CBCS SCHEME 18CS822

Eighth Semester B.E. Degree Examination, July/August 2022

Storage Area Networks Time 3 hrs. Max. Marks: 100

Note: Answer *any FIVE full questions, choosing ONE full question from each module.*

Module-1

1 a. Explain the core elements of data centre along with its key characteristics. (10 Marks)

Core elements of a data center infrastructure:

Five core elements are essential for the basic functionality of a data center:

- Application: An application is a computer program that provides the logic for computing operations. Applications, such as an order processing system, can be layered on a database, which in turn uses operating system services to perform read/write operations to storage devices.
- Database: More commonly, a database management system (DBMS) provides a structured way to store data in logically organized tables that are interrelated. A DBMS optimizes the storage and retrieval of data.
- Server and operating system: A computing platform that runs applications and databases.
- Network: A data path that facilitates communication between clients and servers or between servers and storage.
- Storage array: A device that stores data persistently for subsequent use. These core elements are typically viewed and managed as separate entities, but all the elements must work together to address data processing requirements.

Key Requirements for Data Center Elements

- Availability: All data center elements should be designed to ensure accessibility for ever.
- Security: Polices, procedures, and proper integration of the data center core elements that will prevent unauthorized access to information must be established.
- Scalability: Data center operations should be able to allocate additional processing capabilities or storage on demand, without interrupting business operations. Business growth often requires deploying more servers, new applications, and additional databases.
- Performance: All the core elements of the data center should be able to provide optimal performance and service all processing requests at high speed. The

infrastructure should be able to support performance requirements.

· Data integrity: Data integrity refers to mechanisms such as error correction codes or parity bits which ensure that data is written to disk exactly as it was received

· Capacity: Data center operations require adequate resources to store and process large amounts of data efficiently. When capacity requirements increase, the data center must be able to provide additional capacity without interrupting availability, or, at the very least, with minimal disruption. Capacity may be managed by reallocation of existing resources, rather than by adding new resources.

· Manageability: A data center should perform all operations and activities in the most efficient manner. Manageability can be achieved through automation and the reduction of human (manual) intervention in common tasks.

b. Discuss the process of host access to storage. (06 Marks)

Host

Users store and retrieve data through applications. The computers on which these applications run are referred to as hosts. Hosts can range from simple laptops to complex clusters of servers.

Physical Components

A host has three key physical components:

■■ Central processing unit (CPU)
■■ Storage, such as internal memory and disk devices
■■ Input/Output (I/O) devices

CPU

The CPU consists of four main components:

■■Arithmetic Logic Unit (ALU): This is the fundamental building block of the CPU. It performs arithmetical and logical operations such as addition, subtraction, and Boolean functions (AND, OR, and NOT).
■■ Control Unit: A digital circuit that controls CPU operations and coordinates the functionality of the CPU.
■■ Register: A collection of high-speed storage locations. The registers store intermediate data that is required by the CPU to execute an instruction and provide fast access because of their proximity to the ALU. CPUs typically have a small number of registers.

I/O Devices

I/O devices enable sending and receiving data to and from a host. This communication may be one of the following types:

■■User to host communications: Handled by basic I/O devices, such as the keyboard, mouse, and monitor. These devices enable users to enter data and view the results of operations.
■■Host to host communications: Enabled using devices such as a Network Interface Card (NIC) or modem.
■■Host to storage device communications: Handled by a Host Bus Adaptor (HBA).

c. Write a short note on evolution of storage architecture. (04 Marks)

Redundant Array of Independent Disks (RAID): This technology was developed to address the cost, performance, and availability requirements of data. It continues to evolve today and is used in all storage architectures such as DAS, SAN, and so on.
Direct-attached storage (DAS): This type of storage connects directly to a server (host) or a group of servers in a cluster. Storage can be either internal or external to the server. External DAS alleviated the challenges of limited internal storage capacity.
Storage area network (SAN): This is a dedicated, high-performance Fibre Channel (FC) network to facilitate block-level communication between servers and storage. Storage is partitioned and assigned to a server for accessing its data. SAN offers scalability, availability, performance, and cost benefits compared to DAS.

Network-attached storage (NAS): This is dedicated storage for file serving applications. Unlike a SAN, it connects to an existing communication network (LAN) and provides file access to heterogeneous clients. Because it is purposely built for providing storage to file server applications, it offers higher scalability, availability, performance, and cost benefits compared to general purpose file servers.
Internet Protocol SAN (IP-SAN): One of the latest evolutions in storage architecture, IP-SAN is a convergence of technologies used in SAN and NAS. IP-SAN provides block-level communication across a local or wide area network (LAN or WAN), resulting in greater consolidation
and availability of data.

OR
2 a. Explain in detail disk drive components with suitable diagram.
(10 Marks)
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.

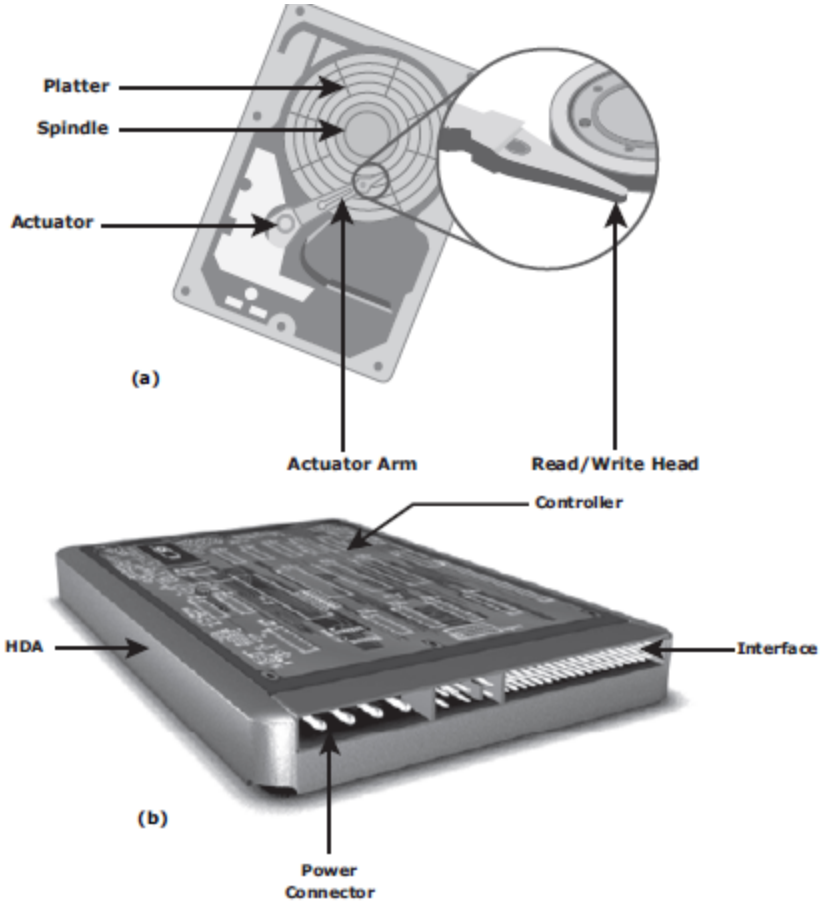

Platter
A typical HDD consists of one or more flat circular disks called platters
Spindle
A spindle connects all the platters, as shown in the figure, and is connected to a motor. The motor of the spindle rotates with a constant speed.
Read/Write Head
Read/Write (R/W) heads, shown in Figure 2-4, read and write data from or to a platter. Drives have two R/W heads per platter, one for each surface of the platter.
Actuator Arm Assembly
The R/W heads are mounted on the actuator arm assembly

b. Discuss the concept of DAS with advantages and disadvantages. (06 Marks)

DAS requires a relatively lower initial investment than storage networking. Storage networking

architectures are discussed later in this book. DAS configuration is simple and can be deployed easily and rapidly. Setup is managed using host-based tools, such as the host OS, which makes storage management tasks easy for small and medium enterprises. DAS is the simplest solution when compared to other storage networking models and requires fewer management tasks, and less hardware and software elements to set up and operate.

However, DAS does not scale well. A storage device has a limited number of ports, which restricts the number of hosts that can directly connect to the storage. A limited bandwidth in DAS restricts the available I/O processing capability.

When capacities are being reached, the service availability may be compromised, and this has a ripple effect on the performance of all hosts attached to that specific device or array. The distance limitations associated with implementing DAS because of direct connectivity requirements can be addressed by using Fibre Channel connectivity. DAS does not make optimal use of resources due to its limited ability to share front end ports. In DAS environments, unused resources cannot be easily re-allocated, resulting in islands of over-utilized and under-utilized storage pools.

c. Explain disk partitioning and concatenation. (04 Marks)

HDDs. In partitioning, an HDD is divided into logical containers called logical volumes (LVs). For example, a large physical drive can be partitioned into multiple LVs to maintain data according to the file system's and applications' requirements. The partitions are created from groups of contiguous cylinders when the hard disk is initially set up on the host. The host's file system accesses the partitions without any knowledge of partitioning and the physical structure of the disk.

Concatenation is the process of grouping several smaller physical drives and presenting them to the host as one logical drive

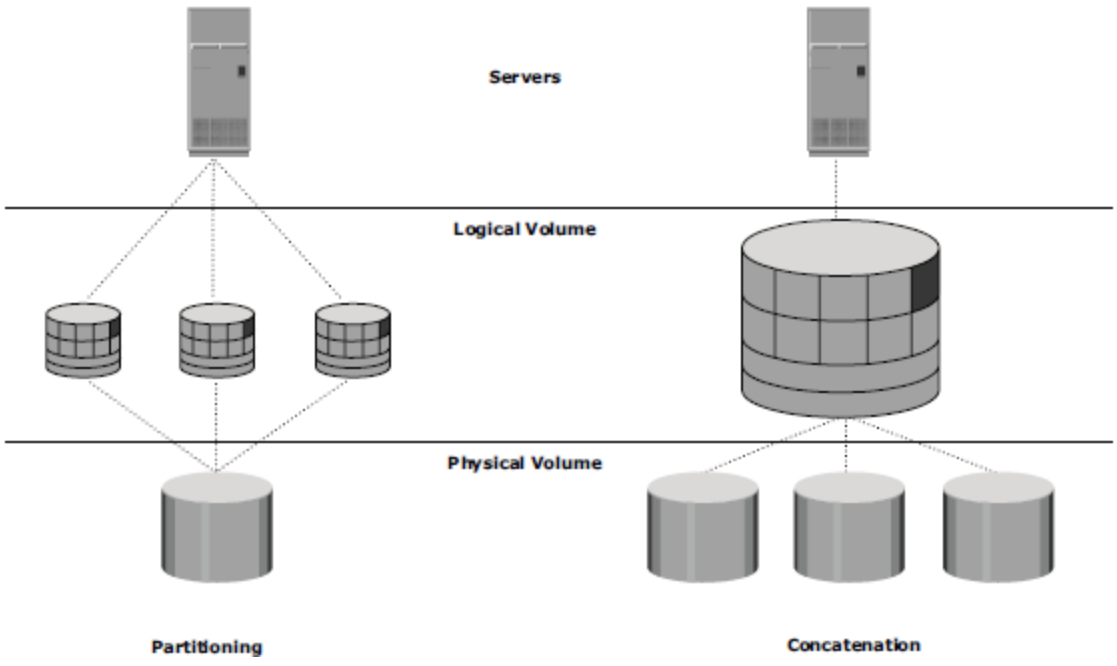

Module 2

3 a. What is RAID? List different RAD levels where parity technique has been adopted. Explain nested RAID RAID 3, RAID 5 with diagram. (10 Marks)

Raid Levels

1. RAID 0 - Striped array with no fault tolerance
2. RAID 1 - Disk mirroring
3. RAID 3 - Parallel access array with dedicated parity disk
4. RAID 4 - Striped array with independent disks and a dedicated parity disk
5. RAID 5 - Striped array with independent disks and distributed parity
6. RAID 6 - Striped array with independent disks and dual distributed parity
7. Nested Combinations of RAID levels - Example: RAID 1 + RAID 0

i)Striping

A RAID set is a group of disks. Within each disk, a predefined number of contiguously addressable disk blocks are defined as *strips*. The set of aligned strips that spans across all the disks within the RAID set is called a *stripe*. Figure: Striped RAID set

- *Strip size* (also called *stripe depth*) describes the number of blocks in a *strip*, and is the maximum amount of data that can be written to or read from a single HDD in the set before the next HDD is accessed, assuming that the accessed data starts at the beginning of the strip. Note that all strips in a stripe have the same number of blocks, and decreasing strip size means that data is broken into smaller pieces when spread across the disks.
- Stripe size is a multiple of strip size by the number of HDDs in the RAID set.
- *Stripe width* refers to the number of data strips in a stripe.
- Striped RAID does not protect data unless parity or mirroring is used.
- However, striping may significantly improve I/O performance.

Depending on the type of RAID implementation, the RAID controller can be configured to access data across multiple HDDs simultaneously.

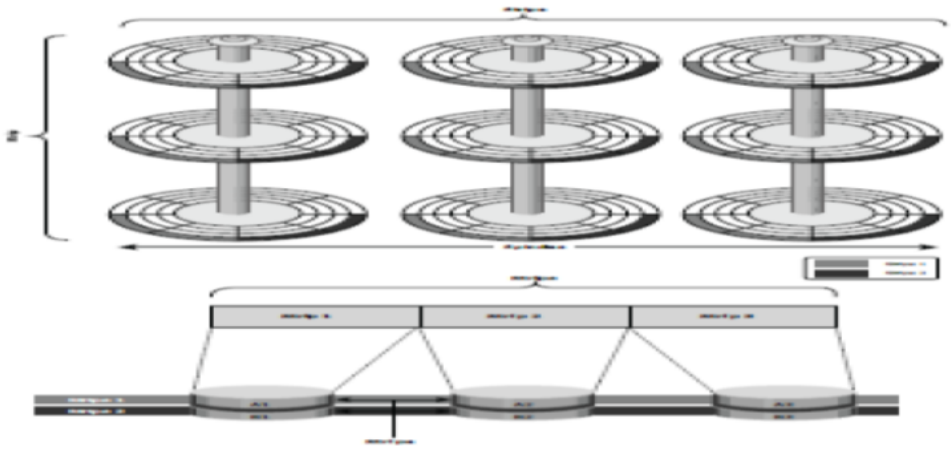

Figure: Striped RAID set

ii) Mirroring

· *Mirroring* is a technique whereby data is stored on two different HDDs, yielding two copies of data.

· In the event of one HDD failure, the data is intact on the surviving HDD and the controller continues to service the host's data requests from the surviving disk of a mirrored pair.

· When the failed disk is replaced with a new disk, the controller copies the data from the surviving disk of the mirrored pair.

· This activity is transparent to the host.

· In addition to providing complete data redundancy, mirroring enables faster recovery from disk failure.

· However, disk mirroring provides only data protection and is not a substitute for data backup. Mirroring constantly captures changes in the data, whereas a backup captures point-in-time images of data.

· Mirroring involves duplication of data — the amount of storage capacity needed is twice the amount of data being stored.

· Therefore, mirroring is considered expensive and is preferred for mission-critical applications that cannot afford data loss.

· Mirroring improves read performance because read requests can be serviced by both disks. However, write performance deteriorates, as each write request

manifests as two writes on the HDDs.

· In other words, mirroring does not deliver the same levels of write performance as a striped RAID.

iii Parity

· *Parity* is a method of protecting striped data from HDD failure without the cost of mirroring.

· An additional HDD is added to the stripe width to hold parity, a mathematical construct that allows re-creation of the missing data.

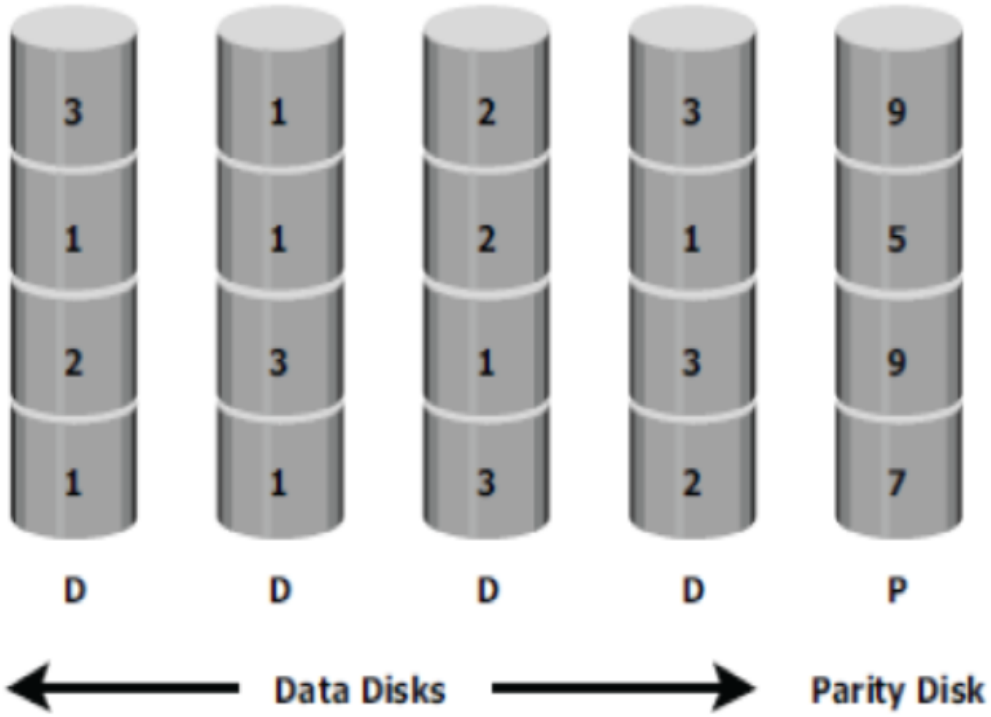

Figure: Parity RAID

· Parity is a redundancy check that ensures full protection of data without maintaining a full set of duplicate data.

· Parity information can be stored on separate, dedicated HDDs or distributed across all the drives in a RAID set.

· The first four disks, labeled *D,* contain the data.

· The fifth disk, labeled *P*, stores the parity information, which in this case is the sum of the elements in each row.

· Now, if one of the *D*s fails, the missing value can be calculated by subtracting the sum of the rest of the elements from the parity value.

· Parity calculation is a *bitwise XOR* operation.

· Calculation of parity is a function of the RAID controller.

· Compared to mirroring, parity implementation considerably reduces the cost associated with data protection.

· Consider a RAID configuration with five disks.

· Four of these disks hold data, and the fifth holds parity information.

· Parity requires 25 percent extra disk space compared to mirroring, which requires 100 percent extra disk space.

b.Write a short note on: i) Node ports ii) Cable and connectors. (06 Marks)

Node Ports

In fibre channel, devices such as hosts, storage and tape libraries are all referred to as nodes. Each node is a source or destination of information for one or more nodes. Each node requires one or more ports to provide a physical interface for communicating with other nodes.

Cabling

SAN implementations use optical fiber cabling. Copper can be used for shorter distances for back-end connectivity, as it provides a better signal-to-noise ratio for distances up to 30 meters. Optical fiber cables carry data in the form of light. There are two types of optical cables, multi-mode and single-mode

Interconnect Devices

Hubs, switches, and directors are the interconnect devices commonly used in SAN.

Hubs are used as communication devices in FC-AL implementations. Hubs physically connect nodes in a logical loop or a physical star topology. All the nodes must share the bandwidth because data travels through all the connection points. Because of availability of low cost and high performance switches, hubs are no longer used in 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, resulting in bandwidth aggregation.

C. Discuss RAID impact on disk performance. (04 Marks)

When choosing a RAID type, it is imperative to consider the impact to disk

performance and application IOPS.

In both mirrored and parity RAID configurations, every write operation translates into more I/O overhead for the disks which is referred to as write penalty. In a RAID 1 implementation, every write operation must be performed on two disks configured as a mirrored pair while in a RAID 5 implementation, a write operation may manifest as four I/O operations. When performing small I/Os to a disk configured with RAID 5, the controller has to read, calculate, and write a parity segment for every data write operation.

RAID 5 that contains a group of five disks. Four of these disks are used for data and one is used forparity.

The parity (P) at the controller is calculated as follows:

Ep = E1 + E2 + E3 + E4 (XOR operations)

Here, D1 to D4 is striped data across the RAID group of five disks.

Whenever the controller performs a write I/O, parity must be computed by reading the old parity (Ep old) and the old data (E4 old) from the disk, which means two read I/Os. The new parity (Ep new) is computed as follows:

Ep new = Ep old – E4 old + E4 new (XOR operations)

4 a. Explain structure of cache with operations. (10 Marks)

b. List and explain RAID techniques. (06 Marks)
Striping
A RAID set is a group of disks. Within each disk, a predefined number of contiguously

addressable disk blocks are defined as strips. The set of aligned strips that spans across all the disks within the RAID set is called a stripe

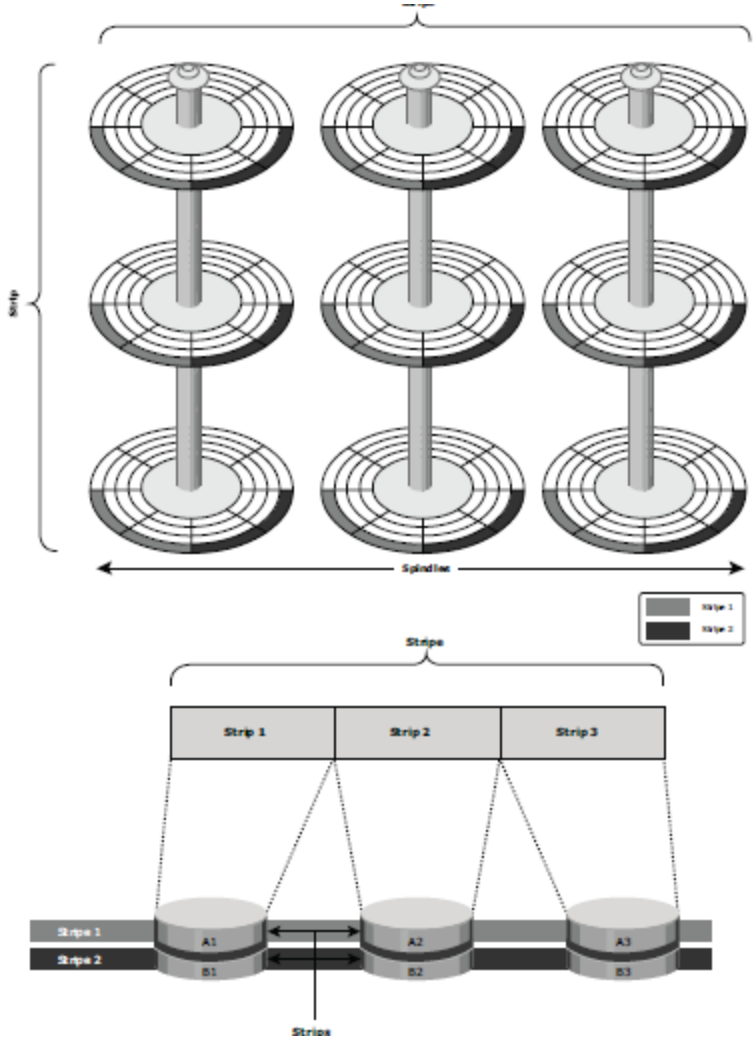

Mirroring

Mirroring is a technique whereby data is stored on two different HDDs, yielding two copies of data. In the event of one HDD failure, the data is intact on the surviving HDD and the controller continues to service the host's data requests from the surviving disk of a mirrored pair.

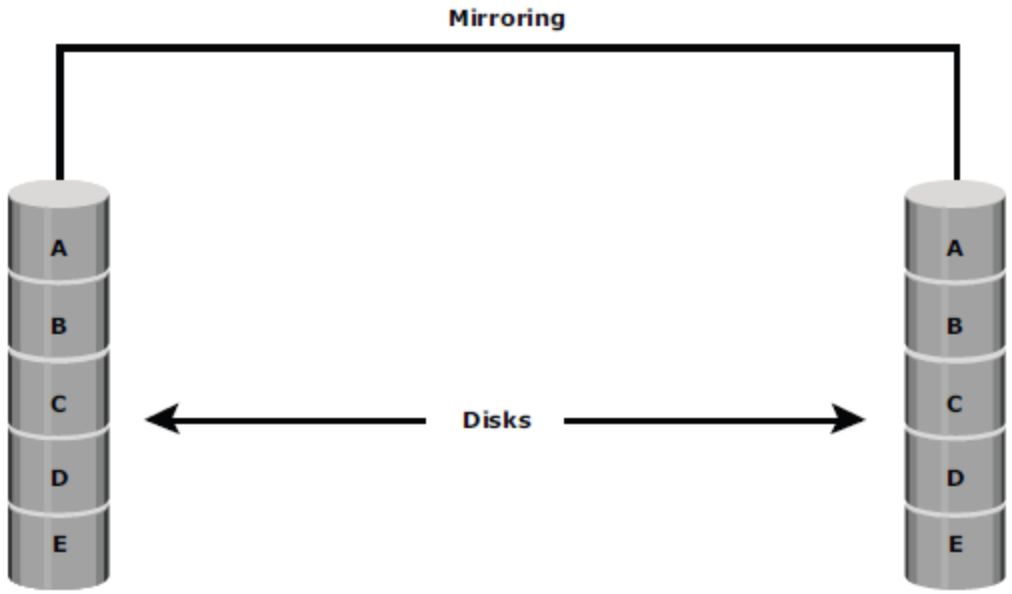

Parity

Parity is a method of protecting striped data from HDD failure without the cost of mirroring. An additional HDD is added to the stripe width to hold parity, a mathematical construct that allows re-creation of the missing data. Parity is a redundancy check that ensures full protection of data without maintaining a full set of duplicate data

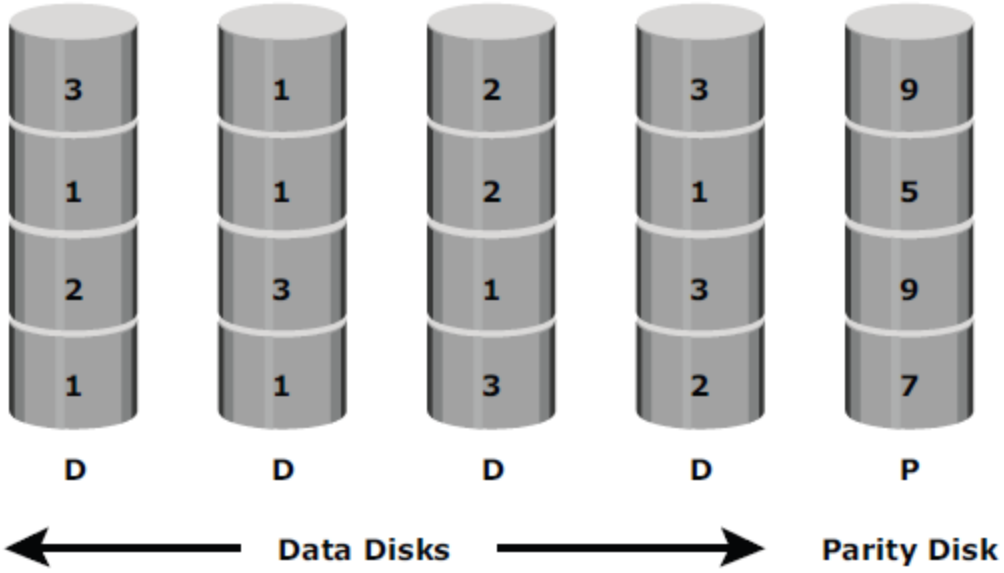

C. List types of intelligent storage systems and explain any one in detail. (04 Marks)

Module 3

5 a. Write a note on iSCSI and its topologies. (08 Marks)

iSCSI is the host-based encapsulation of SCSI I/O over IP using an Ethernet NIC card or an iSCSI HBA in the host. As illustrated in Figure 8-2 (a), IP traffic is routed over a network either to a gateway device that extracts the SCSI I/O from the IP packets or to an iSCSI storage array. The gateway can then send the SCSI I/O to an FC-based external storage array, whereas an iSCSI storage array can handle the extraction and I/O natively.

Native iSCSI Connectivity

If an iSCSI-enabled array is deployed, FC components are not needed for iSCSI connectivity in the native topology. In the example shown in Figure, the array has one or more Ethernet NICs that are connected to a standard Ethernet switch and configured with an IP address and listening port.

Bridged iSCSI Connectivity

A bridged iSCSI implementation includes FC components in its configuration. Figure illustrates an existing FC storage array used to service hosts connected through iSCSI

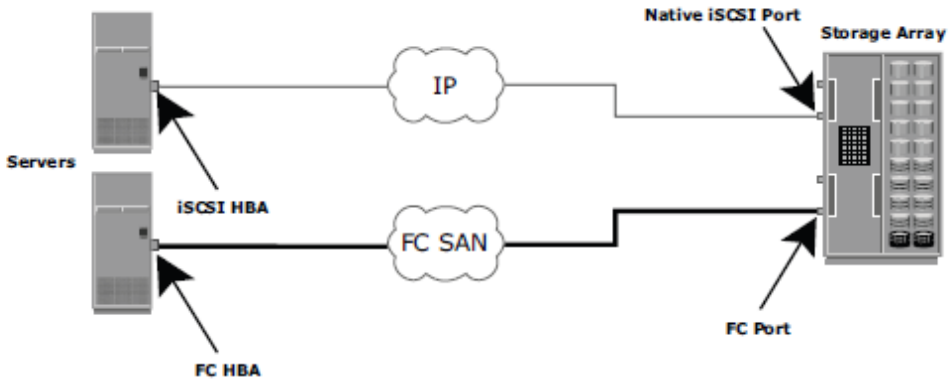

(a) Native iSCSI Connectivity

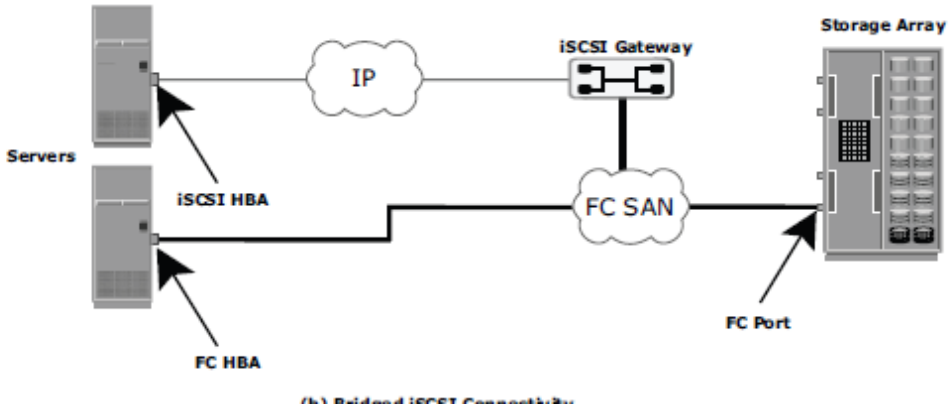

b. Explain the advantages of NAS. (08 Marks)

Block-Level Storage Virtualization

Block-level storage virtualization provides a translation layer in the SAN, between the hosts and the storage arrays, Instead of being directed to the LUNs on the individual storage arrays, the hosts are directed to the virtualized LUNs on the virtualization device. The virtualization device translates

File-Level Virtualization

File-level virtualization addresses the NAS challenges by eliminating the dependencies between the data accessed at the file level and the location where the files are physically stored. This provides opportunities to optimize storage utilization and server consolidation and to perform nondisruptive file migrations.

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.

C. Compare CIFS and NFS protocols. (04 Marks)

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.

NFS is a client/server protocol for file sharing that is most commonly used on UNIX systems. NFS was originally based on the connectionless User Datagram Protocol (UDP). It uses a machine-independent model to represent user data. It also uses Remote Procedure Call (RPC) as a method of interprocess communication between two computers. 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

OR

6 a. Explain fibre channel protocol stack with neat diagram and write a short note on its

performance and security. (08 Marks)

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.

FC-4 Upper Layer Protocol

- FC-4 is the uppermost layer in the FCP stack.
- This layer defines the application interfaces and the way Upper Layer Protocols (ULPs) are mapped to the lower FC layers.
- The FC standard defines several protocols that can operate on the FC-4 layer.
- Some of the protocols include SCSI, HIPPI Framing Protocol, Enterprise Storage Connectivity (ESCON), ATM, and IP.

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.

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.

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.

b. Explain NAS components with diagram. (06 Marks)

c. With a neat diagram explain Gateway network attached storage connectivity. (06 Marks)

Module-4

7 a. Discuss the life cycle of BC Planning. (10 Marks)

BC planning must follow a disciplined approach like any other planning process. Organizations today dedicate specialized resources to develop and maintain BC plans. From the conceptualization to the realization of the BC plan, a lifecycle of activities can be defined for the BC process. The BC planning lifecycle includes five stages

1. Establishing objectives
2. Analyzing
3. Designing and developing
4. Implementing
5. Training, testing, assessing, and maintaining

Several activities are performed at each stage of the BC planning lifecycle including the following key activities:

1. Establishing objectives

Determine BC requirements.

Estimate the scope and budget to achieve requirements.

Select a BC team by considering subject matter experts from all areas of the business, whether internal or external.

Create BC policies.

2. Analyzing

Collect information on data profiles, business processes, infrastructure support, dependencies, and frequency of using business infrastructure.

Identify critical business needs and assign recovery priorities.

Create a risk analysis for critical areas and mitigation strategies.

Conduct a Business Impact Analysis (BIA).

Create a cost and benefit analysis based on the consequences of data unavailability.

Evaluate options.

3. Designing and developing

Define the team structure and assign individual roles and responsibilities.

Design data protection strategies and develop infrastructure.

Develop contingency scenarios.

Develop emergency response procedures.

Detail recovery and restart procedures.

4. Implementing

Implement risk management and mitigation procedures that include backup, replication, and management of resources.

Prepare the disaster recovery sites that can be utilized if a disaster affects the primary data center.

Implement redundancy for every resource in a data center to avoid single points of failure.

5. Training, testing, assessing, and maintaining

Train the employees who are responsible for backup and replication of business-critical data on a regular basis or whenever there is a modification in the BC plan.

Train employees on emergency response procedures when disasters are declared.

Train the recovery team on recovery procedures based on contingency scenarios.

b. List some important BC technology solutions. Explain the failure analysis in BC. (10 Marks)
Backup and recovery: Backup to tape is the predominant method of ensuring data availability. These days, low-cost, high-capacity disks are used for backup, which considerably speeds up the backup and recovery process. The frequency of backup is determined based on RPO, RTO, and the frequency of data changes.
■■ Storage array-based replication (local): Data can be replicated to a separate location within the same storage array. The replica is used independently for BC operations. Replicas can also be used for restoring operations if data corruption occurs.
■■ Storage array-based replication (remote): Data in a storage array can be replicated to another storage array located at a remote site. If the storage array is lost due to a disaster, BC operations start from the remote storage array.
■■Host-based replication: The application software or the LVM ensures that a copy of the data managed by them is maintained either locally or at a remote site for recovery purposes.

Failure analysis in BC

- Single Point of Failure
- Fault Tolerance
- Multipathing Software

OR
8 a. Describe backup and restore operation. (10 Marks)
When a backup process is initiated, significant network communication takes place between the different components of a backup infrastructure.
The backup server initiates the backup process for different clients based on the backup schedule configured for them.
For example, the backup process for a group of clients may be scheduled to start at 3:00 am every day.
The backup server coordinates the backup process with all the components in a backup configuration.
The backup server maintains the information about backup clients to be contacted and storage nodes to be used in a backup operation.

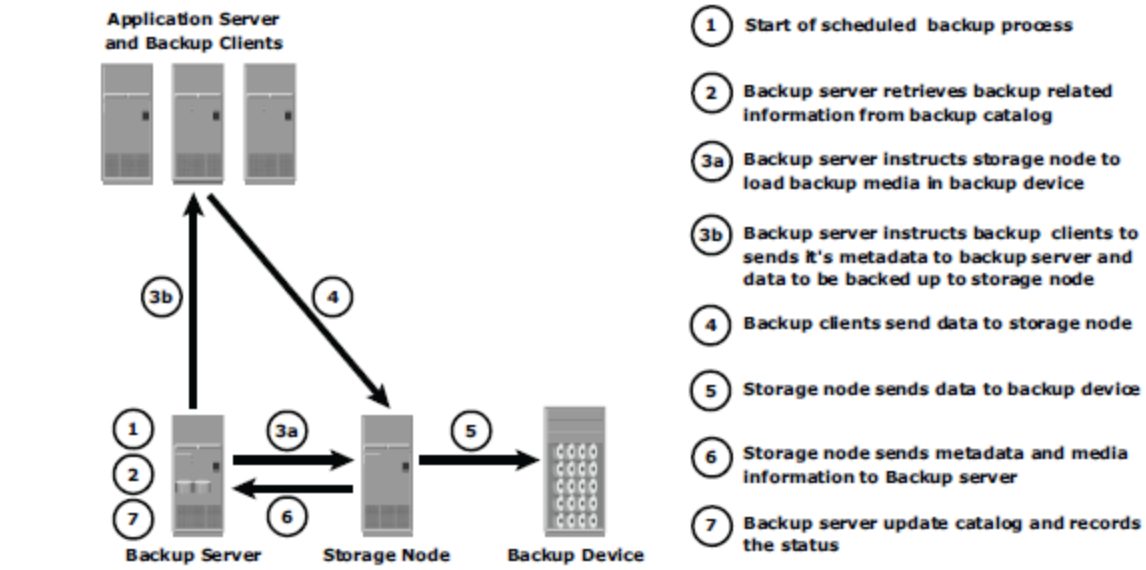

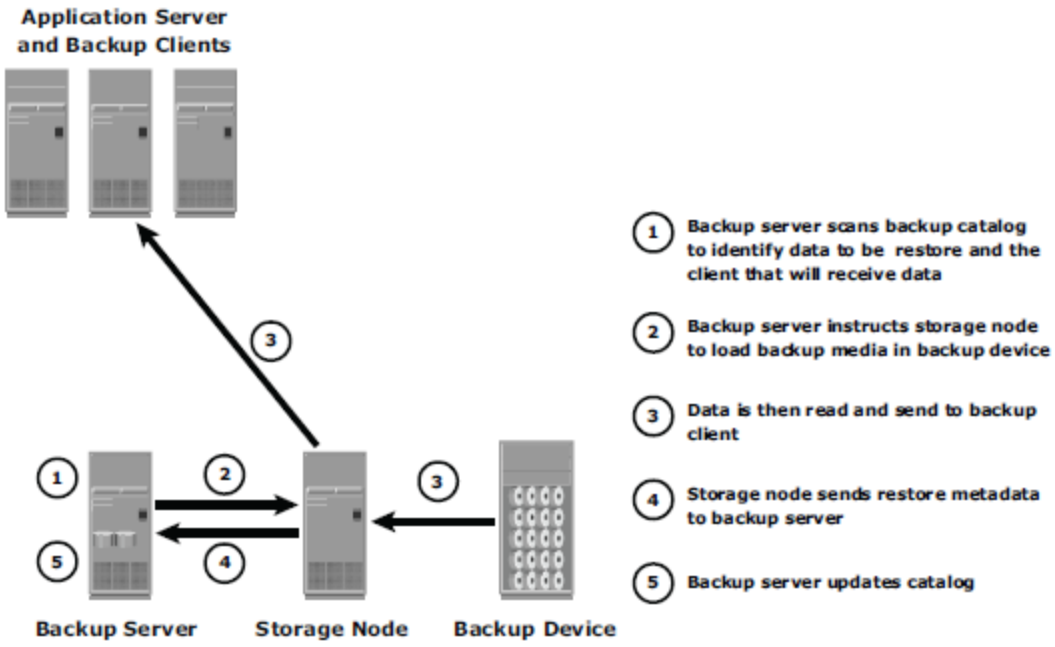

b. Write a short note on: i) Backup architecture

A backup system uses client/server architecture with a backup server and multiple backup clients.
The backup server manages the backup operations and maintains the backup catalog, which contains information about the backup process and backup metadata.
The backup server depends on backup clients to gather the data to be backed up.

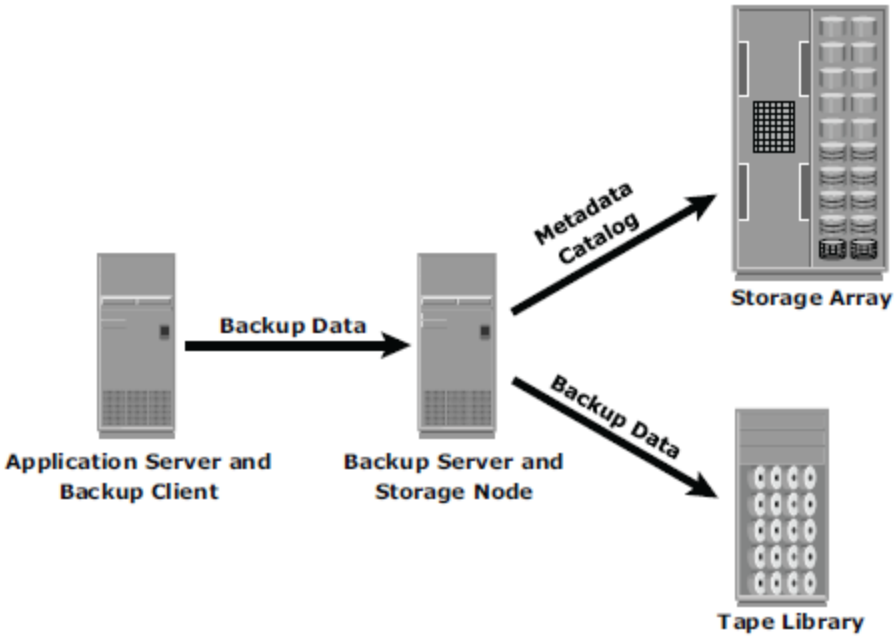

ii) Backup purpose. (10 Marks)

Backup Purpose

Backups are performed to serve three purposes: disaster recovery, operational backup, and archival.

i) Disaster Recovery

Backups can be performed to address disaster recovery needs.

The backup copies are used for restoring data at an alternate site when the primary site is incapacitated due to a disaster.

Based on RPO and RTO requirements, organizations use different backup strategies for disaster recovery.

When a tape-based backup method is used as a disaster recovery strategy, the backup tape media is shipped and stored at an offsite location.

These tapes can be recalled for restoration at the disaster recovery site.

Organizations with stringent RPO and RTO requirements use remote replication technology to replicate data to a disaster recovery site.

This allows organizations to bring up production systems online in a relatively short period of time in the event of a disaster

ii)Operational Backup

Data in the production environment changes with every business transaction and operation.

Operational backup is a backup of data at a point in time and is used to restore data in the event of data loss or logical corruptions that may occur during routine processing.

Archival

Backups are also performed to address archival requirements.

Although CAS has emerged as the primary solution for archives, traditional backups are still

used by small and medium enterprises for long-term preservation of transaction records, e-mail messages, and other business records required for regulatory compliance.
Apart from addressing disaster recovery, archival, and operational

Module-5

9 a. Mention major local replication technologies. Explain network based local replication. (10 Marks)

Alternate source for backup: Under normal backup operations, data is read from the production volumes (LUNs) and written to the backup device.

Decision-support activities such as reporting: Running the reports using the data on the replicas greatly reduces the I/O burden placed on the production device.

Testing platform: A local replica can be used for testing critical business data or applications. For example, when planning an application upgrade, it can be tested using the local replica. If the test is successful, it can be restored to the source volumes.

Data migration: Local replication can also be used for data migration.

Data migration may be performed for various reasons, such as migrating from a small LUN to a larger LUN.

In storage array-based local replication, the array operating environment performs the local replication process. The host resources such as CPU and memory are not used in the replication process.

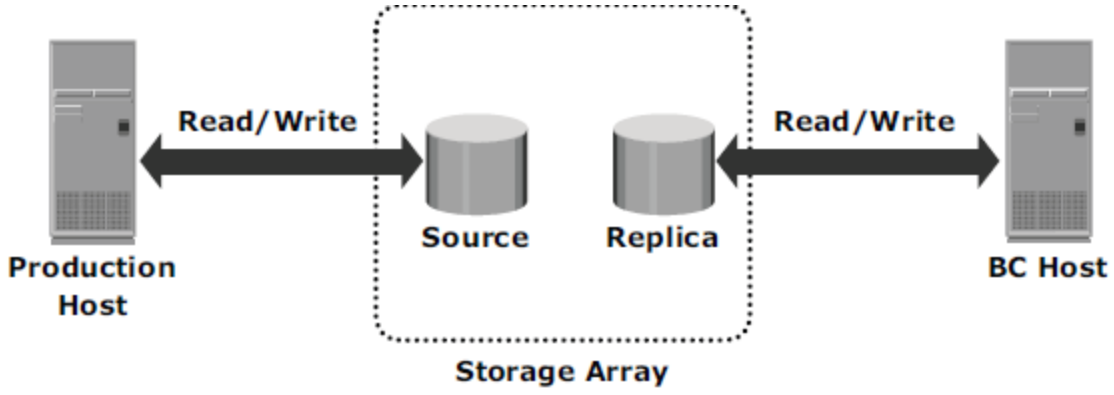

b. Discuss flushing the file system buffer. *(*05 Marks)

Flushing The process of committing data from cache to disk.

Force flushing In case of a large I/O burst, this process forcibly flushes dirty pages onto the disk.

Double buffering The buffering of data in two places. For example RDBMs use their own buffering along with file system buffering.

c. Explain the uses of local Replicas. (05 Marks)

Alternate source for backup: Under normal backup operations, data is read from the production volumes (LUNs) and written to the backup device.

Decision-support activities such as reporting: Running the reports using the data on the replicas greatly reduces the I/O burden placed on the production device.

Testing platform: A local replica can be used for testing critical business data or applications. For example, when planning an application upgrade, it can be tested using the local replica. If the test is successful, it can be restored to the source volumes.

Data migration: Local replication can also be used for data migration.

Data migration may be performed for various reasons, such as migrating from a small LUN to a larger LUN.

In storage array-based local replication, the array operating environment performs the local replication process. The host resources such as CPU and memory are not used in the replication process.

OR

10 a. Write a short pote on array based synchronous remote replication. (06 Marks)

The two basic modes of remote replication are synchronous and asynchronous.

In synchronous remote replication, writes must be committed to the source and the target, prior to acknowledging “write complete” to the host.

Additional writes on the source cannot occur until each preceding write has been completed and acknowledged. This ensures that data is identical on the source and the replica at all times.

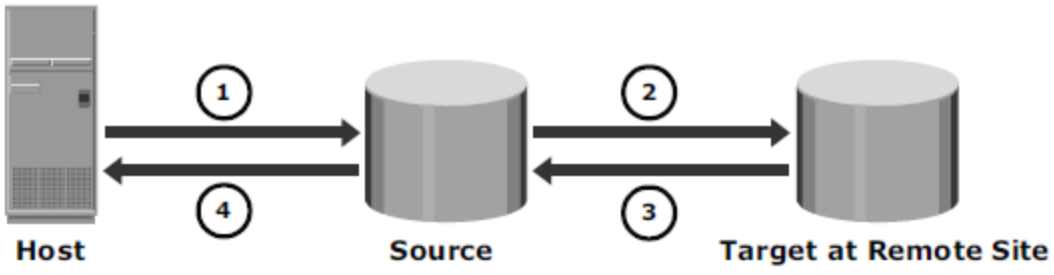

1. Host writes data to source
2. Data from source is replicated to target at remote site
3. Target acknowledges back to source
4. Source acknowledges write complete to host

b. Explain security threats in backup, replication and archive environment. (06 Marks)
Threats are the potential attacks that can be carried out on an IT infrastructure.
These attacks can be classified as active or passive. Passive attacks are attempts to gain unauthorized access into the system. They pose threats to confidentiality of information. Active attacks include data modification, Denial of Service (DoS), and repudiation attacks. They pose threats to data integrity and availability.
In a modification attack, the unauthorized user attempts to modify information for malicious purposes. A modification attack can target data at rest or data in transit. These attacks pose a threat to data integrity.
Denial of Service (DoS) attacks denies the use of resources to legitimate users.
These attacks generally do not involve access to or modification of information on the computer system. Instead, they pose a threat to data availability. The intentional flooding of a network or website to prevent legitimate access to authorized users is one example of a DoS attack.
Repudiation is an attack against the accountability of the information. It attempts to provide false information by either impersonating someone or denying that an event or a transaction has taken place.

c. Write a note on: i) Assets

Information is one of the most important assets for any organization. Other assets include hardware, software, and the network infrastructure required to access this information. To protect these assets, organizations must develop a set of parameters to ensure the availability of the resources to authorized users and trusted networks. These parameters apply to storage resources, the network infrastructure, and organizational policies.
Several factors need to be considered when planning for asset security. Security methods have two objectives. First objective is to ensure that the network is easily accessible to authorized users. It should also be reliable and stable under disparate environmental conditions and volumes of usage.

ii) Vulnerability. (08 Marks)
The paths that provide access to information are the most vulnerable to potential attacks. Each of these paths may contain various access points, each of which provides different levels of access to the storage resources. It is very important to implement adequate security controls at all the access points on an access path. Implementing security controls at each access point of every access path is termed as defense in depth.
Defense in depth recommends protecting all access points within an environment.
This reduces vulnerability to an attacker who can gain access to storage resources by bypassing inadequate security controls implemented at the vulnerable single point of access.

Such an attack can jeopardize the security of information assets. For example, a failure to properly authenticate a user may
put the confidentiality of information at risk. Similarly, a DoS attack against a storage device can jeopardize information availability.