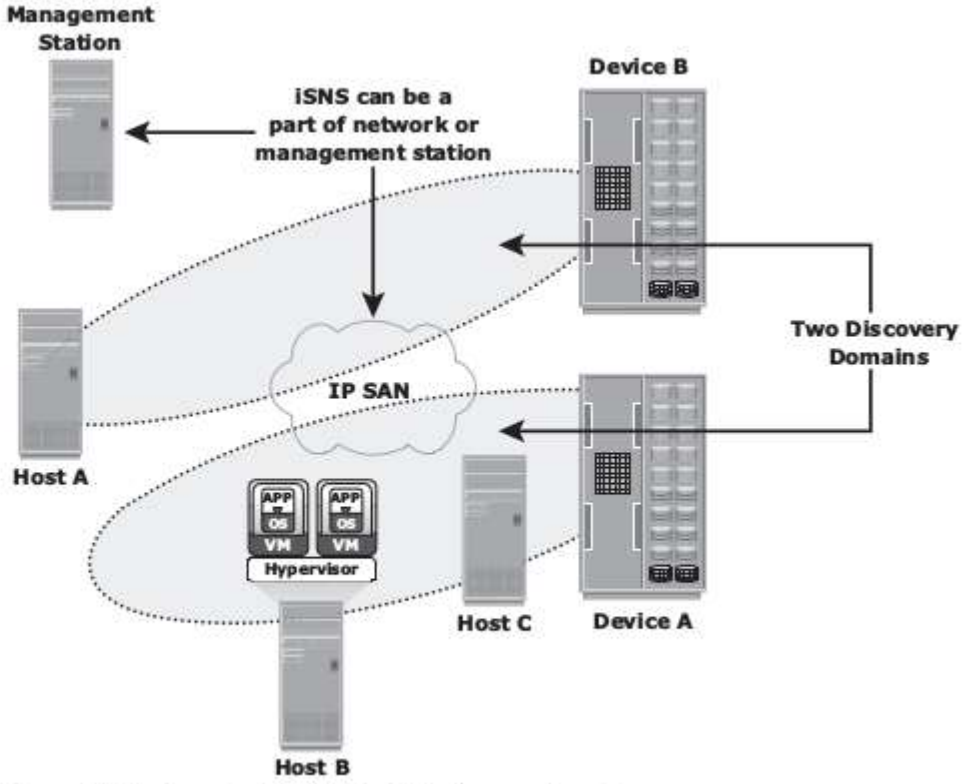

Figure 14-11: Securing IPSAN with iSNS discovery domains

9.b Explain the following(10)

(i) uses of local replicas

(ii) LVM based replications

(iii) Full Volume Mirrorring

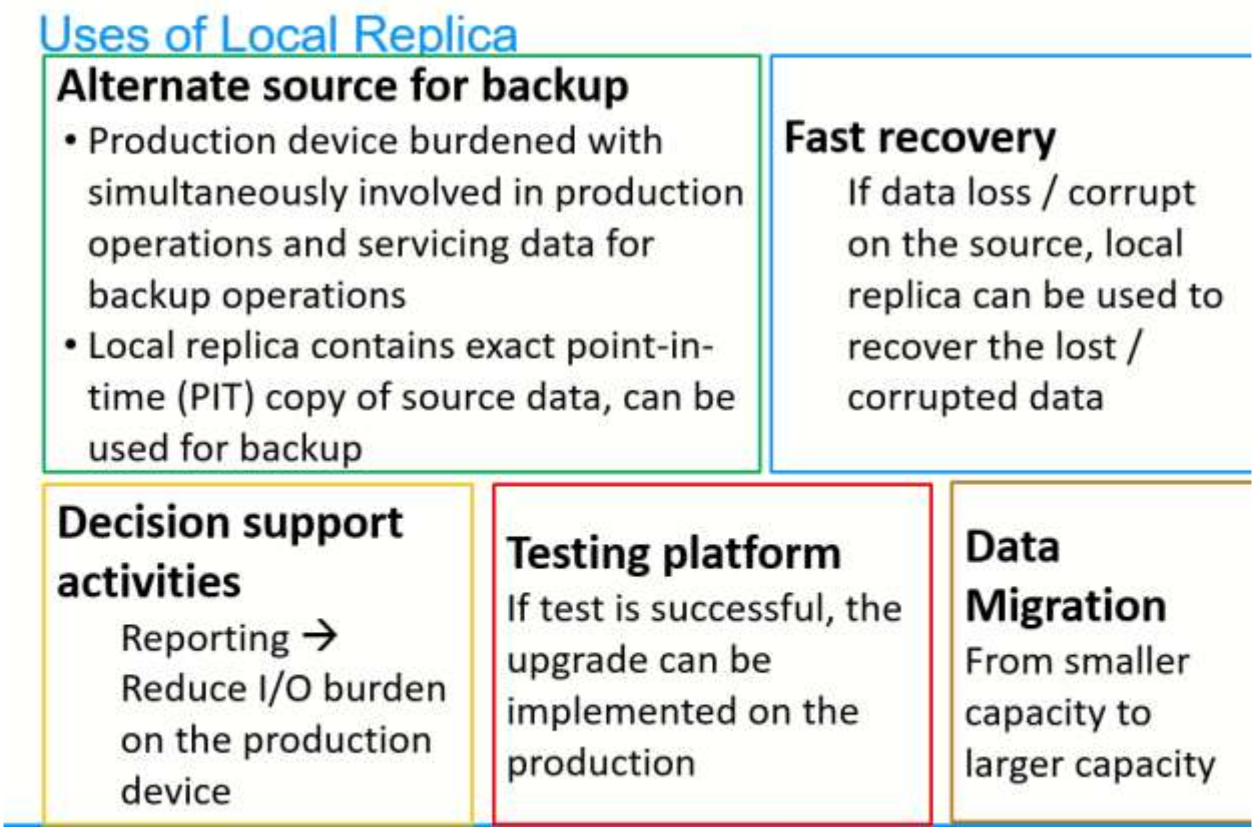

Uses of Local Replica

Alternate source for backup

- Production device burdened with simultaneously involved in production operations and servicing data for backup operations
- Local replica contains exact point-in-time (PIT) copy of source data, can be used for backup

Fast recovery

If data loss / corrupt on the source, local replica can be used to recover the lost / corrupted data

Decision support activities

Reporting → Reduce I/O burden on the production device

Testing platform

If test is successful, the upgrade can be implemented on the production

Data Migration

From smaller capacity to larger capacity

(ii) **LVM based replications**

- LVM based preplicas **add overhead** on host CPUs
  - Each write is translated into two writes on the disk
  - Can degrade application performance
- If host volumes are already **storage array LUNs** then the added **redundancy** provided by LVM mirroring is unnecessary
  - The devices will have some RAID protection already
- Both replica and source are stored within the **same volume** group
  - Replica **cannot be accessed** by another host
  - If server fails, both source and replica would be unavailable
- Keeping track of changes on the mirrors is a **challenge**

(iii) Full Volume Mirroring

- Target is a full physical copy of the source device
- Target is attached to the source and data from source is copied to the target
- Target is unavailable while it is attached

- Target device is as large as the source device
- Good for full backup, decision support, development, testing and restore to last PIT

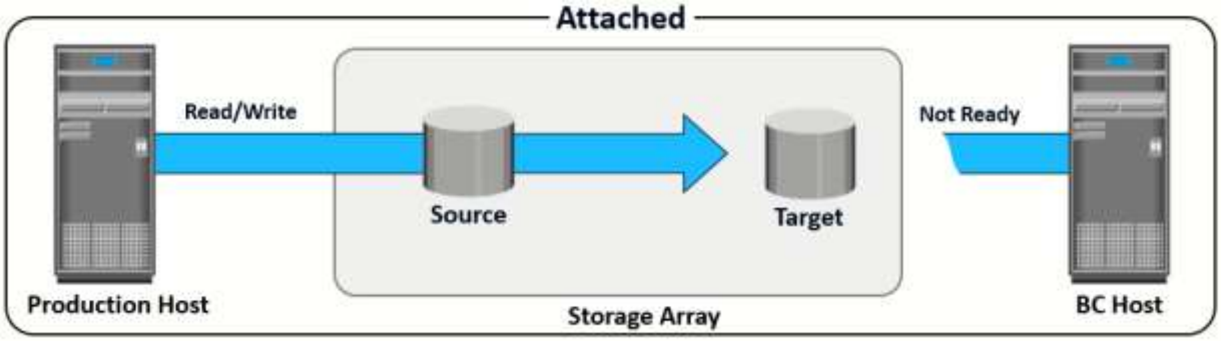

10.a Explain Storage security domains

Storage Security Domains

Storage devices connected to a network raise the risk level and are more exposed to security threats via networks. However, with increasing use of networking in storage environments, storage devices are becoming highly exposed to security threats from a variety of sources. Specific controls must be implemented to secure a storage networking environment.

To identify the threats that apply to a storage network, access paths to data storage can be categorized into three security domains: application access, management access, and backup, replication, and archive. Figure 14-1 depicts the three security domains of a storage system environment.

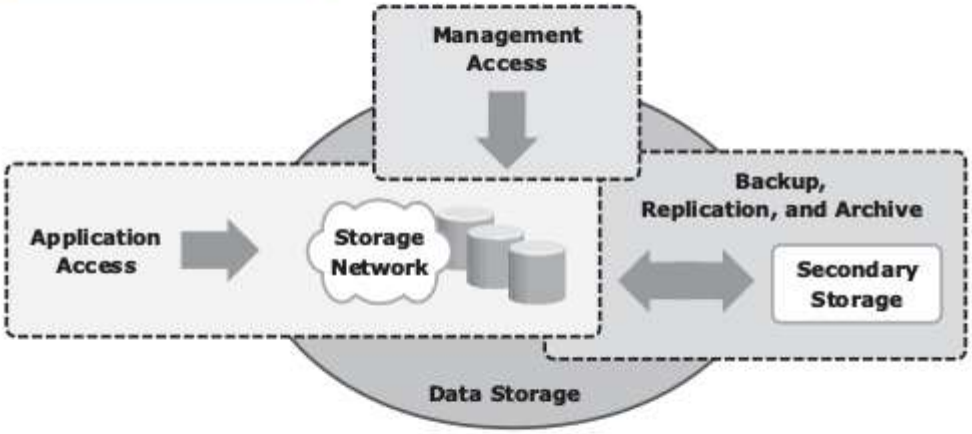

Figure 14-1: Storage security domains

To secure the storage networking environment, identify the existing threats within each of the security domains and classify the threats based on the type of security services — availability, confidentiality, integrity, and accountability. The next step is to select and implement various controls as countermeasures to the threats.

Securing the Application Access Domain

The application access domain may include only those applications that access the data through the file system or a database interface.

An important step to secure the application access domain is to identify the threats in the environment and appropriate controls that should be applied. Implementing physical security is also an important consideration to prevent media theft.

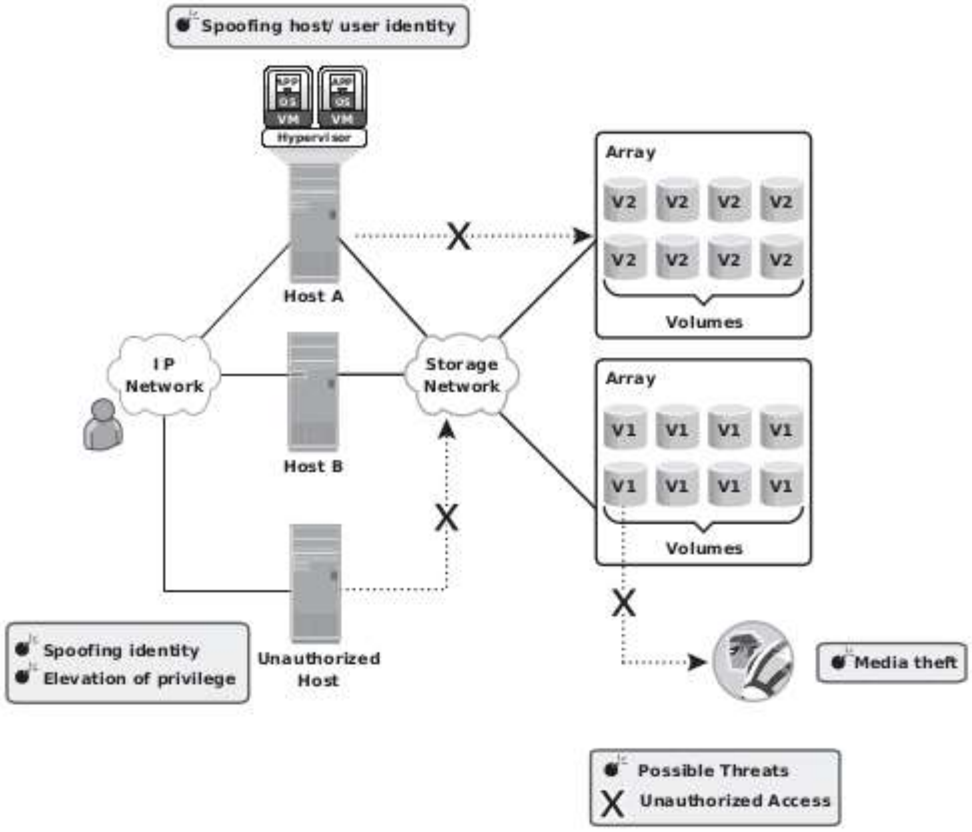

Figure 14-2: Security threats in an application access domain

10.b What is remote replication? Explain storage array based replication.(10)

- Process of creating replicas at remote sites

storage array based replication

- Replication is performed by array-operating environment

- Three replication methods: synchronous, asynchronous, and disk buffered
- **Synchronous**
  - Writes are committed to both source and replica before it is acknowledged to host
- **Asynchronous**
  - Writes are committed to source and immediately acknowledged to host
  - Data is buffered at source and transmitted to remote site later

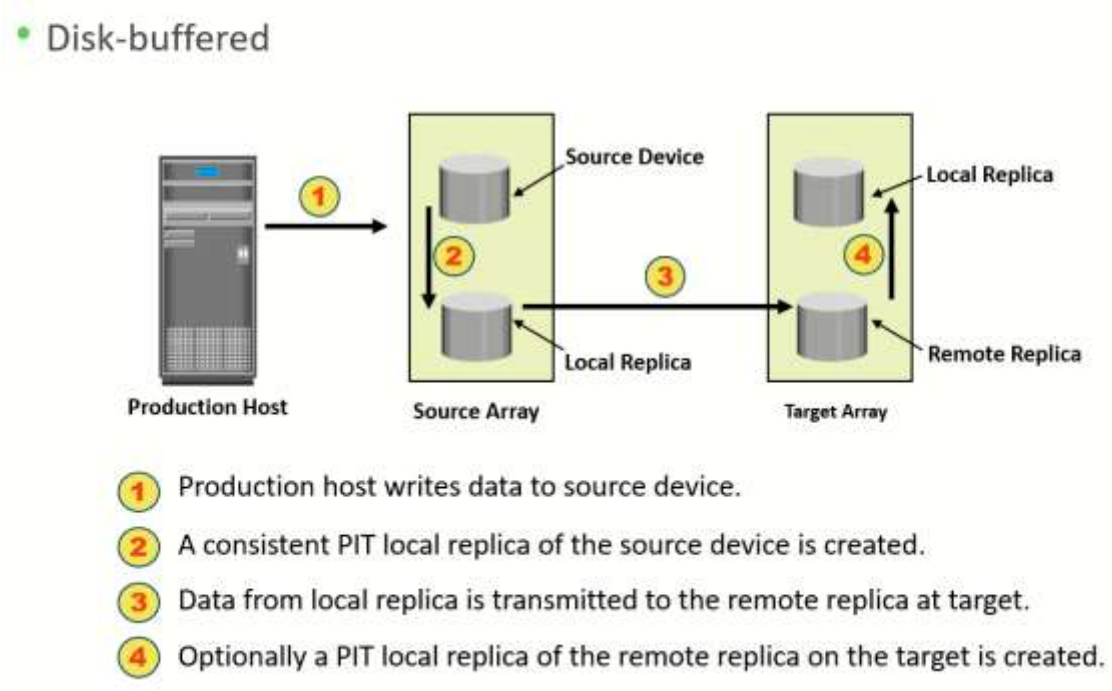

- Disk-buffered

1. Production host writes data to source device.
2. A consistent PIT local replica of the source device is created.
3. Data from local replica is transmitted to the remote replica at target.
4. Optionally a PIT local replica of the remote replica on the target is created.