CBCS SCHEME

USN □□□□□□□□□□

18CS822

Eighth Semester B.E. Degree Examination, June/July 2024

Storage Area Networks

Time: 3 hrs. Max. Marks: 100

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

Important Note : 1. On completing your answers, compulsorily draw diagonal cross lines on the remaining blank pages.
2. Any revealing of identification, appeal to evaluator and /or equations written eg, 42+8 = 50, will be treated as malpractice.

Module-1

1 a. Discuss the various factors that have contributed to the growth of digital data. **(06 Marks)**
b. Illustrate with a figure, the evolution of storage architecture. **(08 Marks)**
c. Explain the key characteristics of a data center. **(06 Marks)**

OR

2 a. Explain with a figure, the process of mapping user files to the disk storage subsystem with an LVM. **(08 Marks)**
b. Discuss the most popular interface protocols used for host to storage communication. **(08 Marks)**
c. Explain with a figure, the platter. **(04 Marks)**

Module-2

3 a. Discuss the limitations of software RAID. **(04 Marks)**
b. Discuss the various RAID techniques. **(10 Marks)**
c. Explain with a figure, the RAIDO configuration. **(06 Marks)**

OR

4 a. Discuss the various components of an intelligent storage system. **(08 Marks)**
b. Discuss the various components of Fibre Channel (FC) SAN. **(08 Marks)**
c. Explain with a figure, the active-passive configuration. **(04 Marks)**

Module-3

5 a. Explain with a figure, the FCIP protocol stack. **(06 Marks)**
b. Discuss the various benefits offered by the Network Attached Storage (NAS). **(08 Marks)**
c. Explain the process of handling I/Os in a network attached storage environment. **(06 Marks)**

OR

6 a. Explain with a figure, the iSCSI command sequencing. **(06 Marks)**
b. Explain the various factors affecting Network Attached Storage (NAS) performance. **(06 Marks)**
c. Explain with a figure, the various components of Network Attached Storage (NAS). **(08 Marks)**

Module-4

7 a. What is information availability? Discuss the causes of information availability. **(06 Marks)**
b. Explain the various stages involved in the Business Continuity (BC) planning life cycle. **(08 Marks)**
c. List the various tasks involved in business impact analysis. **(06 Marks)**

OR

8 a. Explain with a figure, the backup and restore operation. (10 Marks)
b. Explain with a figure, the two different types of backup topologies. (10 Marks)

Module-5

9 a. Explain the various terms which are used to represent entities and operations in a replication environment. (06 Marks)
b. Explain with a figure, the storage array based local replication. (06 Marks)
c. Illustrate with a figure, the synchronous and asynchronous replication. (08 Marks)

OR

10 a. Illustrate with a figure, the synchronous and asynchronous array based remote replication. (08 Marks)
b. Discuss the four security goals which are achieved through information security framework. (06 Marks)
c. Explain with a figure, the storage security domains. (06 Marks)

* * * * *

1. **a. Discuss the various factors that have contributed to the growth of digital data.**

The following is a list of some of the factors that have contributed to the growth of digital data:

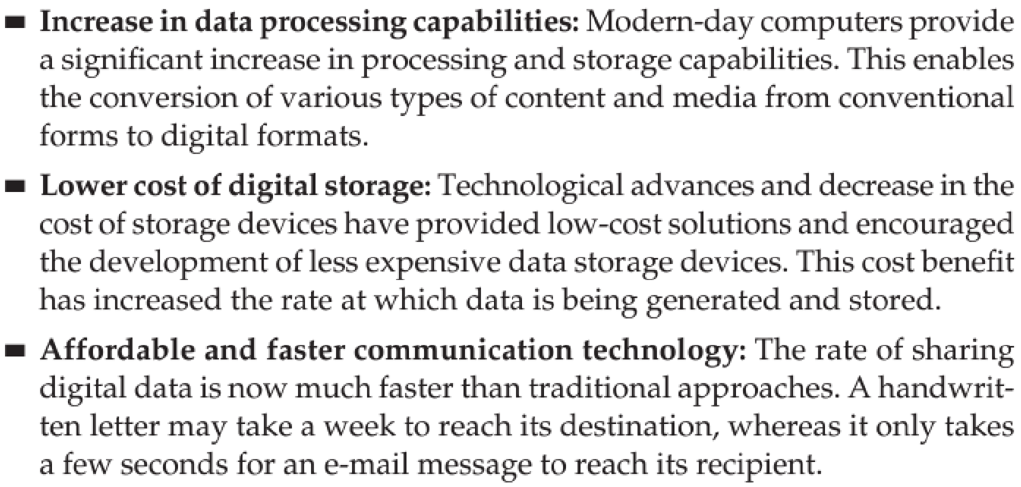

- **Increase in data processing capabilities:** Modern-day computers provide a significant increase in processing and storage capabilities. This enables the conversion of various types of content and media from conventional forms to digital formats.
- **Lower cost of digital storage:** Technological advances and decrease in the cost of storage devices have provided low-cost solutions and encouraged the development of less expensive data storage devices. This cost benefit has increased the rate at which data is being generated and stored.
- **Affordable and faster communication technology:** The rate of sharing digital data is now much faster than traditional approaches. A handwritten letter may take a week to reach its destination, whereas it only takes a few seconds for an e-mail message to reach its recipient.

Inexpensive and easier ways to create, collect, and store all types of data, coupled with increasing individual and business needs, have led to accelerated data growth, popularly termed the data explosion. Data has different purposes and criticality, so both individuals and businesses have contributed in varied proportions to this data explosion. The importance and the criticality of data vary with time. Most of the data created holds significance in the short-term but becomes less valuable over time. This governs the type of data storage solutions used. Individuals store data on a variety of storage devices, such as hard disks, CDs, DVDs, or Universal Serial Bus (USB) flash drives.

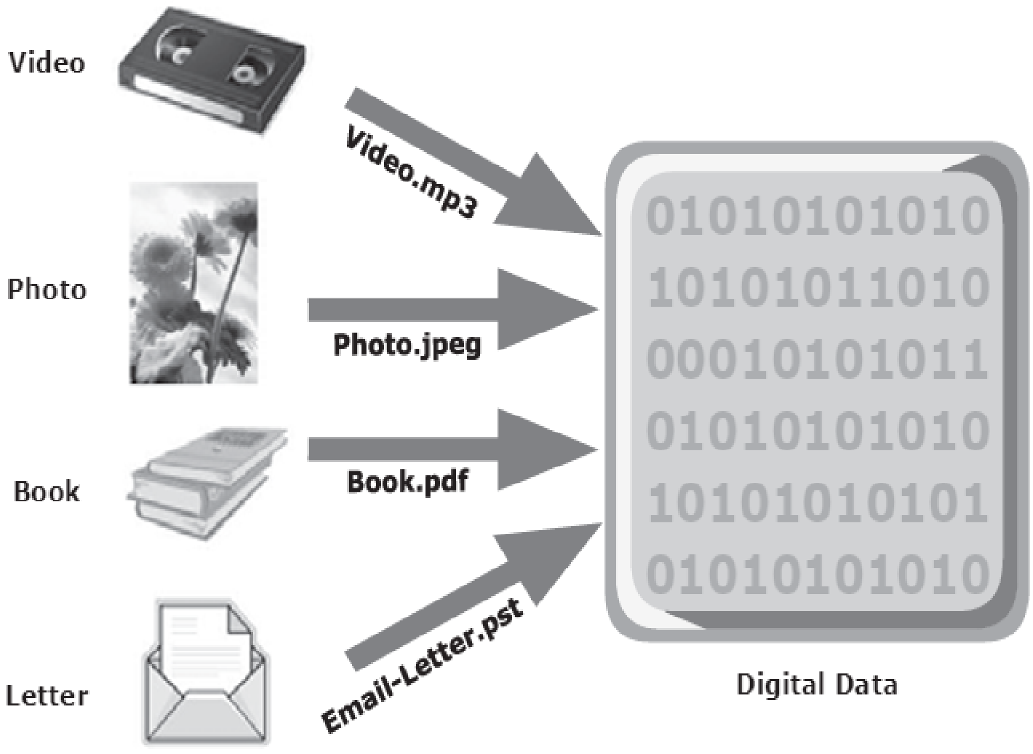

b. Illustrate with a figure, the evolution of storage architecture.

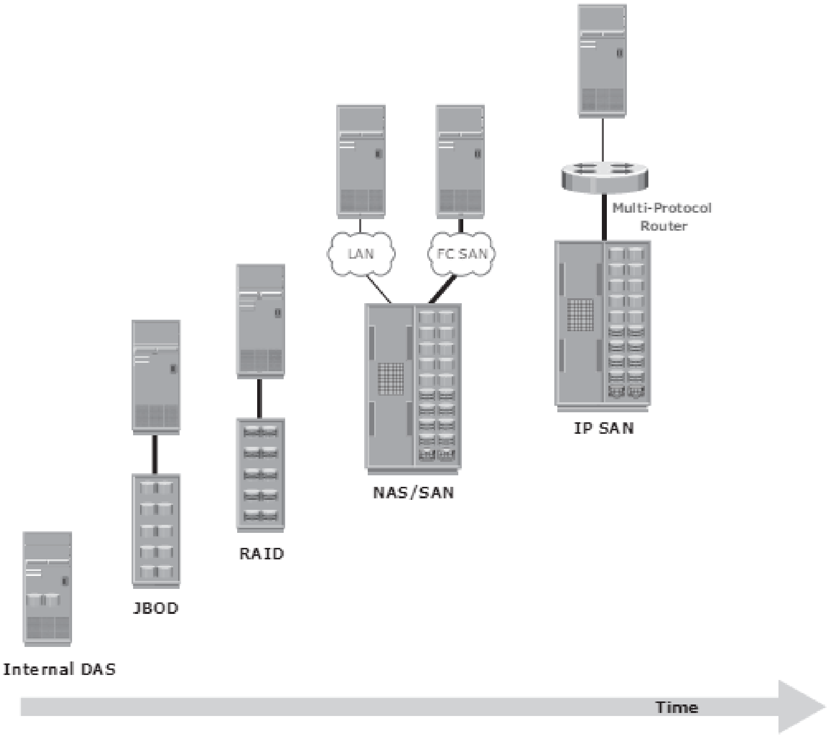

Figure 1-4: Evolution of storage architectures

- Historically, organizations had centralized computers (mainframe) and information storage devices (tape reels and disk packs) in their data center. The evolution of open systems and the affordability and ease of deployment that they offer 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. The proliferation of departmental servers in an enterprise resulted in unprotected, unmanaged, fragmented islands of information and increased operating cost. Originally, there were very limited policies and processes for managing these servers and the data created. To overcome these challenges, storage technology evolved from non-intelligent internal storage to intelligent networked storage (see Figure 1-4). Highlights of this technology evolution include:
- **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.

Storage technology and architecture continues to evolve, which enables organizations to consolidate, protect, optimize, and leverage their data to achieve the highest return on information assets.

c. Explain the key characteristics of a data center.

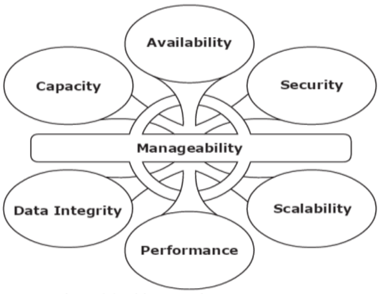

Figure 1-6: Key characteristics of data center elements

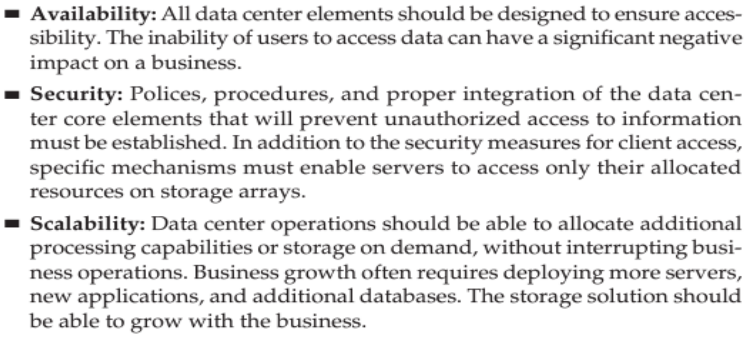

- **Availability:** All data center elements should be designed to ensure accessibility. The inability of users to access data can have a significant negative impact on a business.
- **Security:** Polices, procedures, and proper integration of the data center core elements that will prevent unauthorized access to information must be established. In addition to the security measures for client access, specific mechanisms must enable servers to access only their allocated resources on storage arrays.
- **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. The storage solution should be able to grow with the business.

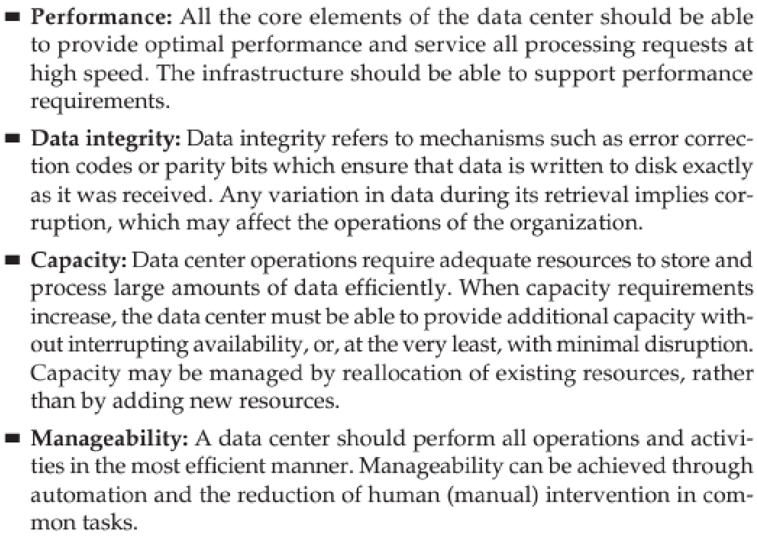

- **Performance:** All the core elements of the data center should be able to provide optimal performance and service all processing requests at high speed. The infrastructure should be able to support performance requirements.
- **Data integrity:** Data integrity refers to mechanisms such as error correction codes or parity bits which ensure that data is written to disk exactly as it was received. Any variation in data during its retrieval implies corruption, which may affect the operations of the organization.
- **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.

2. A. Explain with a figure, the process of mapping user files to the disk storage subsystem with an LVM (Logical Volume Manager)

The process of mapping user files to the disk storage subsystem using a **Logical Volume Manager (LVM)** can be understood as follows:

Logical Volume Manager (LVM) Overview

LVM is a device-mapping layer that abstracts the physical storage devices (such as hard drives) and presents them as logical volumes. It provides flexibility in managing storage space, allowing for dynamic resizing and easier management of storage resources.

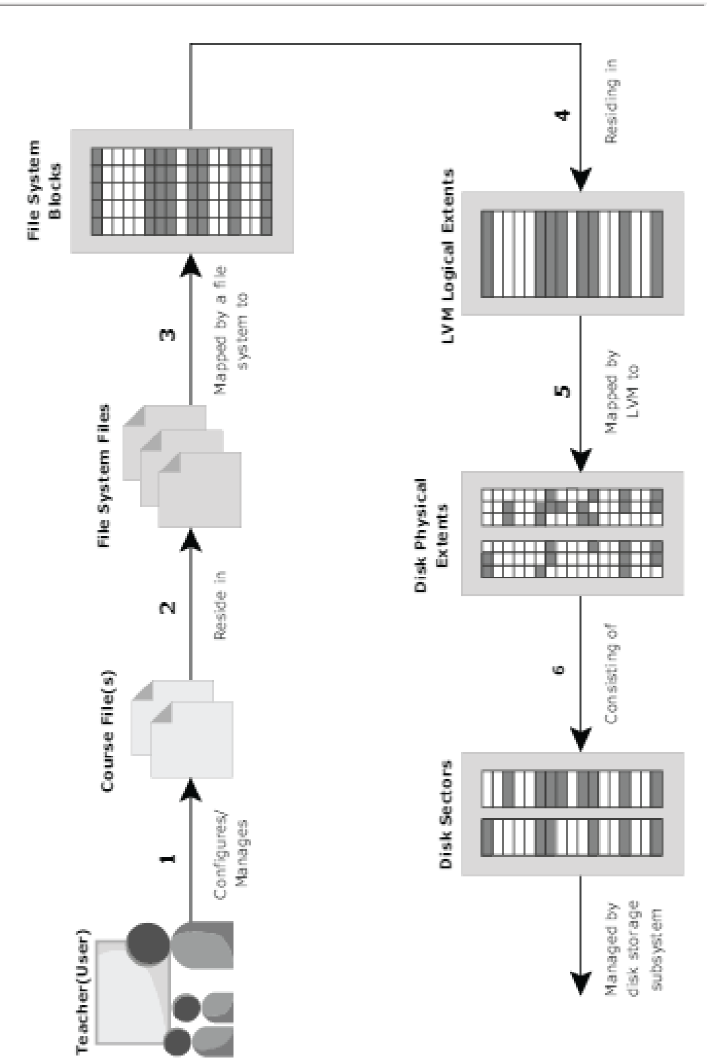

Figure 2-12: Process of mapping user files to disk storage

Steps in Mapping User Files to Disk Storage Using LVM:

1. **Physical Volumes (PVs):**
   The underlying physical storage devices (e.g., hard disks or partitions) are first assigned to the LVM as **physical volumes**. These physical volumes are the basic building blocks of the storage.

2. **Volume Groups (VGs):**
   Multiple physical volumes are grouped together to form a **volume group (VG)**. A volume group provides a pool of storage that can be used to create logical volumes. The VG abstracts the underlying physical volumes, providing flexibility for adding or removing physical storage devices.
3. **Logical Volumes (LVs):**
   From the volume group, **logical volumes (LVs)** are created. These logical volumes are flexible storage spaces that can be resized, extended, or reduced without directly affecting the underlying physical volumes. Logical volumes are presented to the operating system as if they were regular disk partitions.
4. **File System Creation:**
   On the logical volumes, file systems are created. User files are stored within these file systems, which are mounted and accessed by the operating system.
5. **Mapping Process:**
   The **mapping** of user files happens when the operating system writes data to the file system on the logical volumes. The LVM translates these writes into appropriate operations on the underlying physical volumes. The LVM handles the distribution of data across the physical volumes in the volume group, allowing for transparent storage management.

Benefits of LVM:

- **Flexibility in Storage Management:** Logical volumes can be resized or moved across physical volumes without impacting user data.
- **Dynamic Disk Allocation:** LVM allows administrators to dynamically allocate and manage storage as needed, reducing downtime.
- **Snapshot Capabilities:** LVM enables the creation of snapshots (point-in-time copies), which are useful for backups and data recovery.

- **Physical Volumes (PVs):** Multiple disks or partitions represented as physical volumes.
- **Volume Group (VG):** A logical grouping of the physical volumes.
- **Logical Volumes (LVs):** Created from the volume group.
- **File System:** Created on top of the logical volume, where user files are stored.

b. Discuss the most popular interface protocols used for host-to-storage communication.

There are several popular **interface protocols** used for host-to-storage communication, each suited for specific environments and performance needs. The most common protocols include:

- **Fibre Channel (FC)** for high-performance SANs.
- **iSCSI** for cost-effective block-level access over IP networks.
- **FCoE** for converged network environments.
- **NFS and CIFS** for file-level access in NAS systems.
- **AoE** for cost-effective Ethernet-based block storage.

Each of these protocols plays a critical role depending on the environment (SAN vs. NAS), performance needs, and budget.

1. Fibre Channel (FC)

- **Description**:
  Fibre Channel (FC) is a high-speed networking protocol primarily used in **Storage Area Networks (SANs)**. It provides a dedicated, high-performance path between the server (host) and storage devices, often with speeds ranging from 2 Gbps to 32 Gbps.
- **Key Features**:
  - **Block-level storage** protocol, allowing direct communication between the server and the storage at the block level.
  - **Low Latency**: FC is known for low latency and high throughput, making it ideal for mission-critical applications such as databases, virtualization, and high-transaction environments.
  - **Highly Reliable**: Fibre Channel supports redundancy and multiple paths for fault tolerance.
- **Use Cases**:
  Enterprise-level data centers, large-scale SAN environments, and high-performance computing where uptime and speed are critical.

---

2. Internet Small Computer Systems Interface (iSCSI)

- **Description**:
  iSCSI is a popular protocol used to connect hosts (servers) to storage over an **IP-based network**. It encapsulates **SCSI commands** over TCP/IP networks, making it possible to use standard Ethernet infrastructure for SAN-like storage communication.
- **Key Features**:
  - **Cost-Effective**: iSCSI leverages existing Ethernet networks, reducing the need for dedicated Fibre Channel infrastructure.
  - **Block-Level Storage**: Like FC, iSCSI provides block-level access to storage devices.
  - **Scalability**: iSCSI can scale across distances and networks, making it a practical solution for remote data centers and disaster recovery setups.
- **Use Cases**:
  iSCSI is often used in **small to medium-sized enterprises (SMEs)** or data centers where cost is a consideration but where SAN-level performance is still required. It's also a good choice for remote replication and backup scenarios.

---

3. Fibre Channel over Ethernet (FCoE)

- **Description**:
  FCoE combines the benefits of **Fibre Channel** and **Ethernet** by encapsulating Fibre Channel frames over a **10 Gigabit Ethernet (GbE)** network. It allows organizations to converge storage and IP traffic onto a single network, reducing the need for separate SAN and LAN infrastructures.
- **Key Features**:
  - **Converged Network**: FCoE provides a unified infrastructure for both storage and network traffic, simplifying management and reducing costs.

- **Compatibility**: It uses Ethernet networks, but the data is processed using the Fibre Channel protocol, ensuring compatibility with existing FC-based storage systems.
    - **High Performance**: FCoE offers high bandwidth (10GbE or higher) while maintaining the low-latency and high-reliability characteristics of FC.
- **Use Cases**:
  FCoE is commonly used in **data centers** that want to reduce network complexity and costs by converging storage and networking traffic. It is suitable for virtualized environments and high-performance storage needs.

---

4. Network File System (NFS) and Common Internet File System (CIFS)

- **Description**:
  NFS (primarily used in Unix/Linux environments) and **CIFS** (also known as SMB in Windows environments) are file-level protocols used in **Network Attached Storage (NAS)** environments.
- **Key Features**:
    - **File-Level Storage**: Unlike block-level protocols (like FC and iSCSI), NFS and CIFS provide file-level access. This means that files are shared over the network, and multiple users can access the same files simultaneously.
    - **Standard Network Use**: Both protocols work over standard IP networks, allowing clients and servers to communicate using Ethernet.
    - **Ease of Use**: NFS and CIFS allow easy file sharing across different devices and users, making them practical for collaborative environments.
- **Use Cases**:
  NFS and CIFS are typically used in NAS systems, where file-sharing capabilities are critical. They are commonly used in environments requiring **centralized file storage** and collaboration, such as **office networks** or **content management systems**.

---

5. ATA over Ethernet (AoE)

- **Description**:
  AoE is a network protocol that allows **Advanced Technology Attachment (ATA)** storage devices to be accessed over Ethernet networks. It is a simpler and less expensive protocol compared to iSCSI and Fibre Channel.
- **Key Features**:
    - **Block-Level Access**: Like iSCSI, AoE provides block-level access over Ethernet.
    - **Cost-Effective**: It is particularly attractive for low-cost, high-capacity storage solutions.
    - **Less Overhead**: AoE operates at the Ethernet layer and does not rely on TCP/IP, reducing some of the overhead seen in IP-based protocols.
- **Use Cases**:
  AoE is used in storage systems where cost is a major concern, such as **small businesses** or environments requiring cheap, bulk storage. It is also suitable for applications that don’t require the complex features offered by protocols like iSCSI or FC.

---

6. HyperSCSI

- **Description**:
  HyperSCSI is a network protocol designed for transporting **SCSI commands** directly over **Ethernet**, bypassing the TCP/IP stack. It is a lesser-known protocol compared to iSCSI.
- **Key Features**:
  - **Low Overhead**: By skipping the TCP/IP layer, HyperSCSI reduces the processing overhead and improves efficiency in data transport.
  - **Block-Level Access**: Like other SCSI-based protocols, it provides block-level access to storage devices.
- **Use Cases**:
  HyperSCSI is generally used in **high-performance storage networks**, where minimizing overhead and maximizing throughput is critical.

---

c. Explain, with a figure, the platter.

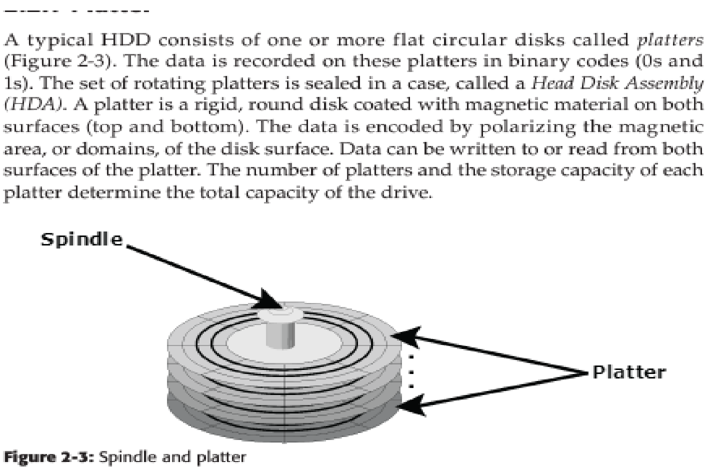

A typical HDD consists of one or more flat circular disks called *platters* (Figure 2-3). The data is recorded on these platters in binary codes (0s and 1s). The set of rotating platters is sealed in a case, called a *Head Disk Assembly (HDA)*. A platter is a rigid, round disk coated with magnetic material on both surfaces (top and bottom). The data is encoded by polarizing the magnetic area, or domains, of the disk surface. Data can be written to or read from both surfaces of the platter. The number of platters and the storage capacity of each platter determine the total capacity of the drive.

Figure 2-3: Spindle and platter

3. a. Discuss the limitations of software RAID.

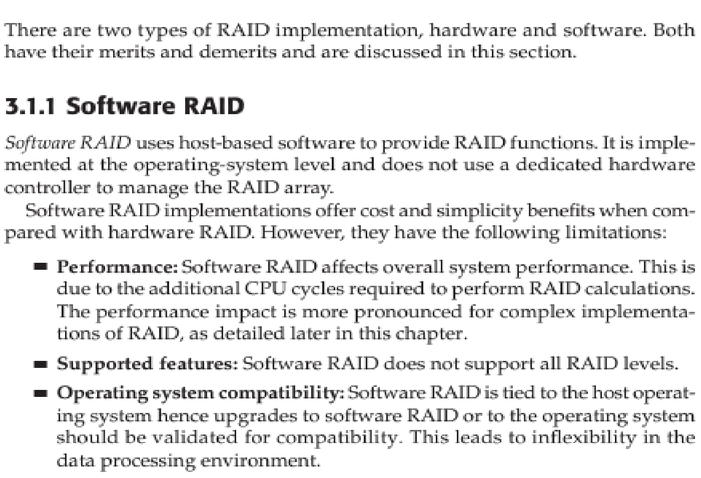

There are two types of RAID implementation, hardware and software. Both have their merits and demerits and are discussed in this section.

3.1.1 Software RAID

Software RAID uses host-based software to provide RAID functions. It is implemented at the operating-system level and does not use a dedicated hardware controller to manage the RAID array.

Software RAID implementations offer cost and simplicity benefits when compared with hardware RAID. However, they have the following limitations:

- **Performance:** Software RAID affects overall system performance. This is due to the additional CPU cycles required to perform RAID calculations. The performance impact is more pronounced for complex implementations of RAID, as detailed later in this chapter.
- **Supported features:** Software RAID does not support all RAID levels.
- **Operating system compatibility:** Software RAID is tied to the host operating system hence upgrades to software RAID or to the operating system should be validated for compatibility. This leads to inflexibility in the data processing environment.

b. Discuss various RAID techniques.

RAID (Redundant Array of Independent Disks) provides different levels of data redundancy and performance based on how data is distributed across multiple disks. Common RAID techniques include:

Table 3-1: Raid Levels

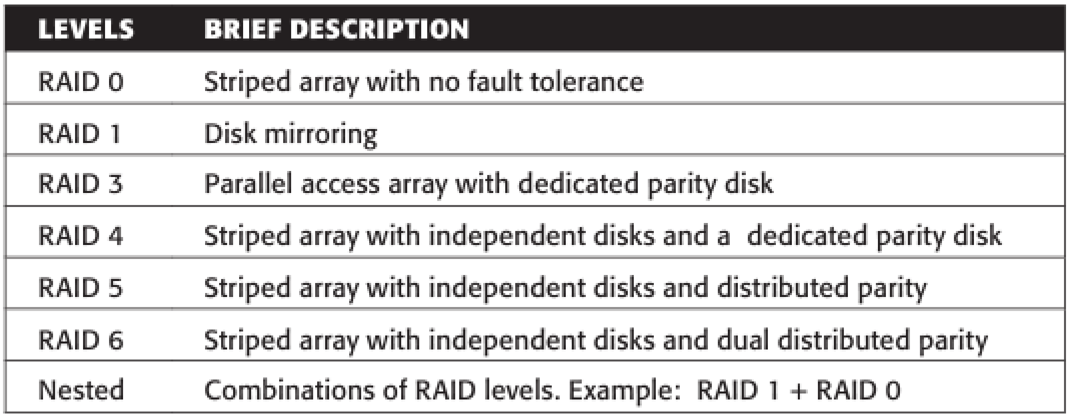

LEVELS	BRIEF DESCRIPTION
RAID 0	Striped array with no fault tolerance
RAID 1	Disk mirroring
RAID 3	Parallel access array with dedicated parity disk
RAID 4	Striped array with independent disks and a dedicated parity disk
RAID 5	Striped array with independent disks and distributed parity
RAID 6	Striped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

1. **RAID 0 (Striping)**:
    - **Data distribution**: Splits data evenly across two or more disks without redundancy.
    - **Advantages**: Improves performance since multiple disks are used simultaneously.
    - **Disadvantages**: No fault tolerance. If one disk fails, all data is lost.
2. **RAID 1 (Mirroring)**:
    - **Data distribution**: Duplicates data across two disks (or more).
    - **Advantages**: High fault tolerance. If one disk fails, the other contains a copy of the data.
    - **Disadvantages**: Expensive since it requires double the disk space.
3. **RAID 5 (Striping with Parity)**:
    - **Data distribution**: Distributes data and parity information across three or more disks. Parity is used to rebuild data if a disk fails.
    - **Advantages**: Fault tolerance and efficient use of disk space (only one disk used for parity).
    - **Disadvantages**: Write performance is lower due to the need to calculate parity.
4. **RAID 6 (Double Parity)**:
    - **Data distribution**: Similar to RAID 5 but with two parity blocks. Requires at least four disks.
    - **Advantages**: Can tolerate two simultaneous disk failures.
    - **Disadvantages**: Slower performance due to additional parity calculations.
5. **RAID 10 (1+0)**:
    - **Data distribution**: Combines RAID 1 (mirroring) and RAID 0 (striping). Requires at least four disks.
    - **Advantages**: Combines high fault tolerance with good performance.
    - **Disadvantages**: Expensive since it uses mirroring, requiring double the storage capacity.

c. Explain, with a figure, RAID 0 configuration.

In a RAID 0 configuration, data is striped across the HDDs in a RAID set. It utilizes the full storage capacity by distributing strips of data over multiple HDDs in a RAID set. To read data, all the strips are put back together by the controller. The stripe size is specified at a host level for software RAID and is vendor specific for hardware RAID. Figure 3-5 shows RAID 0 on a storage array in which data is striped across 5 disks. When the number of drives in the array increases, performance improves because more data can be read or written simultaneously. RAID 0 is used in applications that need high I/O throughput. However, if these applications require high availability, RAID 0 does not provide data protection and availability in the event of drive failures.

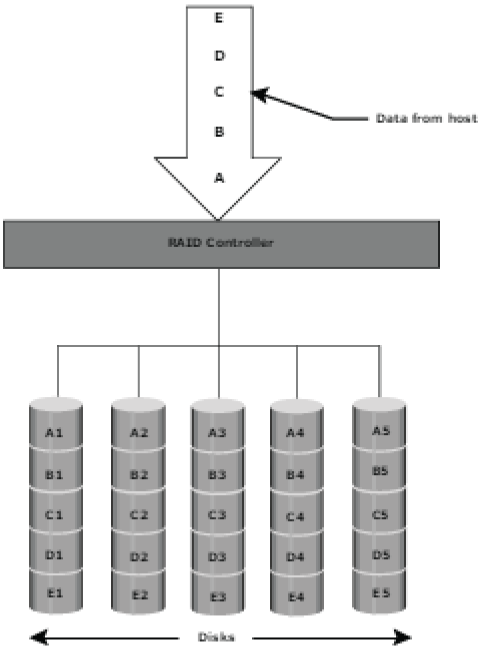

Figure 3-5: RAID 0

4. a. Discuss the components of an intelligent storage system.

An intelligent storage system consists of four key components: front end, cache, back end, and physical disks. Figure 4-1 illustrates these components and their interconnections. An I/O request received from the host at the front-end port is processed through cache and the back end, to enable storage and retrieval of data from the physical disk. A read request can be serviced directly from cache if the requested data is found in cache.

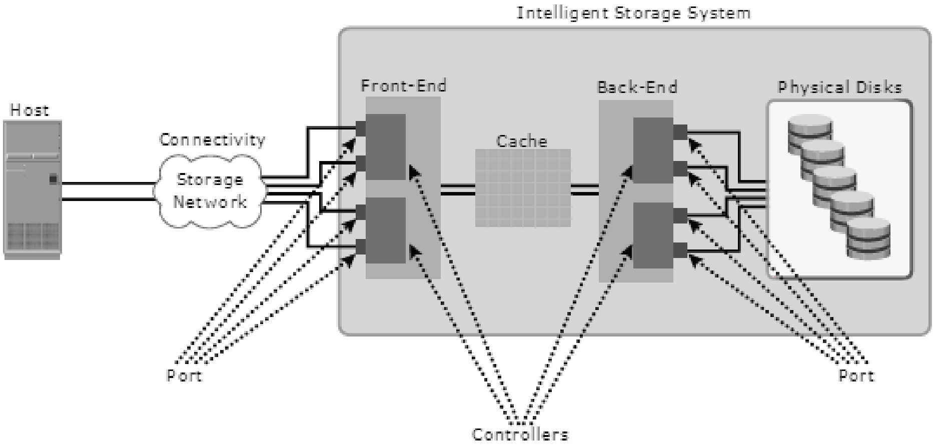

Figure 4-1: Components of an intelligent storage system

An **intelligent storage system** includes several advanced features to manage and optimize data storage, performance, and reliability. The major components are:

1. **Front-end (I/O Interface)**:
    - Handles communication between the hosts and the storage system.
    - Manages the data transfer requests from servers.
    - Interfaces like Fibre Channel or iSCSI are used.
2. **Cache**:
    - Temporary storage that holds frequently accessed data, speeding up read and write operations.
    - Reduces latency by serving requests faster compared to disk-based storage.
3. **Back-end (Disk Controller)**:
    - Manages the read and write operations to the physical disk drives.
    - Communicates with the disks using protocols like SATA, SAS, or Fibre Channel.
4. **Physical Disks**:
    - Stores the data in RAID groups, ensuring redundancy and fault tolerance.
5. **Management Software**:
    - Software layer that automates tasks like provisioning, snapshots, replication, and monitoring of storage.

b. Discuss Fibre Channel (FC) SAN components.

A **Fibre Channel SAN** is a high-speed network that connects storage devices to servers. The main components include:

1. **Node Ports (N_Ports)**:
    - Interfaces on the host and storage devices used to send and receive data. Host N_Ports are on HBAs, and storage N_Ports are on storage controllers.
2. **Fabric (F_Ports)**:
    - Switches or directors that interconnect multiple devices in the SAN.
    - Provides a central point for routing data between servers and storage devices.
3. **HBA (Host Bus Adapter)**:

- A hardware component installed on the server, allowing the server to connect to the Fibre Channel network.

4. **Storage Arrays**:
   - Storage devices or disk arrays that store the data. These are connected to the SAN and accessed by multiple hosts.
5. **Zoning**:
   - Logical separation of devices in the SAN fabric, providing security by allowing certain hosts to communicate only with specific storage devices.

c. Explain, with a figure, the active-passive configuration.

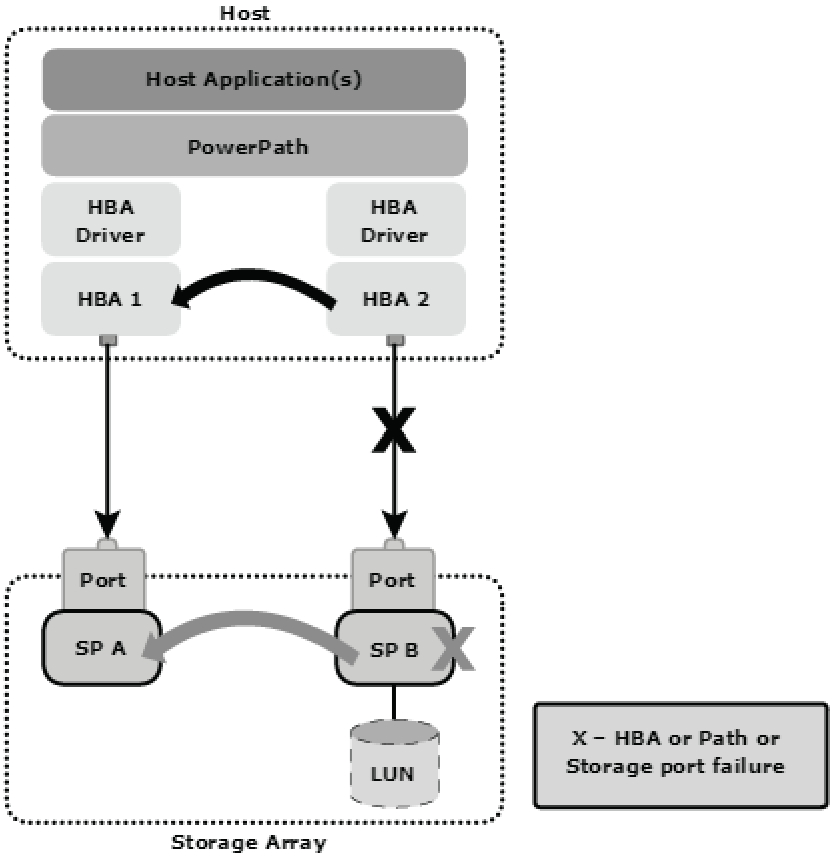

Figure 11-10: Path failover with PowerPath for an active-passive array

Path failure can occur due to a failure of the link, HBA, or storage processor (SP). In the event of a path failure, PowerPath with an active-passive configuration performs the path failover operation in the following way: If an active I/O path to SP B through HBA 2 fails, PowerPath uses a passive path to SP B through HBA 1. If HBA 2 fails, the application uses HBA 1 to access the logical device. If SP B fails, PowerPath stops all I/O to SP B and trespasses the device over to SP A. All I/O will be sent down the paths to SP A, this process is referred as LUN trespassing. When SP B is brought back online, PowerPath recognizes that it is available and resumes sending I/O down to SP B.

5. a. Explain, with a figure, the FCIP protocol stack.

Application layer generate SCSI commands and data, which are processed by various layers of the protocol stack.

The higher layer protocol SCSI includes the SCSI driver program that executes the read-and-write commands. Below the SCSI layer is that the Fibre Channel Protocol(FCP) layer, that is just a Fibre Channel frame whose payload is SCSI. The FCP layer rides on the top of the FC transport layer. This allows the FC frames to run natively inside theSAN environment. Additionally, the FC frames are often encapsulated into an IP packet and sent to a distant SAN over an IP Network.

FCIP Protocol Stack:

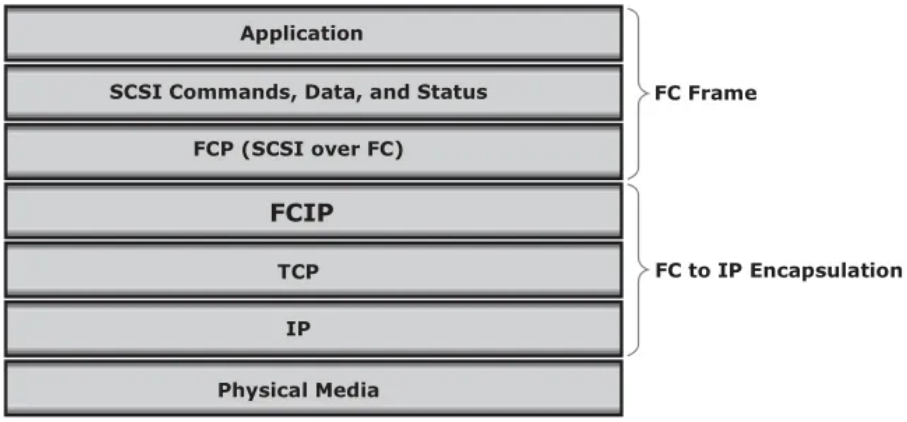

Fig 2: FCIP Protocol Stack

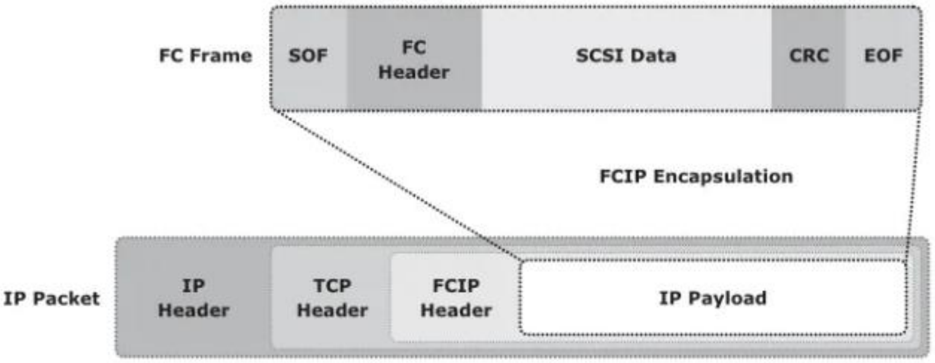

Fig 3: FCIP Encapsulation

Encapsulation of FC frame into an IP packet could cause the IP packet to be fragmented when the data link cannot support the maximum transmission unit (MTU) size of an IP packet. When the IP packet is fragmented, the required parts of the header must be copied by all fragments. When the TCP packet is segmented, normal TCP operations are used for receiving and re-sequencing the data before passing it to the FC processing portion of the device.

b. Discuss the benefits of Network Attached Storage (NAS).

NAS offers the following benefits:

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

■ Improved efficiency: Eliminates bottlenecks that occur during file access from a general-purpose file server because NAS uses an operating system specialized for file serving. It improves the utilization of general-purpose servers by relieving them of file-server operations.

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

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

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

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

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

■ Security: Ensures security, user authentication, and file locking in conjunction.

c. Explain how I/Os are handled in a network-attached storage environment.

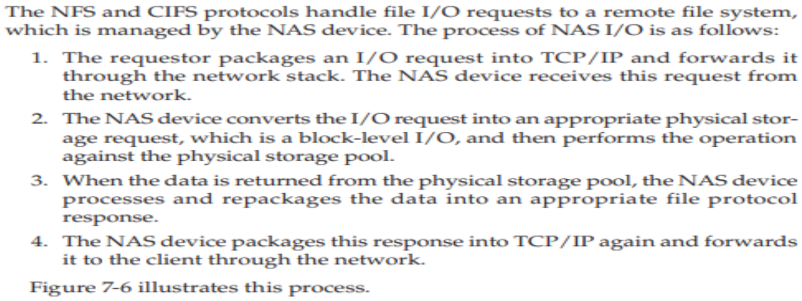

The NFS and CIFS protocols handle file I/O requests to a remote file system, which is managed by the NAS device. The process of NAS I/O is as follows:

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.

Figure 7-6 illustrates this process.

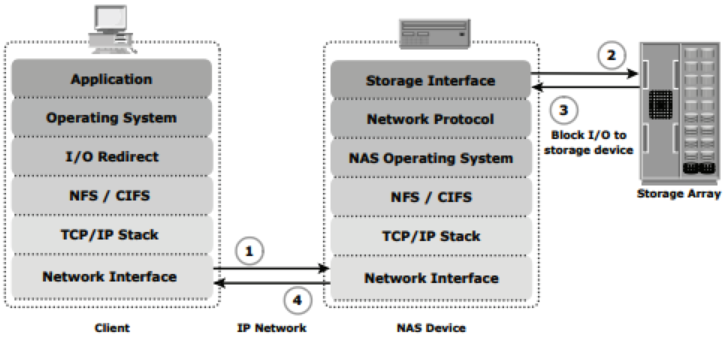

Figure 7-6: NAS I/O operation

6. a. Explain, with a figure, iSCSI command sequencing.

iSCSI communication between initiators and targets is based on the requestresponse command sequences. A command sequence may generate multiple PDUs. A command sequence number (CmdSN) within an iSCSI session is used to number all initiator-to-target command PDUs belonging to the session. This number is used to ensure that every command is delivered in the same order in which it is transmitted, regardless of the TCP connection that carries the command in the session. Command sequencing begins with the first login command and the CmdSN is incremented by one for each subsequent command. The iSCSI target layer is responsible for delivering the commands to the SCSI layer in the order of their CmdSN. This ensures the correct order of data and commands at a target even when there are multiple TCP connections between an initiator and the target using portal groups. Similar to command numbering, a status sequence number (StatSN) is used to sequentially number status responses, as shown in Figure 8-8. These unique numbers are established at the level of the TCP connection. A target sends the request-to-transfer (R2T) PDUs to the initiator when it is ready to accept data. Data sequence number (DataSN) is used to ensure in-order delivery of data within the same command. The DataSN and R2T sequence numbers are used to sequence data PDUs and R2Ts, respectively. Each of these sequence numbers is stored locally as an unsigned 32-bit integer counter defined by iSCSI. These numbers are communicated between the initiator and target in the appropriate iSCSI PDU fields during command, status, and data exchanges. In the case of read operations, the DataSN begins at zero and is incremented by one for each subsequent data PDU in that command sequence. In the case of a write operation, the first unsolicited data PDU or the first data PDU in response to an R2T begins with a DataSN of zero and increments by one for each subsequent data PDU. R2TSN is set to zero at the initiation of the command and incremented by one for each subsequent R2T sent by the target for that command.

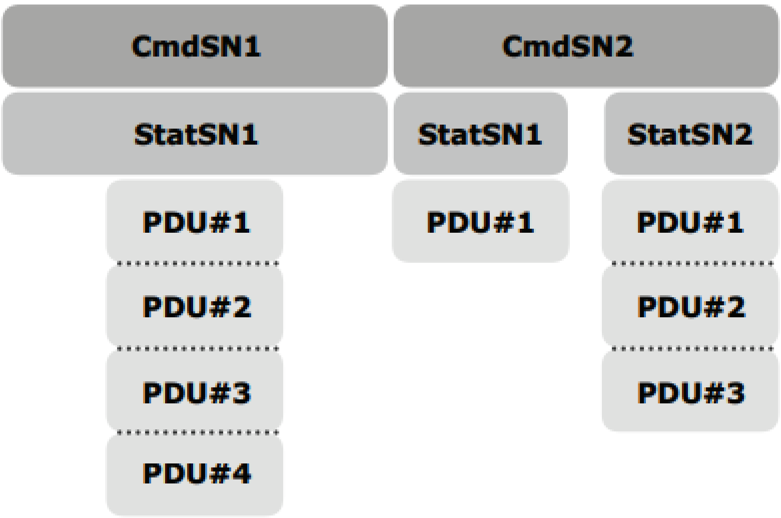

Figure 8-8: Command and status sequence number

b. Explain factors affecting NAS performance.

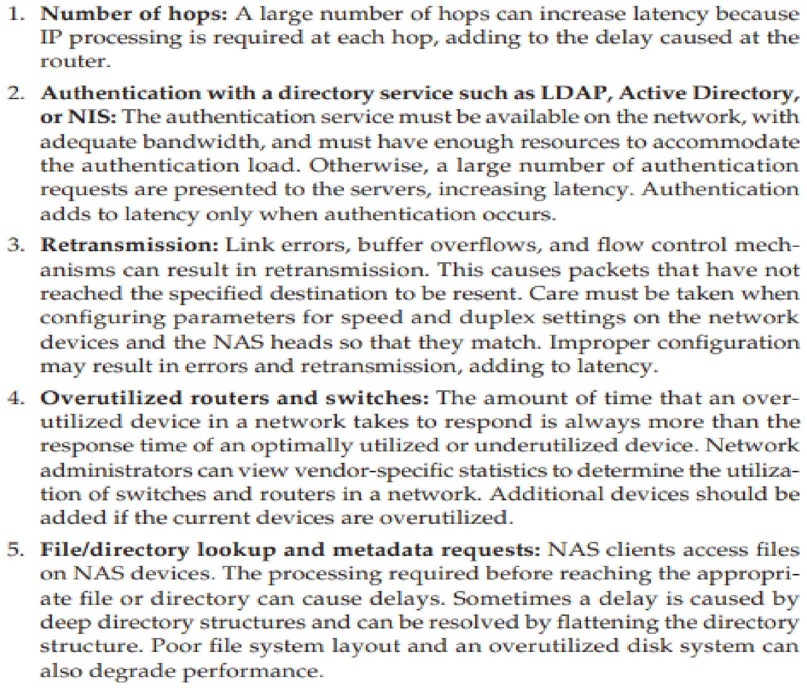

1. **Number of hops:** A large number of hops can increase latency because IP processing is required at each hop, adding to the delay caused at the router.
2. **Authentication with a directory service such as LDAP, Active Directory, or NIS:** The authentication service must be available on the network, with adequate bandwidth, and must have enough resources to accommodate the authentication load. Otherwise, a large number of authentication requests are presented to the servers, increasing latency. Authentication adds to latency only when authentication occurs.
3. **Retransmission:** Link errors, buffer overflows, and flow control mechanisms can result in retransmission. This causes packets that have not reached the specified destination to be resent. Care must be taken when configuring parameters for speed and duplex settings on the network devices and the NAS heads so that they match. Improper configuration may result in errors and retransmission, adding to latency.
4. **Overutilized routers and switches:** The amount of time that an overutilized device in a network takes to respond is always more than the response time of an optimally utilized or underutilized device. Network administrators can view vendor-specific statistics to determine the utilization of switches and routers in a network. Additional devices should be added if the current devices are overutilized.
5. **File/directory lookup and metadata requests:** NAS clients access files on NAS devices. The processing required before reaching the appropriate file or directory can cause delays. Sometimes a delay is caused by deep directory structures and can be resolved by flattening the directory structure. Poor file system layout and an overutilized disk system can also degrade performance.

c. Explain the components of NAS, with a figure.

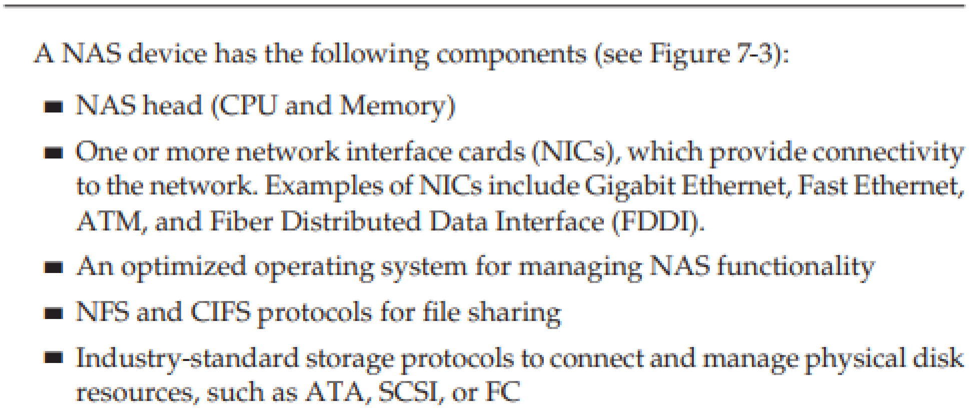

A NAS device has the following components (see Figure 7-3):

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

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

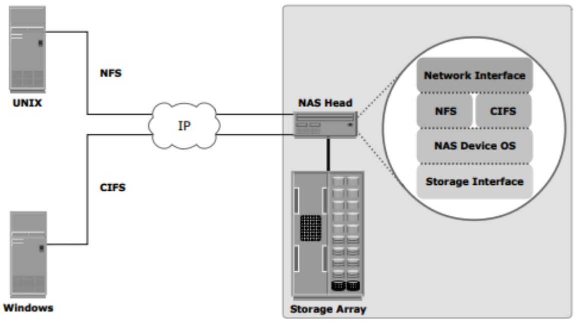

Figure 7-3: Components of NAS

7. **a. What is information availability? Discuss causes of its unavailability.**

Information availability (IA) refers to the ability of the infrastructure to function according to business expectations during its specified time of operation. Information availability ensures that people (employees, customers, suppliers, and partners) can access information whenever they need it. Information availability can be defined with the help of reliability, accessibility and timeliness.

■ Reliability: This reflects a component's ability to function without failure, under stated conditions, for a specified amount of time.

■ Accessibility: This is the state within which the required information is accessible at the right place, to the right user. The period of time during which the system is in an accessible state is termed system uptime; when it is not accessible it is termed system downtime.

■ Timeliness: Defines the exact moment or the time window (a particular time of the day, week, month, and/or year as specified) during which information must be accessible. For example, if online access to an application is required between 8:00 am and 10:00 pm each day, any disruptions to data availability outside of this time slot are not considered to affect timeliness.

Causes of Information Unavailability

Various planned and unplanned incidents result in data unavailability. Planned outages include installation/integration/maintenance of new hardware, software upgrades or patches, taking backups, application and data restores, facility operations (renovation and construction), and refresh/migration of the testing to the production environment. Unplanned outages include failure caused by database corruption, component failure, and human errors. Another type of incident that may cause data unavailability is natural or man-made disasters such as flood, fire, earthquake, and contamination. As illustrated in Figure 11-1, the majority of outages are planned. Planned outages are expected and scheduled, but still cause data to be unavailable. Statistically, less than 1 percent is likely to be the result of an unforeseen disaster.

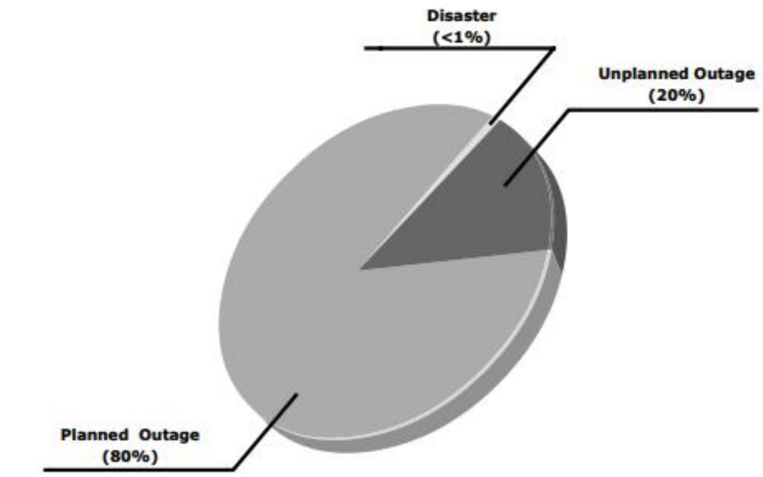

Figure 11-1: Disruptors of data availability

b. Explain stages in the Business Continuity (BC) planning life cycle.

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 (see Figure 11-3): 1. Establishing objectives 2. Analyzing 3. Designing and developing 4. Implementing 5. Training, testing, assessing, and maintaining

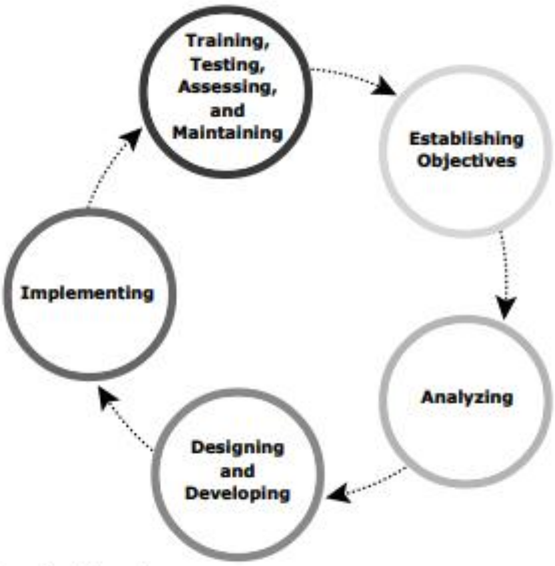

Figure 11-3: BC planning lifecycle

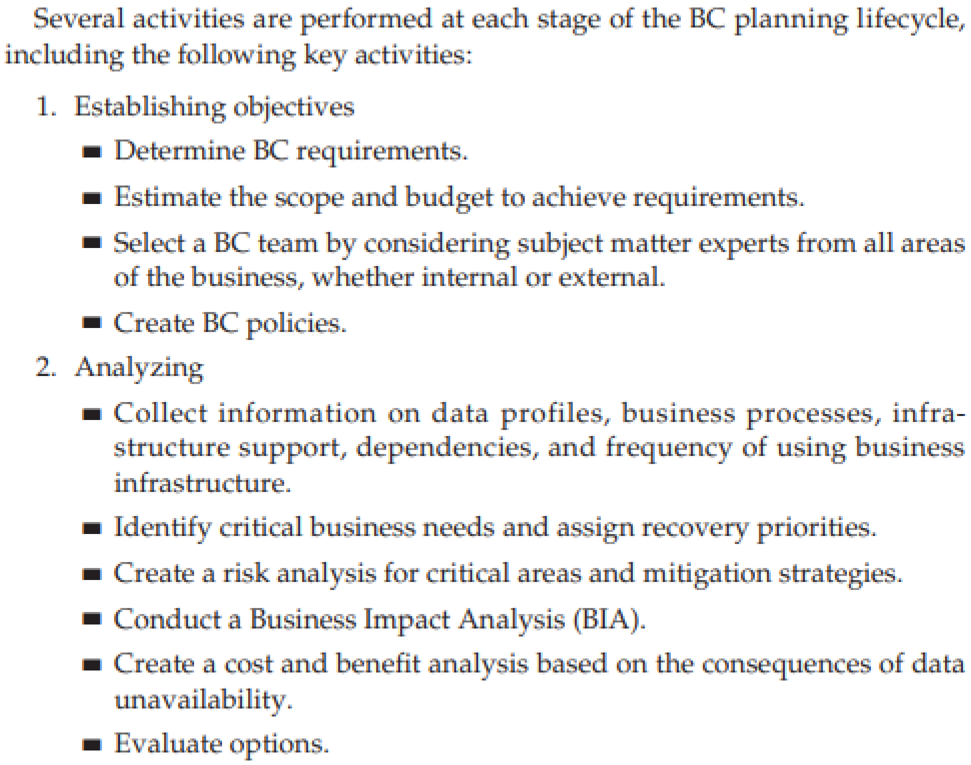

Several activities are performed at each stage of the BC planning lifecycle, including the following key activities:

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.

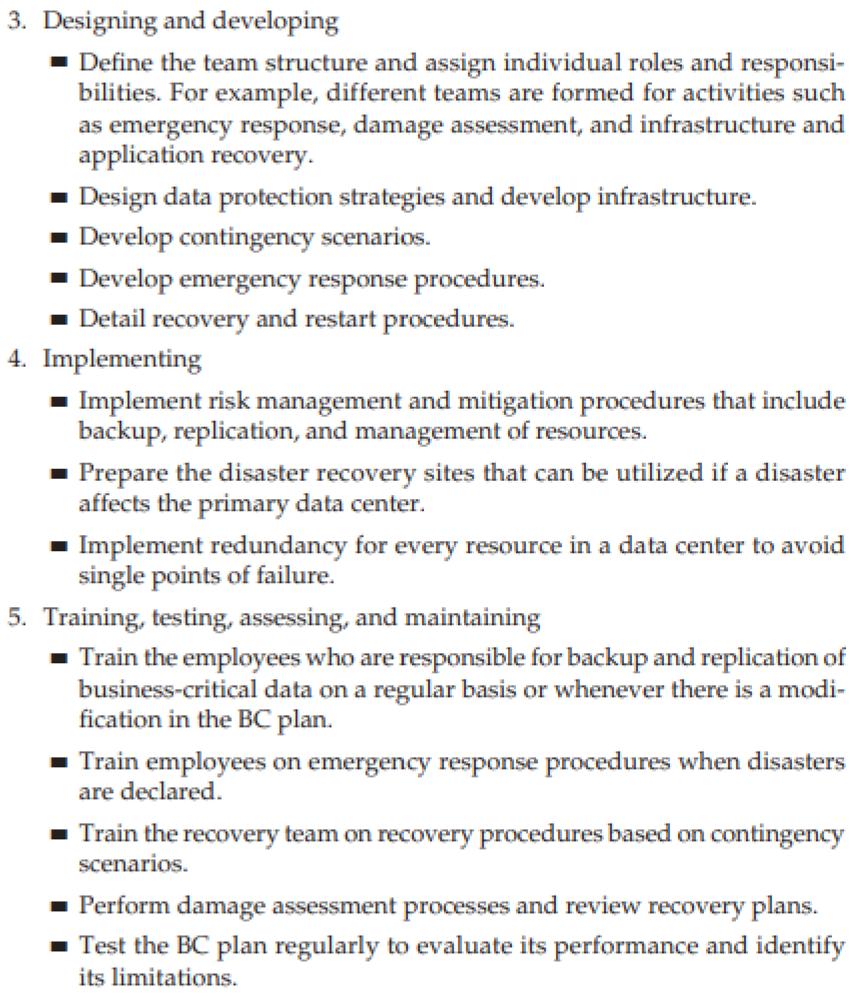

3. Designing and developing
   - Define the team structure and assign individual roles and responsibilities. For example, different teams are formed for activities such as emergency response, damage assessment, and infrastructure and application recovery.
   - Design data protection strategies and develop 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.
   - Perform damage assessment processes and review recovery plans.
   - Test the BC plan regularly to evaluate its performance and identify its limitations.

c. List tasks involved in business impact analysis.

A business impact analysis (BIA) identifies and evaluates financial, operational, and service impacts of a disruption to essential business processes. Selected functional areas are evaluated to determine resilience of the infrastructure to support information availability. The BIA process leads to a report detailing the incidents and their impact over business functions. The impact may be specified in terms of money or in terms of time. Based on the potential impacts associated with downtime, businesses can prioritize and implement countermeasures to mitigate the likelihood of such disruptions. These are detailed in the BC plan. A BIA includes the following set of tasks:

■ Identify the key business processes critical to its operation.

■ Determine the attributes of the business process in terms of applications, databases, and hardware and software requirements.

■ Estimate the costs of failure for each business process.

■ Calculate the maximum tolerable outage and define RTO and RPO for each business process. ■ Establish the minimum resources required for the operation of business processes.

■ Determine recovery strategies and the cost for implementing them. Optimize the backup and business recovery strategy based on business priorities.

■ Analyze the current state of BC readiness and optimize future BC planning.

8. a. Explain with a figure, the backup and restore operation.

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 (see Figure 12-5). The backup server maintains the information about backup clients to be contacted and storage nodes to be used in a backup operation. The backup server retrieves the backup-related information from the backup catalog and, based on this information, instructs the storage node to load the appropriate backup media into the backup devices. Simultaneously, it instructs the backup clients to start scanning the data, package it, and send it over the network to the assigned storage node. The storage node, in turn, sends metadata to the backup server to keep it updated about the media being used in the backup process. The backup server continuously updates the backup catalog with this information.

After the data is backed up, it can be restored when required. A restore process must be manually initiated. Some backup software has a separate application for restore operations. These restore applications are accessible only to the administrators. Upon receiving a restore request, an administrator opens the restore application to view the list of clients that have been backed up. While selecting the client for which a restore request has been made, the administrator also needs to identify the client that will receive the restored data. Data can be restored on the same client for whom the restore request has been made or on any other client. The administrator then selects the data to be restored and the specified point in time to which the data has to be restored based on the RPO. Note that because all of this information comes from the backup catalog, the restore application must also communicate to the backup server. The administrator first selects the data to be restored and initiates the restore process. The backup server, using the appropriate storage node, then identifies the backup media that needs to be mounted on the backup devices. Data is then read and sent to the client that has been identified to receive the restored data. Some restorations are successfully accomplished by recovering only the requested production data. For example, the recovery process of a spreadsheet is completed when the specific file is restored. In database restorations, additional data such as log files and production data must be restored. This ensures application consistency for the restored data. In these cases, the RTO is extended due to the additional steps in the restoration process.

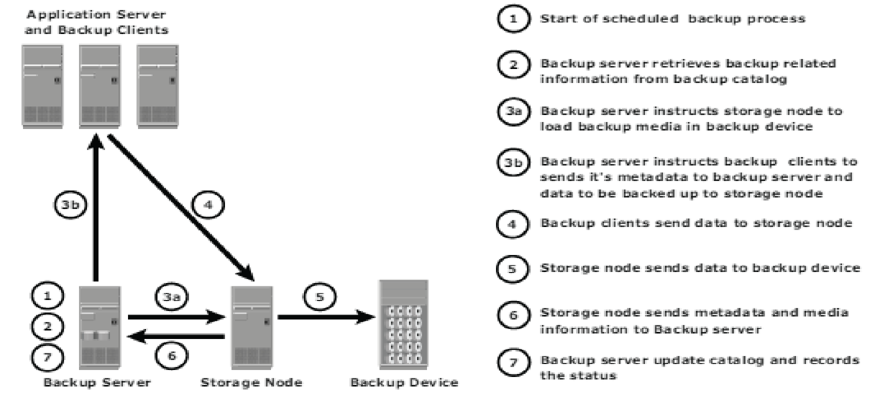
Application Server and Backup Clients
3b
4
1
2
7
3a
6
5
Backup Server
Storage Node
Backup Device
1 Start of scheduled backup process
2 Backup server retrieves backup related information from backup catalog
3a Backup server instructs storage node to load backup media in backup device
3b Backup server instructs backup clients to sends it's metadata to backup server and data to be backed up to storage node
4 Backup clients send data to storage node
5 Storage node sends data to backup device
6 Storage node sends metadata and media information to Backup server
7 Backup server update catalog and records the status

Figure 12-5: Backup operation

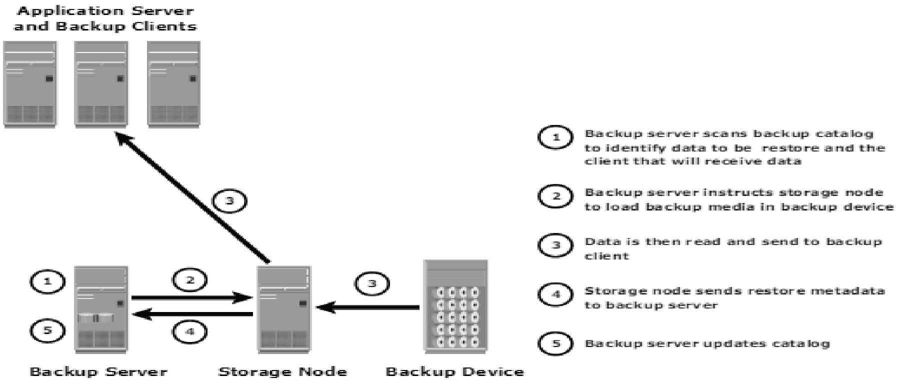
Application Server and Backup Clients
3
1
2
3
5
4
Backup Server
Storage Node
Backup Device
1 Backup server scans backup catalog to identify data to be restore and the client that will receive data
2 Backup server instructs storage node to load backup media in backup device
3 Data is then read and send to backup client
4 Storage node sends restore metadata to backup server
5 Backup server updates catalog

Figure 12-6: Restore operation

b. **Explain with a figure, the two different types of backup topologies.**

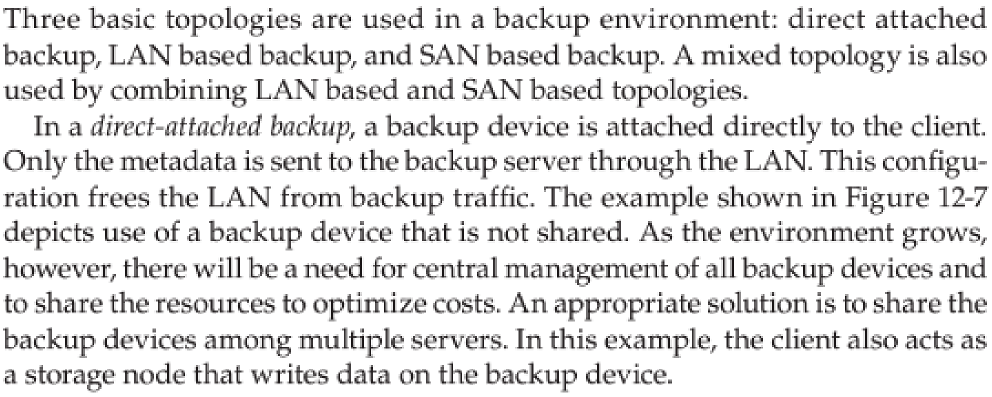

Three basic topologies are used in a backup environment: direct attached backup, LAN based backup, and SAN based backup. A mixed topology is also used by combining LAN based and SAN based topologies.

In a *direct-attached backup*, a backup device is attached directly to the client. Only the metadata is sent to the backup server through the LAN. This configuration frees the LAN from backup traffic. The example shown in Figure 12-7 depicts use of a backup device that is not shared. 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.

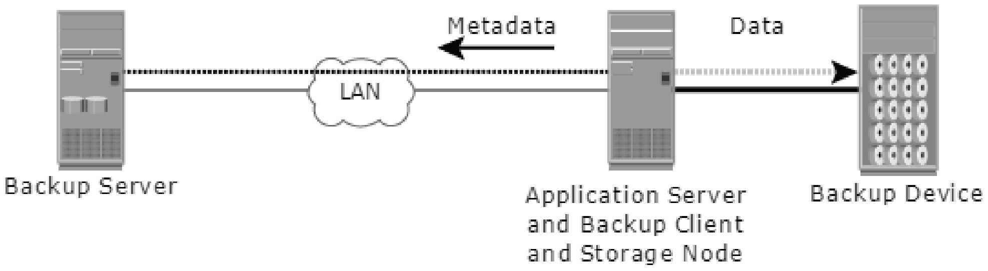

Figure 12-7: Direct-attached backup topology

In LAN-based backup, all servers are connected to the LAN and all storage devices are directly attached to the storage node (see Figure 12-8). 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. Network resources are severely constrained when multiple clients access and share the same tape library unit (TLU). 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.

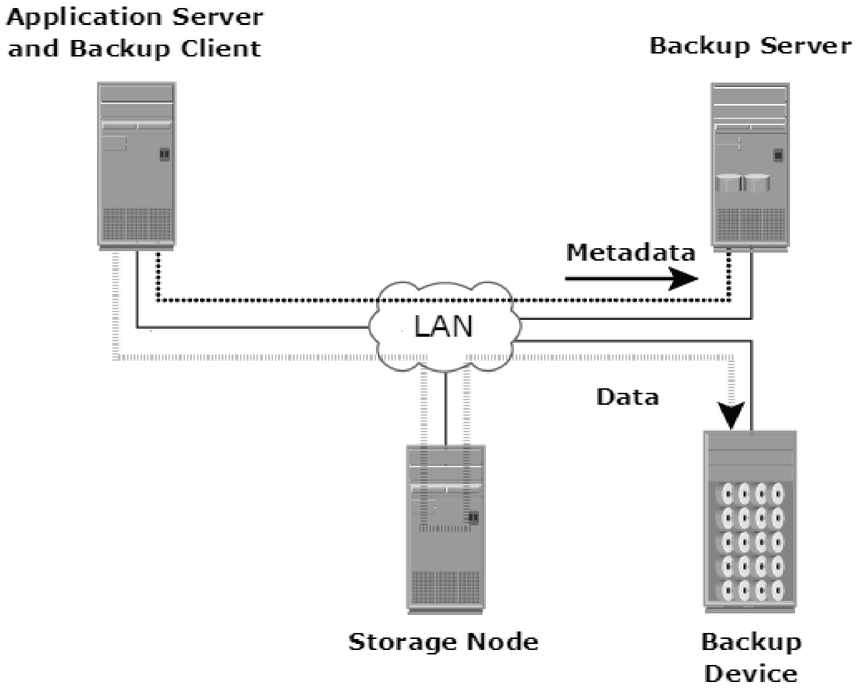

Figure 12-8: LAN-based backup topology

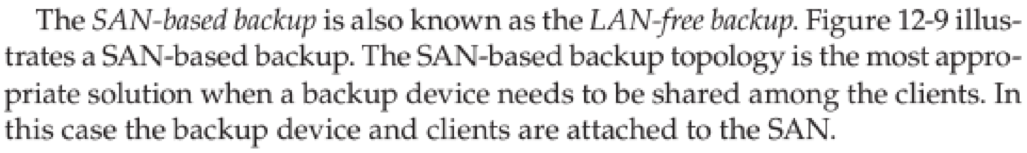

The *SAN-based backup* is also known as the *LAN-free backup*. Figure 12-9 illustrates a SAN-based backup. The SAN-based backup topology is the most appropriate solution when a backup device needs to be shared among the clients. In this case the backup device and clients are attached to the SAN.

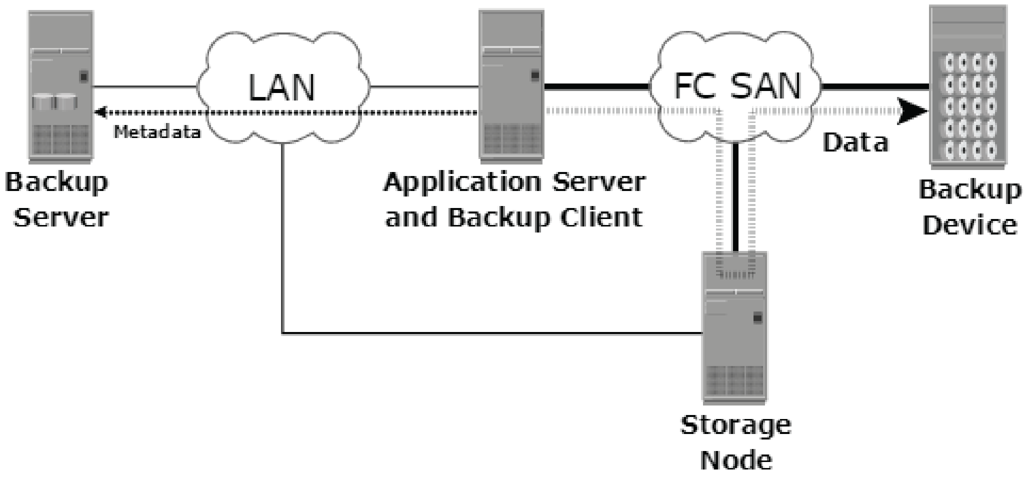

Figure 12-9: SAN-based backup topology

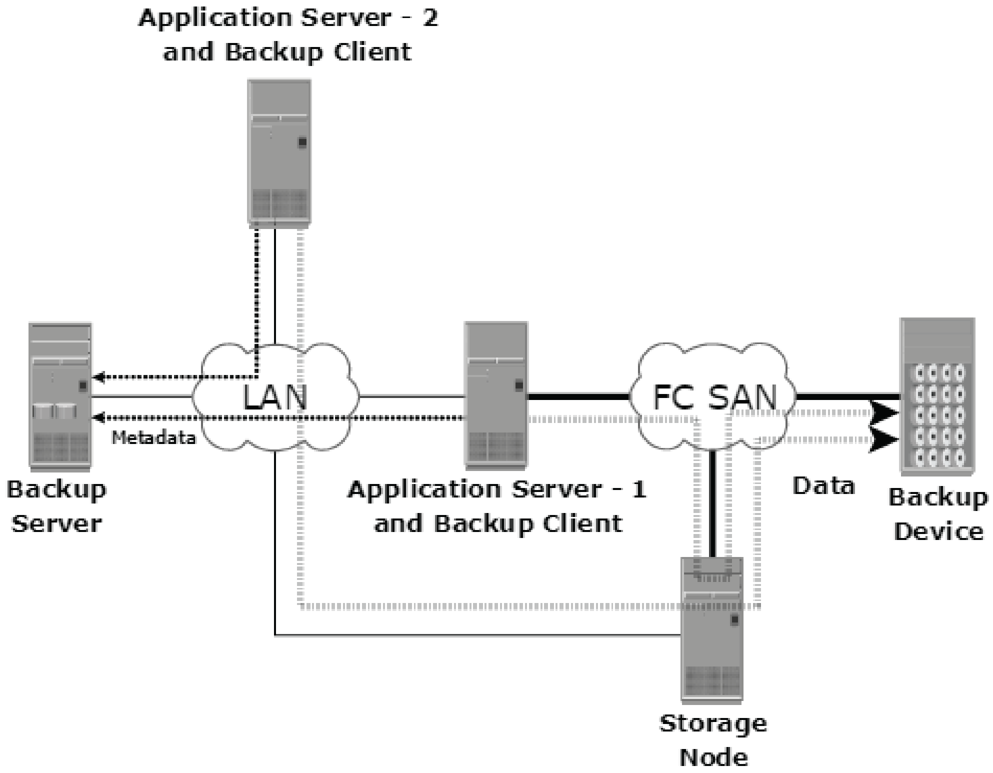

Figure 12-10: Mixed backup topology

9. a. Explain the various terms used to represent entities and operations in a replication environment.

Replication involves copying data from one storage location to another to ensure availability and redundancy. Here are the key terms:

1. **Source/Primary**: The system where the original data resides.
2. **Target/Replica**: The system where the copy of the data is stored.
3. **Consistency Group**: A collection of data that must remain consistent across the source and target.
4. **Write Splitter**: A mechanism that captures write operations and sends them to the replica storage.
5. **RPO (Recovery Point Objective)**: Defines the amount of data loss that can be tolerated (time between backups).
6. **RTO (Recovery Time Objective)**: Defines the time it takes to recover the data.

b. Explain, with a figure, the storage array-based local replication.

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. Consequently, the host is not burdened by the replication operations. The replica can be accessed by an alternate host for any business operations. In this replication, the required number of replica devices should be selected on the same array and then data is replicated between source-replica pairs. A database could be laid out over multiple physical volumes and in that case all the devices must be replicated for a consistent PIT copy of the database. Figure 13-5 shows storage array based local replication, where source and target are in the same array and accessed by different hosts. Storage array-based local replication can be further categorized as full-volume mirroring, pointer-based full-volume replication, and pointer-based virtual replication. Replica devices are also referred as target devices, accessible by business continuity host.

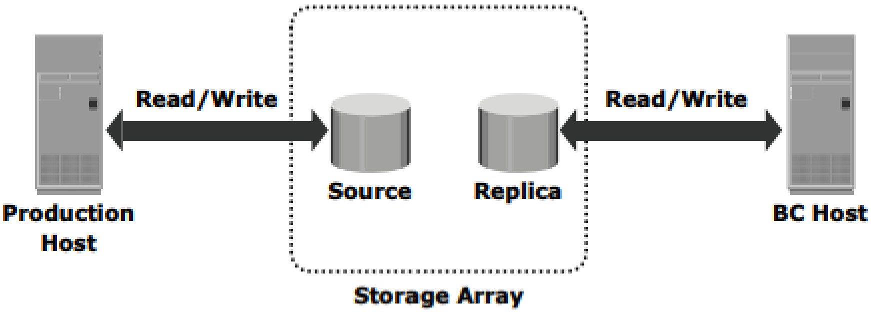

Figure 13-5: Storage array-based replication

c. Illustrate, with a figure, synchronous and asynchronous replication.

Synchronous and Asynchronous Replication are two important methods used in data replication environments to ensure data availability, consistency, and protection in the event of failures or disasters. Here's an explanation and a general illustration of both methods:

1. Synchronous Replication

- **Definition**: In synchronous replication, data is written to both the source (primary storage) and the target (replica) simultaneously, ensuring that both systems always contain the same data in real-time.
- **Process**:
  - A write operation is initiated at the source.

    - The write operation is sent to both the source and target systems.
    - The source waits for an acknowledgment from the target, ensuring that data is successfully written to both systems before completing the write operation.
    - This ensures zero or minimal data loss, as both systems are always in sync.
- **Advantages**:
    - Ensures data consistency with no data loss (RPO = 0).
    - Ideal for critical applications where real-time data synchronization is necessary.
- **Disadvantages**:
    - Higher latency due to waiting for acknowledgments from the target.
    - Requires high bandwidth and low-latency network connections.

2. Asynchronous Replication

- **Definition**: In asynchronous replication, data is written to the source system first, and the write operation is then queued for replication to the target system at a later time. There may be a time lag between the source and the target systems.
- **Process**:
    - A write operation is initiated at the source.
    - The source completes the write operation without waiting for the target.
    - The write operations are then queued and sent to the target asynchronously, which introduces a potential time lag between the source and target.
    - The target will eventually catch up with the source, but there may be some data that is yet to be replicated in case of a failure.
- **Advantages**:
    - Reduces latency and is less dependent on network bandwidth.
    - More suitable for long-distance replication where network latency is higher.
- **Disadvantages**:
    - Potential data loss during a failure, as there is a time gap (RPO > 0).
    - The target system may not always have the most up-to-date data.

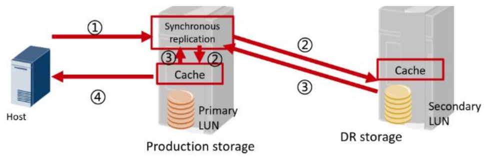

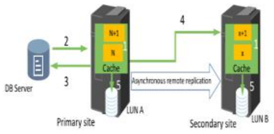

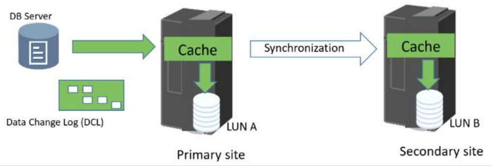

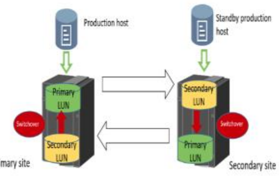

10.a. Illustrate, with a figure, synchronous and asynchronous array-based remote replication.

In storage array-based remote replication, the array operating environment and resources perform and manage data replication. This relieves the burden on the host CPUs, which can be better utilized for running an application. A source and its replica device reside on different storage arrays. In other implementations, the storage controller is used for both the host and replication workload. Data can be transmitted from the source storage array to the target storage array over a shared or a dedicated network. Replication between arrays may be performed in synchronous, asynchronous, or disk-buffered modes. Three-site remote replication can be implemented using a combination of synchronous mode and asynchronous mode, as well as a combination of synchronous mode and disk-buffered mode. Synchronous Replication Mode In array based synchronous remote replication, writes must be committed to the source and the target prior to acknowledging "write complete" to the host. Additional writes on that source cannot occur until each preceding write has been completed and acknowledged. The array-based synchronous replication process is shown in Figure 14-5.

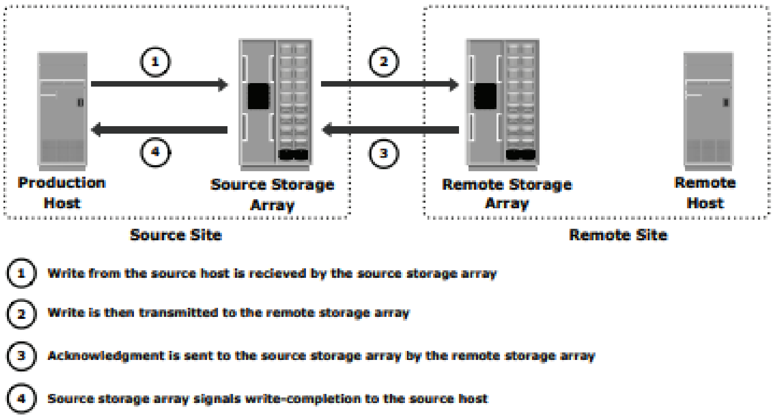
1
2
4
3
Production Host
Source Storage Array
Remote Storage Array
Remote Host
Source Site
Remote Site
1 Write from the source host is recieved by the source storage array
2 Write is then transmitted to the remote storage array
3 Acknowledgment is sent to the source storage array by the remote storage array
4 Source storage array signals write-completion to the source host

Figure 14-5: Array-based synchronous remote replication

In the case of synchronous replication, to optimize the replication process and to minimize the impact on application response time, the write is placed on cache of the two arrays. The intelligent storage arrays can de-stage these writes to the appropriate disks later. If the network links fail, replication is suspended; however, production work can continue uninterrupted on the source storage array. The array operating environment can keep track of the writes that are not transmitted to the remote storage array. When the network links are restored, the accumulated data can be transmitted to the remote storage array. During the time of network link outage, if there is a failure at the source site, some data will be lost and the RPO at the target will not be zero. For synchronous remote replication, network bandwidth equal to or greater than the maximum write workload between the two sites should be provided at all times. Figure 14-6 illustrates the write workload (expressed in MB/s) over time. The "Max" line indicated in Figure 14-6 represents the required bandwidth that must be provisioned for synchronous replication. Bandwidths lower than the maximum write workload results in an unacceptable increase in application response time

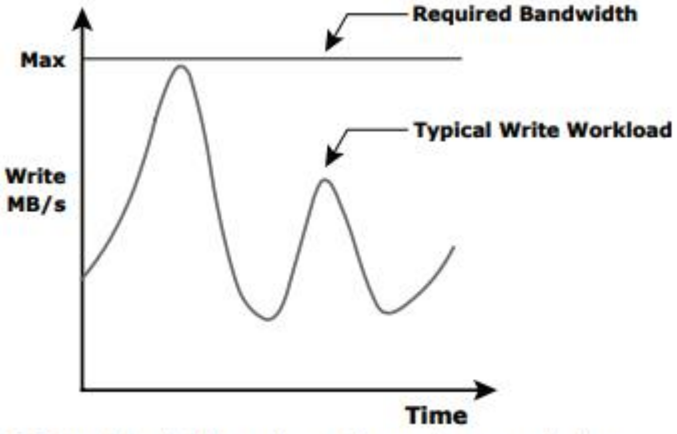

Figure 14-6: Network bandwidth requirement for synchronous replication

Asynchronous Replication Mode: In array-based asynchronous remote replication mode, shown in Figure 14-7, a write is committed to the source and immediately acknowledged to the host. Data is buffered at the source and transmitted to the remote site later. The source and the target devices do not contain identical data at all times. The data on the target device is behind that of the source, so the RPO in this case is not zero.

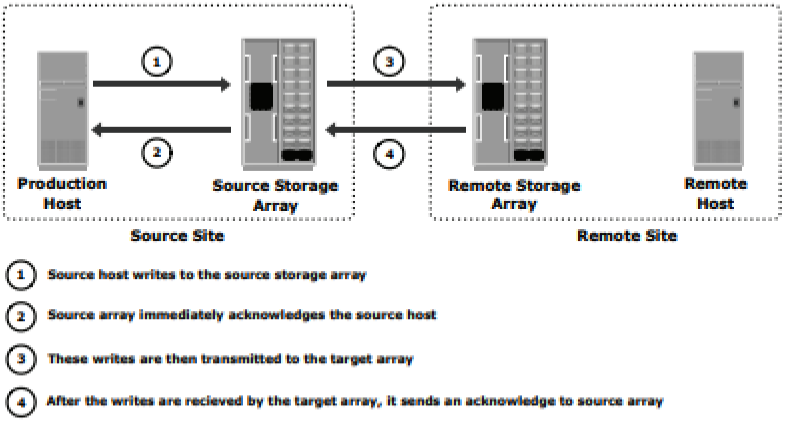

1 Source host writes to the source storage array

2 Source array immediately acknowledges the source host

3 These writes are then transmitted to the target array

4 After the writes are recieved by the target array, it sends an acknowledge to source array

Figure 14-7: Array-based asynchronous remote replication

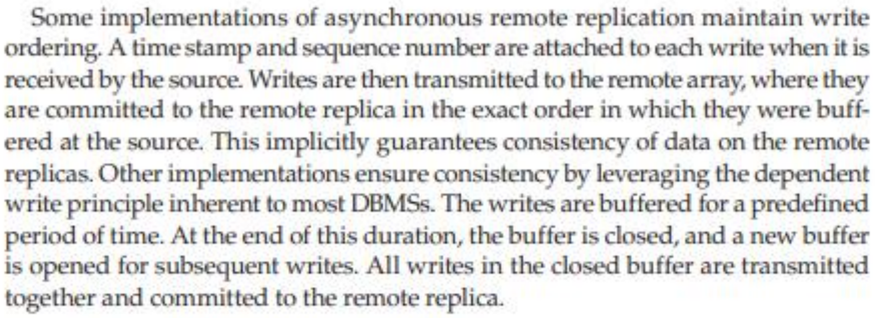

Some implementations of asynchronous remote replication maintain write ordering. A time stamp and sequence number are attached to each write when it is received by the source. Writes are then transmitted to the remote array, where they are committed to the remote replica in the exact order in which they were buffered at the source. This implicitly guarantees consistency of data on the remote replicas. Other implementations ensure consistency by leveraging the dependent write principle inherent to most DBMSs. The writes are buffered for a predefined period of time. At the end of this duration, the buffer is closed, and a new buffer is opened for subsequent writes. All writes in the closed buffer are transmitted together and committed to the remote replica.

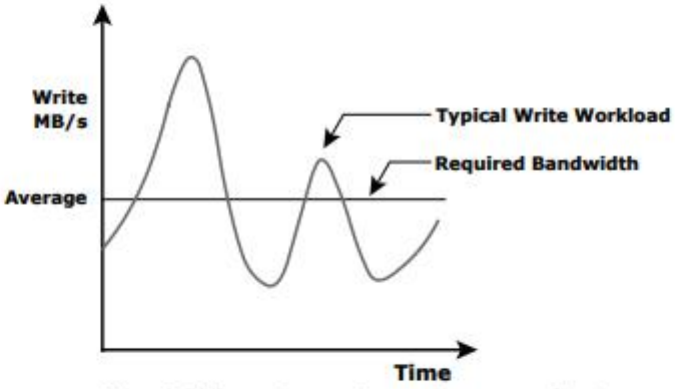

Figure 14-8: Network bandwidth requirement for asynchronous replication

b. Discuss the four security goals achieved through the information security framework.

The basic security framework is built around the four primary services of security: accountability, confidentiality, integrity, and availability. This framework incorporates all security measures required to mitigate threats to these four primary security attributes:

■ Accountability service: Refers to accounting for all the events and operations that takes place in data center infrastructure. The accountability service maintains a log of events that can be audited or traced later for the purpose of security.

■ Confidentiality service: Provides the required secrecy of information and ensures that only authorized users have access to data. This service authenticates users who need to access information and typically covers both data in transit (data transmitted over cables), or data at rest (data on a backup media or in the archives). Data in transit and at rest can be encrypted to maintain its confidentiality. In addition to restricting unauthorized users from accessing information, confidentiality services also implement traffic flow protection measures as part of the security protocol. These protection measures generally include hiding source and destination addresses, frequency of data being sent, and amount of data sent.

■ Integrity service: Ensures that the information is unaltered. The objective of the service is to detect and protect against unauthorized alteration or deletion of information. Similar to confidentiality services, integrity services work in collaboration with accountability services to identify and authenticate the users. Integrity services stipulate measures for both in-transit data and at-rest data.

■ Availability service: This ensures that authorized users have reliable and timely access to data. These services enable users to access the required computer systems, data, and applications residing on these systems. Availability services are also implemented on communication systems used to transmit information among computers that may reside at different locations. This ensures availability of information if a failure in one particular location occurs. These services must be implemented for both electronic data and physical data.

c. Explain, with a figure, the storage security domains.

Storage devices that are not connected to a storage network are less vulnerable because they are not 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. This requires a closer look at storage networking security and a clear understanding of the access paths leading to storage resources. If a particular path is unauthorized and needs to be prohibited by technical controls, one must ensure that these controls are not compromised. If each component within the storage network is considered a potential access point, one must analyze the attack surface that each of these access points provides and identify the associated vulnerability. In order 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 BURA (backup, recovery, and archive). Figure 15-1 depicts the three security domains of a storage system environment. The first security domain involves application access to the stored data through the storage network. The second security domain includes management access to storage and interconnect devices and to the data residing on those devices. This domain is primarily accessed by storage administrators who configure and manage the environment. The third domain consists of BURA access. Along with the access points in the other two domains, backup media also needs to be secured. 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.

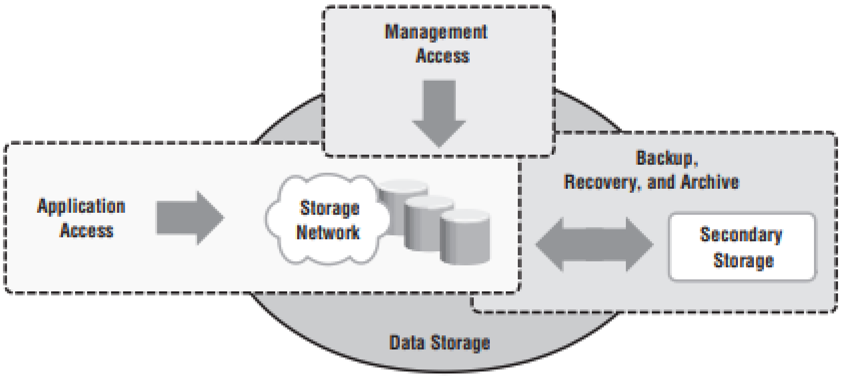

Figure 15-1: Three security domains of data storage