CMR
INSTITUTE OF
TECHNOLOGY

Dr. VAKULA RANI J

Sub:	**Storage Area Networks**	**Sub Code:**	**18MCA554**

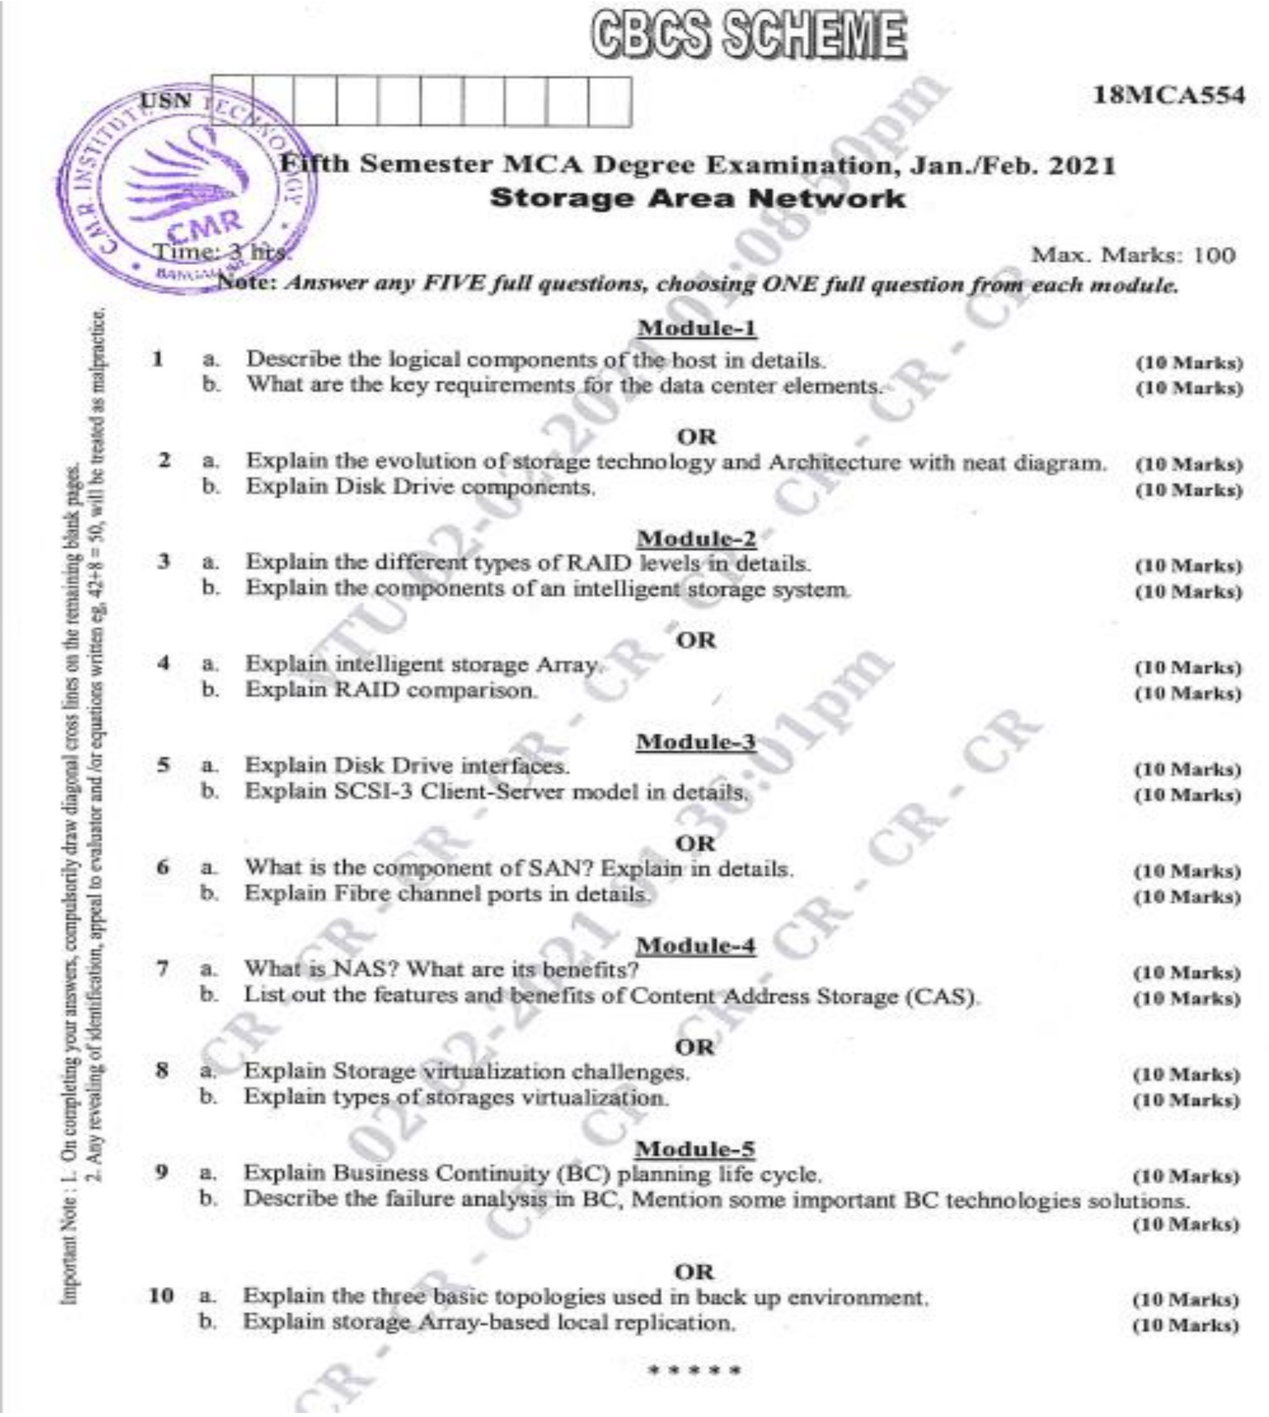

CBCS SCHEME

USN □□□□□□□□□□

18MCA554

Fifth Semester MCA Degree Examination, Jan./Feb. 2021

Storage Area Network

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. Describe the logical components of the host in details. **(10 Marks)**
b. What are the key requirements for the data center elements. **(10 Marks)**

OR

2 a. Explain the evolution of storage technology and Architecture with neat diagram. **(10 Marks)**
b. Explain Disk Drive components. **(10 Marks)**

Module-2

3 a. Explain the different types of RAID levels in details. **(10 Marks)**
b. Explain the components of an intelligent storage system. **(10 Marks)**

OR

4 a. Explain intelligent storage Array. **(10 Marks)**
b. Explain RAID comparison. **(10 Marks)**

Module-3

5 a. Explain Disk Drive interfaces. **(10 Marks)**
b. Explain SCSI-3 Client-Server model in details. **(10 Marks)**

OR

6 a. What is the component of SAN? Explain in details. **(10 Marks)**
b. Explain Fibre channel ports in details. **(10 Marks)**

Module-4

7 a. What is NAS? What are its benefits? **(10 Marks)**
b. List out the features and benefits of Content Address Storage (CAS). **(10 Marks)**

OR

8 a. Explain Storage virtualization challenges. **(10 Marks)**
b. Explain types of storages virtualization. **(10 Marks)**

Module-5

9 a. Explain Business Continuity (BC) planning life cycle. **(10 Marks)**
b. Describe the failure analysis in BC, Mention some important BC technologies solutions. **(10 Marks)**

OR

10 a. Explain the three basic topologies used in back up environment. **(10 Marks)**
b. Explain storage Array-based local replication. **(10 Marks)**

* * * * *

1.a)Describe the Logical components of the Host.

Sol : Applications runs on hosts or Servers. Hosts can range from simple laptops to complex server clusters

i) **Logical Components of the Host**

- **Application**
- **Operating system**
- **Logical Volume Managers**

a) **Application**

Applications is an Interface between user and the host . It has Three-tiered architecture - Application UI, computing logic and underlying databases. Application data access can be classifies as: **Block-level access**: Data stored and retrieved in blocks, specifying the LBA**File-level access**: Data stored and retrieved by specifying the name and path of files.

b) Operating system

Resides between the applications and the hardware. It Controls the environment.

c) Logical Volume Managers

Logical data blocks are mapped to physical data blocks. LVM Components are Physical Volumes and Volume Groups .One or more Physical Volumes form a Volume Group. LVM manages Volume Groups as a single entity. Physical Volumes can be added and removed from a Volume Group as necessary. Physical Volumes are typically divided into contiguous equal-sized disk blocks. A host will always have at least one disk group for the Operating System. Application and Operating System data maintained in separate volume groups.

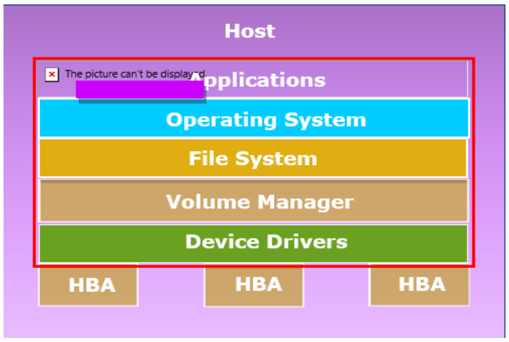

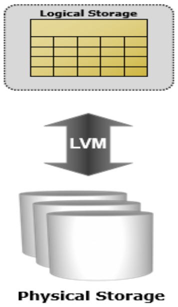

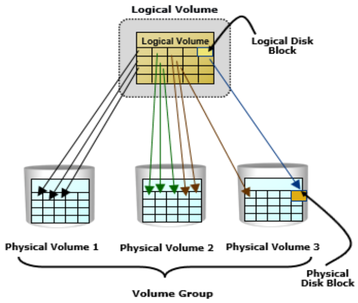

1.b) What are the key requirements of Data Center

Data Center : It is a facility that contains storage , compute network, and other IT resources to provide centralized data – processing capabilities.

KEY ELEMENTS OF A DATA CENTER

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

- ➢ **Application**: A computer program that provides the logic for computing operations.
- ➢ **Database management system (DBMS):** Provides a structured way to store data in logically organized tables that are interrelated.
- ➢ **Host or compute:** A computing platform (hardware, firmware, and software) that runs applications and databases.
- ➢ **Network:** A data path that facilitates communication among various networked devices.
- ➢ **Storage:** A device that stores data persistently for subsequent use
- ➢ The figure below shows an example of an online order transaction system that involves the five core elements of a data center and illustrates their functionality in a business process.

KEY CHARACTERISTICS OF A DATA CENTER

- ➢ Uninterrupted operation of data centers is critical to the survival and success of a business.

- Organizations must have a reliable infrastructure that ensures that data is accessible at all times.
- The following are the key characteristics of Data Center

- **Availability**: A data center should ensure the availability of information when required. Unavailability of information could cost millions of dollars per hour to businesses, such as financial services, telecommunications, and e-commerce.

- **Security**: Data centers must establish policies, procedures, and core element integration to prevent unauthorized access to information.

- **Scalability**: Business growth often requires deploying more servers, new applications, and additional databases. Data center resources should scale based on requirements, without interrupting business operations.

- **Performance**: All the elements of the data center should provide optimal performance based on the required service levels.

- **Data integrity**: Data integrity refers to mechanisms, such as error correction codes or parity bits, which ensures that data is stored and retrieved 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 provide additional capacity without interrupting availability or with minimal disruption. Capacity may be managed by reallocating the existing resources or by adding new resources.

- **Manageability**: A data center should provide easy and integrated management of all its elements. Manageability can be achieved through automation and reduction of human (manual) intervention in common tasks.

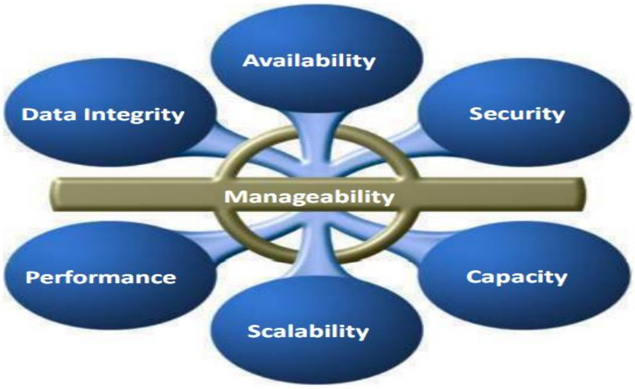

Fig : Key characteristics of a Data Center

2.a) Write a note on Evolution Storage Technology and Architecture

i) **Direct-attached storage (DAS)**: 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.

ii)**Redundant Array of Independent Disks (RAID):** To address the cost, performance, and availability requirements of data.

iii)**Storage area network (SAN):** A dedicated, high-performance Fibre Channel (FC) network to facilitate block-level communication between servers and storage. SAN offers scalability, availability, performance, and cost benefits compared to DAS.

iv)**Network-attached storage (NAS):** A dedicated storage for file serving applications and connects to LAN network and provides file access to heterogeneous clients. it offers higher scalability, availability, performance, and cost benefits

v)**Internet Protocol SAN (IP-SAN):** IP-SAN is a convergence of technology used in SAN and NAS.

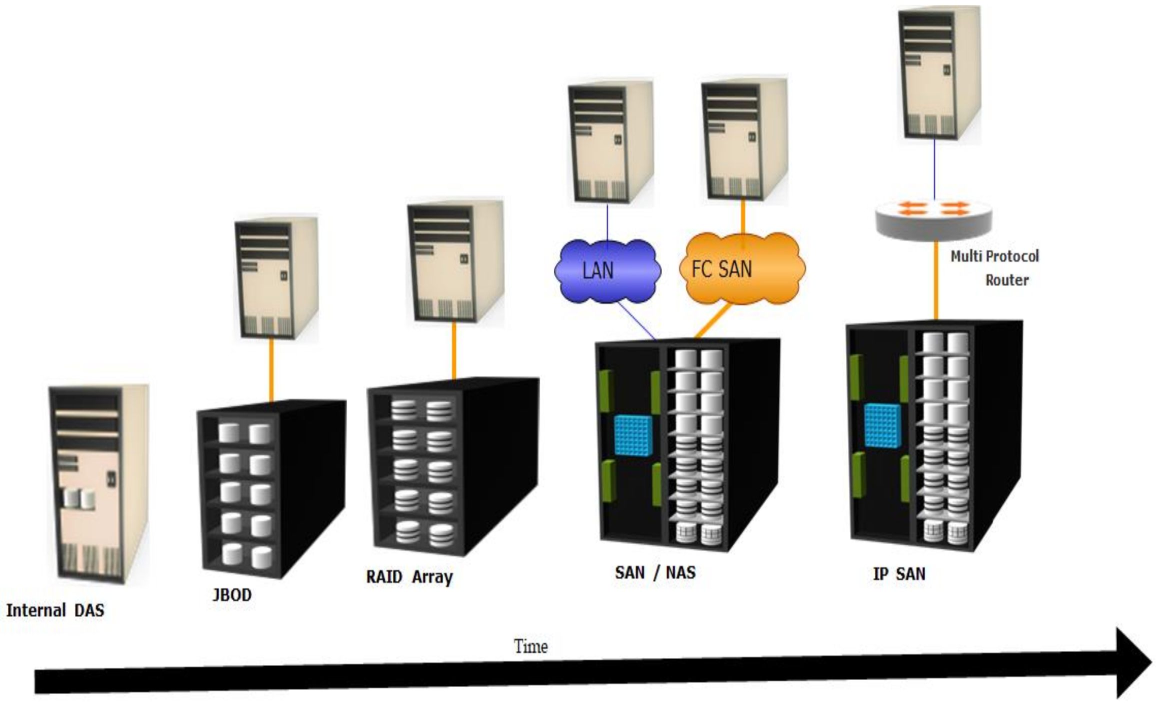

2.b) **Explain disk drive components**

Sol:

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.

Data on the disk is recorded on tracks, which are concentric rings on the platter around the spindle
The tracks are numbered, starting from zero, from the outer edge of the platter. Each track is divided into smaller units called sectors. A sector is the smallest, individually addressable unit of storage.
The key components of disk drive are : i) Platter ii) Spindle iii) Read/Write head iv) actuator arm assembly v) controller

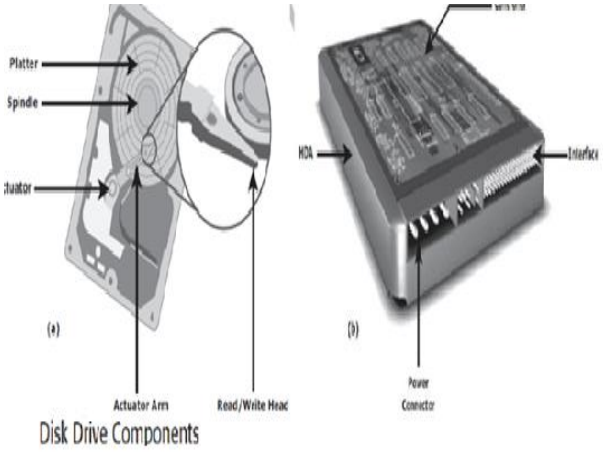

Disk Drive Components

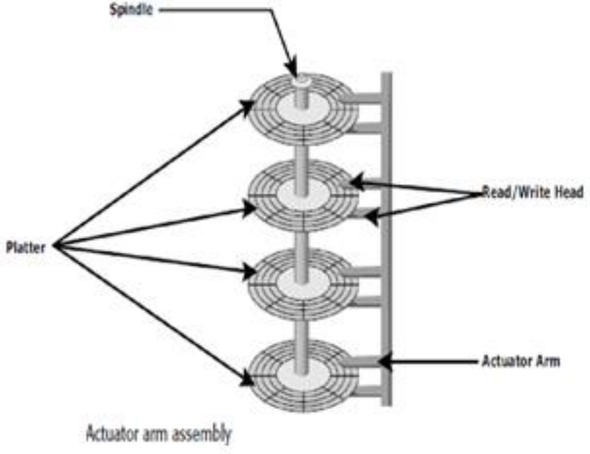

Actuator arm assembly

Platter :

- HDD consists of one or more flat circular disks called platters
- The data is recorded on these platters in binary (0s &1s).
- The set of rotating platters is sealed in a case, called a
- Head Disk Assembly (HDA).
- The data is encoded by polarizing the magnetic area, of the disk surface.
- Data can be written to or read from both surfaces of the platter

Spindle:

- A spindle connects all the platters and to a motor.
- The motor of the spindle rotates with a constant speed.
- The disk platter spins at a speed of several thousands of revolutions per minute (rpm).
- Disk drives have spindle speeds of 7,200 rpm, 10,000 rpm, or 15,000 rpm.
- Platter - diameter of 3.5‖ (90 mm).

Spindle:

- A spindle connects all the platters and to a motor.
- The motor of the spindle rotates with a constant speed.
- The disk platter spins at a speed of several thousands of revolutions per minute (rpm).
- Disk drives have spindle speeds of 7,200 rpm, 10,000 rpm, or 15,000 rpm.
- platter - diameter of 3.5‖ (90 mm).

Read/Write head:

- Drives have two R/W heads per platter to Read and write data one for each surface .
- when writing data- changes the magnetic polarization
- while reading data- detects magnetic polarization
- ensures that heads are moved to the landing zone before they touch the surface.
- If the R/W head accidentally touches the surface of the platter outside the landing zone, a **head crash .**

Actuator Arm Assembly:

- The R/W heads are mounted on the actuator arm assembly
- To read or write data, R/W heads positions at the location on the platter
- The R/W heads for all platters are attached to one actuator arm assembly and move across the platters simultaneously.

Controller:

- The controller is a printed circuit board, mounted at the bottom of a disk drive.
- It consists of a microprocessor, internal memory, circuitry, and firmware.
- The firmware controls power to the spindle motor and the speed of the motor.
- It also manages communication between the drive and the host.

3.a) **Explain different types of Raid levels**

Sol :

- **RAID 0** Striped array with no fault tolerance
- **RAID 1** Disk mirroring
- **Nested RAID** (**Raid 1 + 0, Raid 0 + 1**)
- **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

Most data centers require data redundancy and performance from their RAID arrays.

RAID 0+1 and **RAID 1+0** combine the **performance benefits of RAID 0** with the **redundancy benefits of RAID 1**. They use striping and mirroring techniques and combine their benefits. These types of RAID require an even number of disks, the minimum being four (see Figure 3-7).

RAID 1+0 is also known as RAID 10 (Ten) or RAID 1/0. Similarly, RAID 0+1 is also known as RAID 01 or RAID 0/1.

RAID 1+0 performs well for workloads that use small, random, write-intensive I/O. Some applications that benefit from RAID 1+0 include the following:

1. High transaction rate Online Transaction Processing (OLTP)
2. Large messaging installations
3. Database applications that require high I/O rate, random access, and high availability.

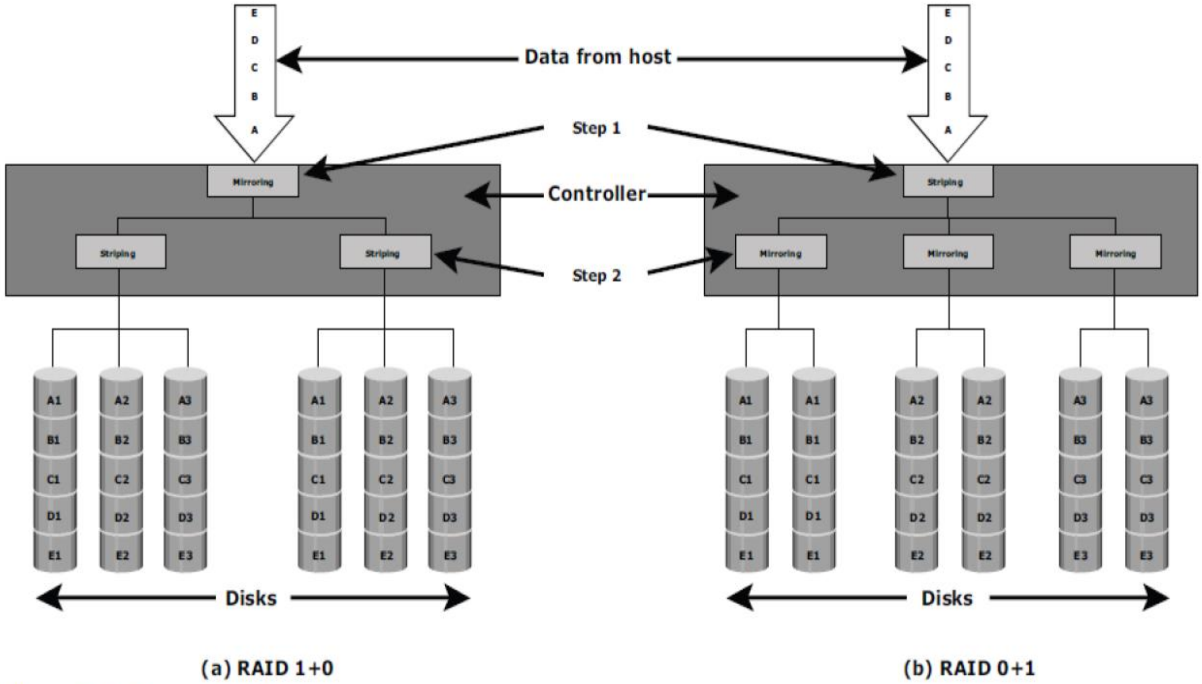

Figure 3-7: Nested RAID

A common misconception is that RAID 1+0 and RAID 0+1 are the same. **RAID 1+0 is also called striped mirror.** The basic element of RAID 1+0 is a mirrored pair, which means that data is first mirrored and then both copies of data are striped across multiple HDDs in a RAID set. When replacing a failed drive, only the mirror is rebuilt, i.e. the disk array controller uses the surviving drive in the mirrored pair for data recovery and continuous operation. Data from the surviving disk is copied to the replacement disk.

RAID 0+1 is also called mirrored stripe. The basic element of RAID 0+1 is a stripe. This means that the process of striping data across HDDs is performed initially and then the entire stripe is mirrored. If one drive fails, then the entire stripe is faulted. A rebuild operation copies the entire stripe, copying data from each disk in the healthy stripe to an equivalent disk in the failed stripe.

Disadv: This causes increased and unnecessary I/O load on the surviving disks and makes the RAID set more vulnerable to a second disk failure.

RAID 5

- RAID 5 is a very versatile RAID implementation. it uses striping and the drives are independently accessible.
- In RAID 5, parity is distributed across all disks. The distribution of parity in RAID 5 overcomes the write bottleneck.
- write I/O operations suffer performance degradation because of the write penalty that manifests with a parity RAID implementation
- The performance degradation also occurs during recovery and reconstruction operations in the event of a disk failure. In addition, multiple disk failures within the array may result in data loss.
- RAID 5 is preferred for messaging, data mining, medium-performance media serving, and relational database management system (RDBMS) implementations in which database administrators (DBAs) optimize data access.

ii)RAID 6

- RAID 6 includes a second parity element to enable survival in the event of the failure of two disks in a RAID group
- Therefore, a RAID 6 implementation requires at least four disks.
- RAID 6 distributes the parity across all the disks.
- The write penalty in RAID 6 is more than that in RAID 5; therefore, RAID 5 writes perform better than RAID 6.
- The rebuild operation in RAID 6 may take longer than that in RAID 5 due to the presence of two parity sets.

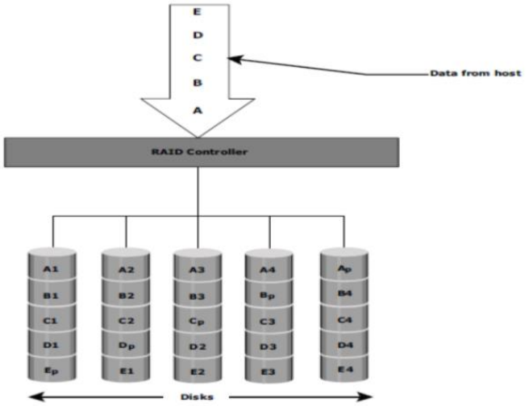

Fig : RAID-5

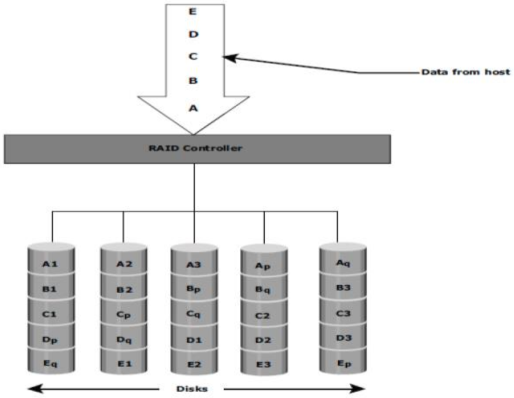

Fig : RAID-6

3.b)) Explain components of an Intelligent Storage System

An intelligent storage system consists of four key components:

i) **Front End ii) Cache iii) Back End iv) Physical Disks.**

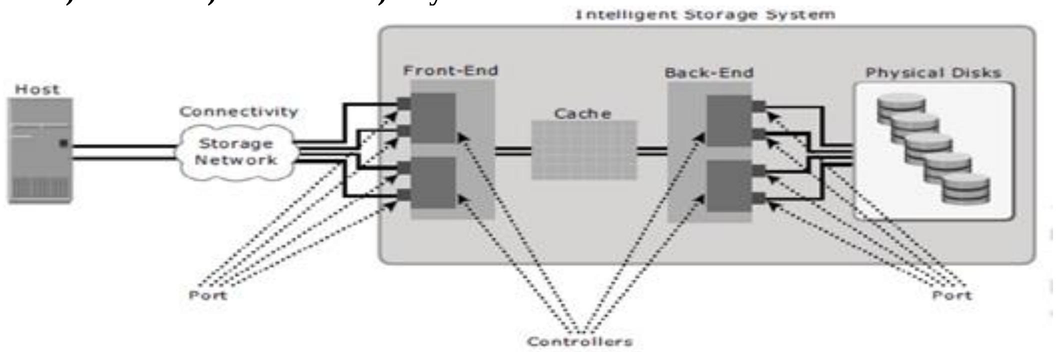

Components of an intelligent storage system

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.

i) Front End

The front end provides the interface between the storage system and the host. It consists of two components: **front-end ports** and **front-end controllers**. **Front-end ports:** enable hosts to connect to the intelligent storage system. Each front-end port has processing logic that executes the appropriate transport protocol, such as SCSI, Fibre Channel, or iSCSI, for storage connections. Redundant ports are provided on the front end for high availability.

Front-end controllers: route data to and from cache via the internal data bus.

ii) Cache

Cache is an important component that enhances the I/O performance in an intelligent storage system. Cache is a memory where data is placed temporarily to reduce the time required to service I/O requests from the host. Accessing data from cache takes less than a millisecond.

Read Operation with Cache: When a host issues a read request, the front-end controller accesses the tag RAM to determine whether the required data is available in cache. If the requested data is found in the cache, it is called a **read cache hit** or **read hit** and data is sent directly to the host, without any disk operation. This provides a fast response time to the host. If the requested data is not found in cache, it is called a **read cache miss** or **read miss** and the data must be read from the disk. Cache misses increase I/O response time.

Write-back cache: Data is placed in cache and an acknowledgment is sent to the host immediately. Later, data from several writes are committed (de-staged) to the disk. Write response times are much faster, as the write operations are isolated from the mechanical delays of the disk. However, uncommitted data is at risk of loss in the event of cache failures.

Write-through cache: Data is placed in the cache and immediately written to the disk, and an acknowledgment is sent to the host. Because data is committed to disk as it arrives, the risks of data loss are low but write response time is longer because of the disk operations.

iii) Back End

The back end provides an interface between cache and the physical disks. It consists of two components: **back-end ports** and **back-end controllers.** The back end controls data transfers between cache and the physical disks. From cache, data is sent to the back end and then routed to the destination disk. Physical disks are connected to ports on the back end. The back end controller communicates with the disks when performing reads and writes.

iv) Physical Disk

A physical disk stores data persistently. Disks are connected to the back-end with either SCSI or a Fibre Channel interface (discussed in subsequent chapters). An intelligent storage system enables the use of a mixture of SCSI or Fibre Channel drives and IDE/ATA drives.

Logical Unit Number: Physical drives or groups of RAID protected drives can be logically split into volumes known as logical volumes, commonly referred to as **Logical Unit Numbers (LUNs)**. The use of LUNs improves disk utilization. For example, without the use of LUNs, a host requiring only 200 GB could be allocated an entire 1TB physical disk. Using LUNs, only the required 200 GB would be allocated to the host, allowing the remaining 800 GB to be allocated to other hosts.

4(a) Explain an Intelligent Storage Array

Sol : An intelligent storage system consists of four key components: i) **Front End ii) Cache iii) Back End iv) Physical Disks.**

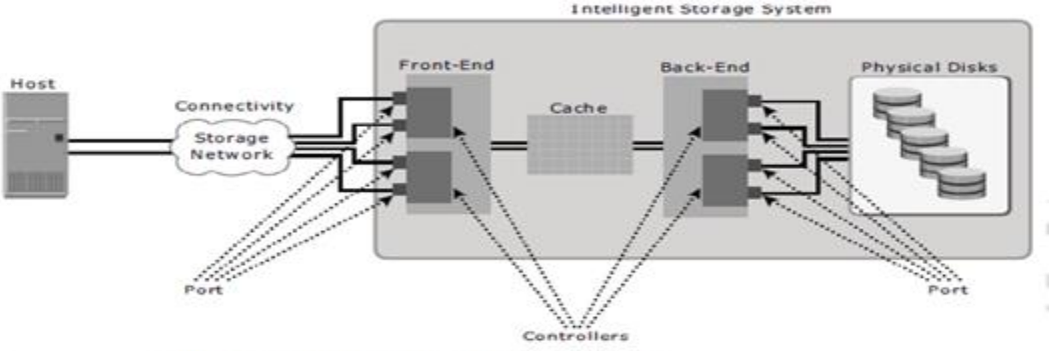

Components of an intelligent storage system

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.

i) Front End

The front end provides the interface between the storage system and the host. It consists of two components: **front-end ports** and **front-end controllers**. **Front-end ports:** enable hosts to connect to the intelligent storage system. Each front-end port has processing logic that executes the appropriate transport protocol, such as SCSI, Fibre Channel, or iSCSI, for storage connections. Redundant ports are provided on the front end for high availability.

Front-end controllers: route data to and from cache via the internal data bus.

ii) Cache

Cache is an important component that enhances the I/O performance in an intelligent storage system. Cache is a memory where data is placed temporarily to reduce the time required to service I/O requests from the host. Accessing data from cache takes less than a millisecond.

Read Operation with Cache: When a host issues a read request, the front-end controller accesses the tag RAM to determine whether the required data is available in cache. If the requested data is found in the cache, it is called a **read cache hit** or **read hit** and data is sent directly to the host, without any disk operation. This provides a fast response time to the host. If the requested data is not found in cache, it is called a **read cache miss** or **read miss** and the data must be read from the disk. Cache misses increase I/O response time.

iii) Back End

The back end provides an interface between cache and the physical disks. It consists of two components: **back-end ports** and **back-end controllers.** The back end controls data transfers between cache and the physical disks. From cache, data is sent to the back end and then routed to the destination disk. Physical disks are connected to ports on the back end. The back end controller communicates with the disks when performing reads and writes.

iv) Physical Disk

A physical disk stores data persistently. Disks are connected to the back-end with either SCSI or a Fibre Channel interface (discussed in subsequent chapters). An intelligent storage system enables the use of a mixture of SCSI or Fibre Channel drives and IDE/ATA drives.

Logical Unit Number: Physical drives or groups of RAID protected drives can be logically split into volumes known as logical volumes, commonly referred to as **Logical Unit Numbers (LUNs)**. The use of LUNs improves disk utilization. For example, without the use of LUNs, a host requiring only 200 GB could be allocated an entire 1TB physical disk. Using LUNs, only the required 200 GB would be allocated to the host, allowing the remaining 800 GB to be allocated to other hosts.

4(b) Explain Raid Comparison

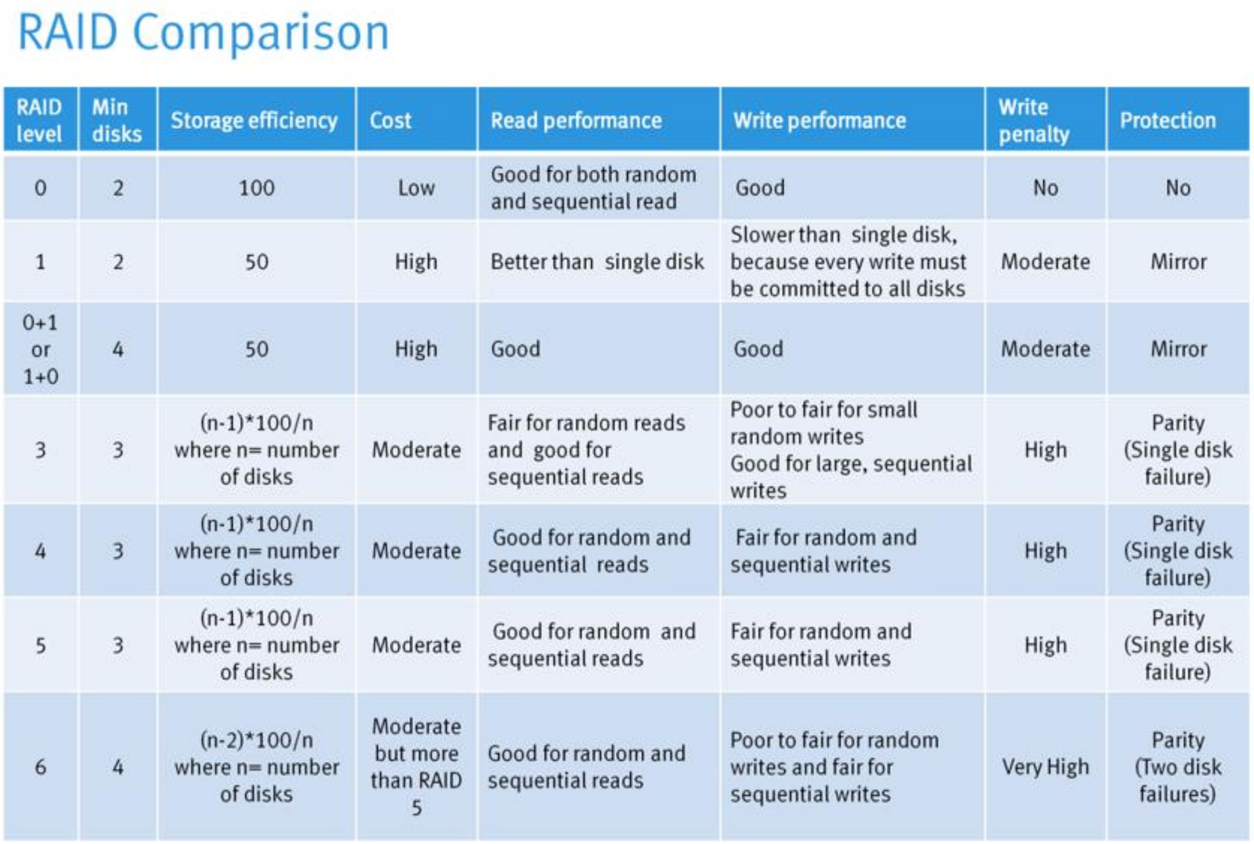

RAID Comparison

RAID level	Min disks	Storage efficiency	Cost	Read performance	Write performance	Write penalty	Protection
0	2	100	Low	Good for both random and sequential read	Good	No	No
1	2	50	High	Better than single disk	Slower than single disk, because every write must be committed to all disks	Moderate	Mirror
0+1 or 1+0	4	50	High	Good	Good	Moderate	Mirror
3	3	(n-1)*100/n where n= number of disks	Moderate	Fair for random reads and good for sequential reads	Poor to fair for small random writes Good for large, sequential writes	High	Parity (Single disk failure)
4	3	(n-1)*100/n where n= number of disks	Moderate	Good for random and sequential reads	Fair for random and sequential writes	High	Parity (Single disk failure)
5	3	(n-1)*100/n where n= number of disks	Moderate	Good for random and sequential reads	Fair for random and sequential writes	High	Parity (Single disk failure)
6	4	(n-2)*100/n where n= number of disks	Moderate but more than RAID 5	Good for random and sequential reads	Poor to fair for random writes and fair for sequential writes	Very High	Parity (Two disk failures)

5.b) SCSI-3 Client-Server Model

SCSI-3 architecture derives its base from the client-server relationship, in which a client directs a service request to a server, which then fulfills the client's request. In a SCSI environment, an initiator-target concept represents the client server model. In a SCSI-3 client-server model, a particular SCSI device acts as a SCSI target device, a SCSI initiator device, or a SCSI target/initiator device. Each device performs the following functions:

1. SCSI initiator device: Issues a command to the SCSI target device, to perform a task. A SCSI host adaptor is an example of an initiator.
2. SCSI target device: Executes commands to perform the task received from a SCSI initiator. Typically a SCSI peripheral device acts as a target device.
3. SCSI-3 client-server model, in which a SCSI initiator, or a client, sends a request to a SCSI target, or a server.
4. The target performs the tasks requested and sends the output to the initiator, using the protocol service interface. A SCSI target device contains one or more logical units. A logical unit has two components, a device server and a task manager
5. The logical unit processes the commands sent by a SCSI initiator. The SCSI initiator device is comprised of an application client and task management function, which initiates device service and task management requests.
6. Each device service request contains a Command Descriptor Block (CDB). The CDB defines the command to be executed and lists command-specific inputs and other parameters specifying how to process the command.

7. The SCSI devices are identified by a specific number called a SCSI ID.

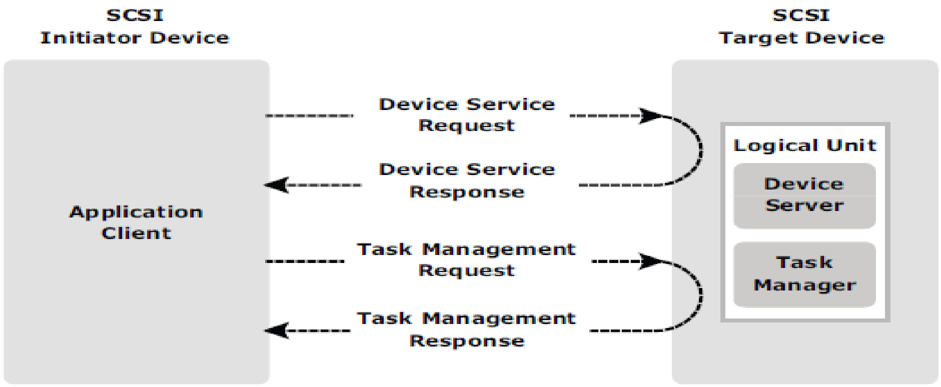

: SCSI-3 client-server model

6 a) Explain the components of FC SAN

Components of FC SAN

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

Node Ports

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

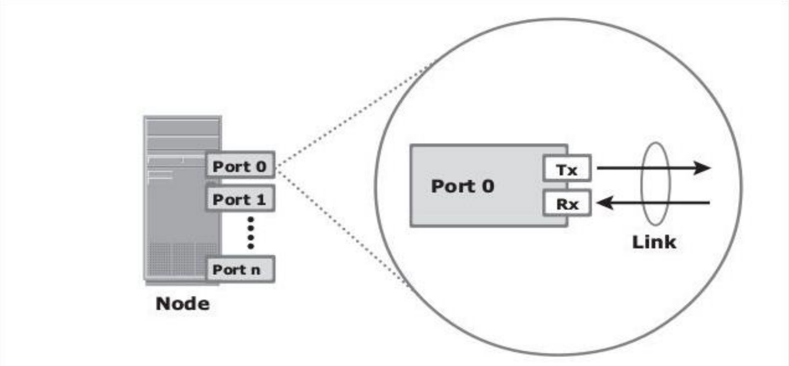

Figure 5-3: Nodes, ports, and links

Cables and Connectors

SAN implementations use optical fiber cabling. Copper can be used for shorter distances for back-end connectivity because it provides an acceptable signal-to-noise ratio for distances up to 30 meters. Optical fiber cables carry data in the form of light. There are two types of optical cables: multimode and single-mode. Multimode fiber (MMF) cable carries multiple beams of light projected at different angles simultaneously onto the core of the cable (see Figure 5-4 [a]). Based on the bandwidth, multimode fibers are classified as OM1 (62.5μm core), OM2 (50μm core), and laser-optimized OM3 (50μm core). In an MMF transmission, multiple light beams traveling inside the cable tend to disperse and collide. This collision weakens the signal strength after it travels a certain distance — a process known as modal dispersion. An MMF cable is typically used for short distances because of signal degradation (attenuation) due to modal dispersion. Single-mode fiber (SMF) carries a single ray of light projected at the center of the core (see Figure 5-4 [b]). These cables are available in core diameters of 7 to 11 microns; the most common size is 9 microns. In an SMF transmission, a single light beam travels in a straight line through the core of the fiber. The small core and the single light wave help to limit modal dispersion. Among all types of fiber cables, single-mode provides minimum signal attenuation over maximum distance (up to 10 km). A single-mode cable is used for long-distance cable runs, and distance usually depends on the power of the laser at the transmitter and sensitivity of the receiver.

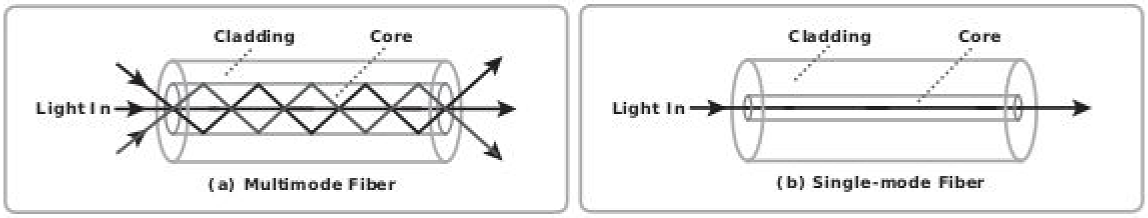

Figure 5-4: Multimode fiber and single-mode fiber

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

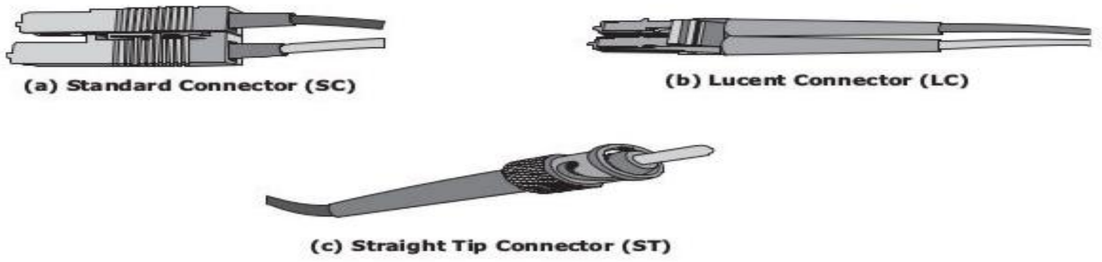
(a) Standard Connector (SC)
(b) Lucent Connector (LC)
(c) Straight Tip Connector (ST)

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

Interconnect Devices

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

Hubs are used as communication devices in FC-AL implementations. Hubs physically connect nodes in a logical loop or a physical star topology. All the nodes must share the loop because

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

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

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

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

SAN Management Software

SAN management software manages the interfaces between hosts, interconnect devices, and storage arrays. The software provides a view of the SAN environment and enables management of various resources from one central console.It provides key management functions, including mapping of storage devices, switches, and servers, monitoring and generating alerts for discovered devices, and zoning.

7.a) What are the benefits of NAS?

Benefits of NAS

- **Supports comprehensive access to information:** Enables efficient file sharing and supports many-to-one and one-to-many configurations. The many-to-one configuration enables a NAS device to serve many clients simultaneously.
- **Improved efficiency:** NAS uses an operating system specialized for file serving. It improves the utilization of general-purpose servers by relieving them of file-server operations.
- **Improved flexibility:** NAS is flexible and can serve requests from different types of clients from the same source.
- **Centralized storage:** Centralizes data storage to minimize data duplication on client workstations, simplify data management, and ensures greater data protection.
- **Simplified management:** Provides a centralized console that makes it possible to manage file systems efficiently.
- **Scalability:** Scales well in accordance with different utilization profiles and types of business applications because of the high performance and low-latency design.
- **High availability:** Offers efficient replication and recovery options, enabling high data availability. NAS uses redundant networking components that provide maximum connectivity options. A NAS device can use clustering technology for failover.
- **Security:** Ensures security, user authentication, and file locking in conjunction with industry-standard security schemas.

7.b) What is Content Addressed Storage? What are the benefits of CAS?

CAS is an object-based system that has been built for storing fixed content data. It is designed for secure online storage and retrieval of fixed content. CAS stores user data and its attributes as separate objects. The stored object is assigned a globally unique address known as a content address (CA). This address is derived from the object's binary representation.

Benefits of CAS

1. Content authenticity: The genuineness of stored content is achieved by generating a unique content address and automating the process for stored objects. Content authenticity is assured because the address assigned to each piece of fixed content is as unique as a fingerprint. Every time an object is read, CAS uses a hashing algorithm to recalculate the object's content address as a validation step and compares the result to its original content address.

2. Content integrity: Refers to the assurance that the stored content has not been altered. Use of hashing algorithm for content authenticity also ensures content integrity in CAS. If the fixed content is altered, CAS assigns a new address to the altered content, rather than overwrite the original fixed content, providing an audit trail and maintaining the fixed content in its original state.

3. Location independence: CAS uses a unique identifier that applications can leverage to retrieve data rather than a centralized directory, path names, or URLs. Using a content address to access fixed content makes the physical location of the data irrelevant to the application requesting the data. Therefore, the location from which the data is accessed is transparent to the application. This yields complete content mobility to applications across locations.

Single-instance storage (SiS): CAS provides an optimized and centrally managed storage solution that can support single-instance storage (SiS) to eliminate multiple copies of the same data.

4. Retention enforcement : Protecting and retaining data objects is a core requirement of an archive system. CAS creates two immutable components: a data object and a meta-object for every object stored. The metaobject stores object's attributes and data handling policies. For systems that support object-retention capabilities, the retention policies are enforced until the policies expire. Record-level protection and disposition

5. Technology independence: The CAS system interface is impervious to technology changes. As long as the application server is able to map the original content address the data remains accessible and ensure compatibility across platforms.

6. Fast record retrieval: CAS maintains all content on disks that provide subsecond **—**time to first byte (200 ms–400 ms) in a single cluster.

8.a) Explain Storage Virtualization Challenges

Storage Virtualization Challenges

Storage networking and feature-rich intelligent storage arrays have addressed and provided specific solutions to business problems. The storage virtualization solution must be capable of addressing issues such as scalability, functionality, manageability, and support.

1. **2Scalability**

- Consider the scalability of an environment with no virtualization. This environment may have several storage arrays that provide storage independently of each other. Each array is managed independently and meets application requirements in terms of IOPS and capacity.

- After virtualization, a storage array can no longer be viewed as an individual entity. The environment as a whole must now be analyzed. As a result, the infrastructure that is implemented both at a physical level and from a virtualization perspective must be able to adequately handle the workload, which may consist of different types of processing and traffic distribution. Greater care must be exercised to ensure that storage devices are performing to meet the appropriate requirements.

2. **Functionality**

- Functionality is another challenge in storage virtualization. Currently, the storage array provides a wide range of advanced functionality necessary for meeting an application's service levels. This includes local replication, extended-distance remote replication and the capability to provide application consistency across multiple volumes and arrays.

- In a virtualized environment, the virtual device must provide the same or better functionality than what is currently available on the storage array, and it must continue to leverage existing functionality on the arrays. It should protect the existing investments in processes, skills, training, and human resources.

i. **Manageability**

- The management of the storage infrastructure in a virtualized environment is an important consideration for storage administrators. A key advantage of today's storage resource management tools in an environment without virtualization is that they provide an end-to-end view, which integrates all the resources in the storage environment. They provide efficient and effective monitoring, reporting, planning, and provisioning services to the storage environment.

- Introducing a virtualization device breaks the end-to-end view into three distinct domains: the server to the virtualization device, the virtualization device to the physical storage, and the virtualization device itself. The virtualized storage environment must be capable of meeting these challenges and must integrate with existing management tools to enable management of an end-to-end virtualized environment.

ii. **Support**

- Virtualization is not a stand-alone technology but something that has to work within an existing environment. This environment may include multiple vendor technologies, such as switch and storage arrays, adding to complexity.

- Without a virtualization solution, many companies try to consolidate products from a single vendor to ease these challenges.

- Introducing a virtualization solution reduces the need to standardize on a single vendor. However, supportability issues in a virtualized heterogeneous environment introduce challenges in coordination and compatibility of products and solutions from different manufacturers and vendors.

8.b) Explain types of Storage Virtualization

Virtual storage is about providing logical storage to hosts and applications independent of physical resources. Virtualization can be implemented in both SAN and NAS storage environments.

In a SAN, virtualization is applied at the block level, whereas in NAS, it is applied at the file level.

iii. **Block-Level Storage Virtualization**

- Block-level storage virtualization provides a translation layer in the SAN, between the hosts and the storage arrays, as shown in Figure 10-6.

- 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 between the virtual LUNs and the physical LUNs on the individual arrays. This facilitates the use of arrays from different vendors simultaneously, without any interoperability issues.

- For a host, all the arrays appear like a single target device and LUNs can be distributed or even split across multiple arrays.

- Block-level storage virtualization extends storage volumes online, resolves application growth requirements, consolidates heterogeneous storage arrays, and enables transparent volume access. It also provides the advantage of non-disruptive data migration.
- In traditional SAN environments, LUN migration from one array to another was an offline event because the hosts needed to be updated to reflect the new array configuration.
- With a block-level virtualization solution in place, the virtualization engine handles the back-end migration of data, which enables LUNs to remain online and accessible while data is being migrated. No physical changes are required because the host still points to the same virtual targets on the virtualization device.

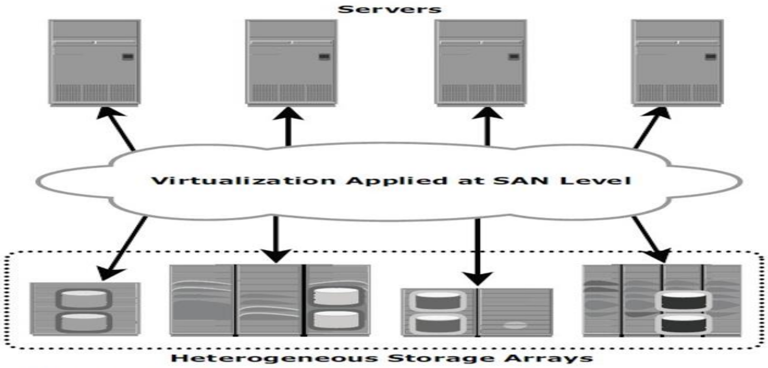
Servers
Virtualization Applied at SAN Level
Heterogeneous Storage Arrays

Figure 10-6: Block-level storage virtualization

iv. 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. Figure 10-7 illustrates a NAS environment before and after the implementation of file-level virtualization.
- Before virtualization, each NAS device or file server is physically and logically independent. Each host knows exactly where its file-level resources are located. Underutilized storage resources and capacity problems result because files are bound to a specific file server.
- It is necessary to move the files from one server to another because of performance reasons or when the file server fills up. Moving files across the environment is not easy and requires downtime for the file servers.
- Moreover, hosts and applications need to be reconfigured with the new path, making it difficult for storage administrators to improve storage efficiency while maintaining the required service level.
- **File-level virtualization simplifies file mobility.** It provides user or application independence from the location where the files are stored.
- File-level virtualization creates a logical pool of storage, enabling users to use a logical path, rather than a physical path, to access files.
- File-level virtualization facilitates the movement of file systems across the online file servers. This means that while the files are being moved, clients can access their files non-disruptively.

- Clients can also read their files from the old location and write them back to the new location without realizing that the physical location has changed.
- Multiple clients connected to multiple servers can perform online movement of their files to optimize utilization of their resources. A global namespace can be used to map the logical path of a file to the physical path names.

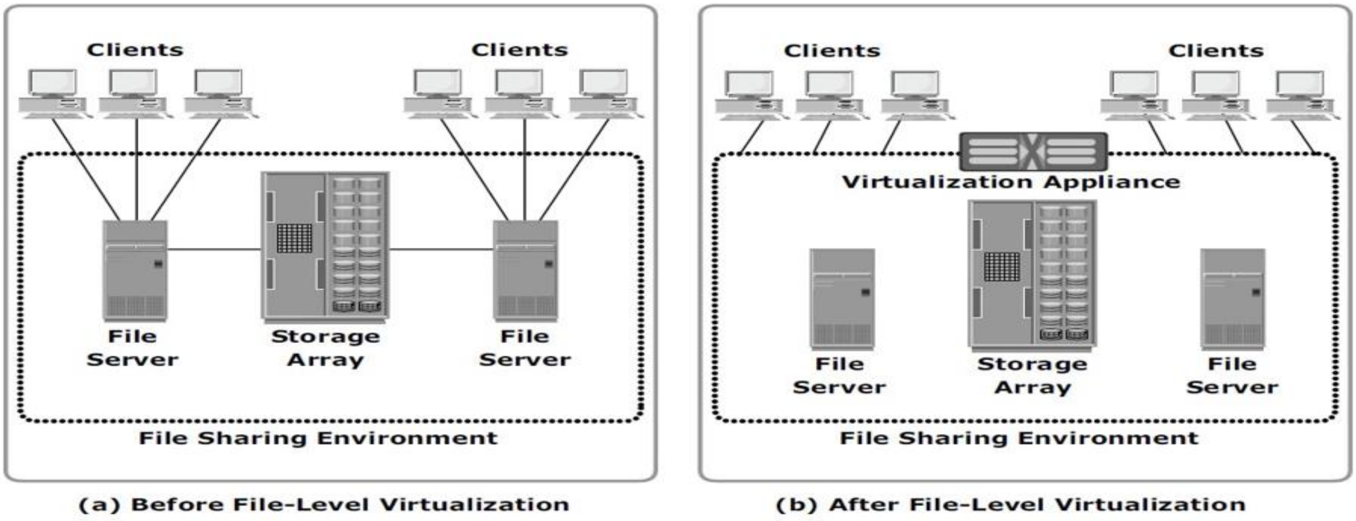

Figure 10-7: NAS device before and after file-level virtualization

9.a) Explain 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

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.
-

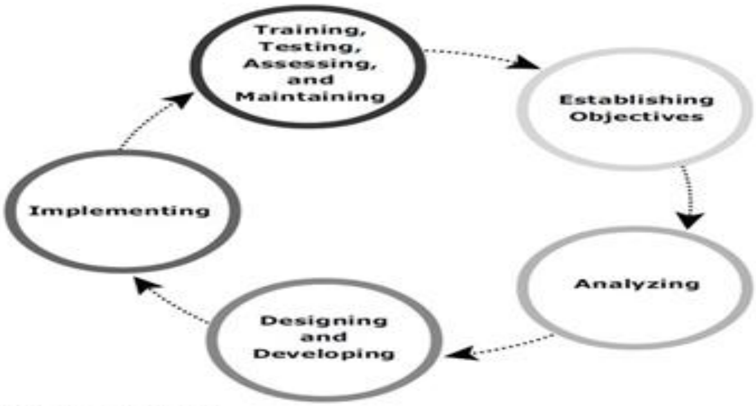

BC planning lifecycle

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 and recovery procedures based on contingency scenarios.
- Test the BC plan regularly to evaluate its performance and identify its limitations.
- Assess the performance reports and identify limitations.
- Update the BC plans and recovery/restart procedures to reflect regular changes within the data center.

9.b) Explain the failure analysis in BC . Mention some important BC technologies Solutions.

Failure analysis involves analyzing the data center to identify systems that are susceptible to a single point of failure and implementing fault-tolerance mechanisms such as redundancy.

1) Single Point of Failure

A single point of failure refers to the failure of a component that can terminate the availability of the entire system or IT service. Figure 11-4 illustrates the possibility of a single point of failure in a system with various components: server, network, switch, and storage array. The figure depicts a system setup in which an application running on the server provides an interface to the client and performs I/O operations. The client is connected to the server through an IP network, the server is connected to the storage array through a FC connection, an HBA installed at the server sends or receives data to and from a storage array, and an FC switch connects the HBA to the storage port.

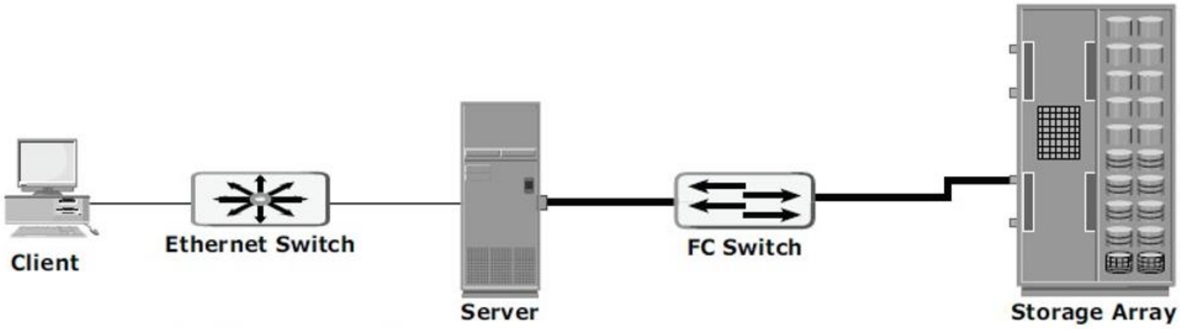

Figure 11-4: Single point of failure

In a setup where each component must function as required to ensure data availability, the failure of a single component causes the failure of the entire data center or an application, resulting in disruption of business operations. In this example, several single points of failure can be identified. The single HBA on the server, the server itself, the IP network, the FC switch, the storage array ports, or even the storage array could become potential single points of failure. To avoid single points of failure, it is essential to implement a fault-tolerant mechanism.

2) Fault Tolerance

To mitigate a single point of failure, systems are designed with redundancy, such that the system will fail only if all the components in the redundancy group fail. This ensures that the failure of a single component does not affect data availability. Figure 11-5 illustrates the fault-tolerant implementation of the system just described (and shown in Figure 11-4).

Data centers follow stringent guidelines to implement fault tolerance. Careful analysis is performed to eliminate every single point of failure. In the example shown in Figure 11-5, all enhancements in the infrastructure to mitigate single points of failures are emphasized:

1. Configuration of multiple HBAs to mitigate single HBA failure.
2. Configuration of multiple fabrics to account for a switch failure.
3. Configuration of multiple storage array ports to enhance the storage array's availability.
4. RAID configuration to ensure continuous operation in the event of disk failure.
5. Implementing a storage array at a remote site to mitigate local site failure.
6. Implementing server (host) clustering, a fault-tolerance mechanism whereby two or more servers in a cluster access the same set of volumes. Clustered servers exchange heartbeats to inform each other about their health. If one of the servers fails, the other server takes up the complete workload.

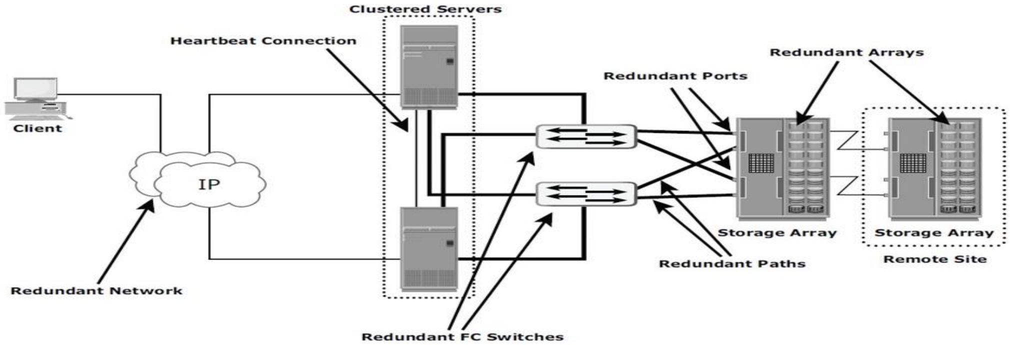

Figure 11-5: Implementation of fault tolerance

3) Multipathing Software

Configuration of multiple paths increases the data availability through path failover. If servers are configured with one I/O path to the data there will be no access to the data if that path fails. Redundant paths eliminate the path to become single points of failure. Multiple paths to data also improve I/O performance through load sharing and maximize server, storage, and data path utilization.

In practice, merely configuring multiple paths does not serve the purpose. Even with multiple paths, if one path fails, I/O will not reroute unless the system recognizes that it has an alternate path. Multipathing software provides the functionality to recognize and utilize alternate I/O path to data. Multipathing software also manages the load balancing by distributing I/Os to all available, active paths.

10.a) Explain the three basic topologies used in backup environment

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.

1. 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.

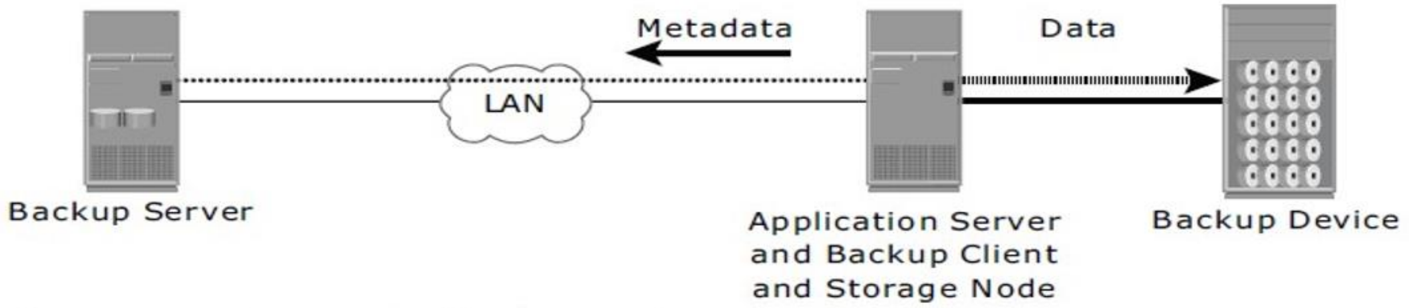

Figure 12-7: Direct-attached backup topology

2. 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).

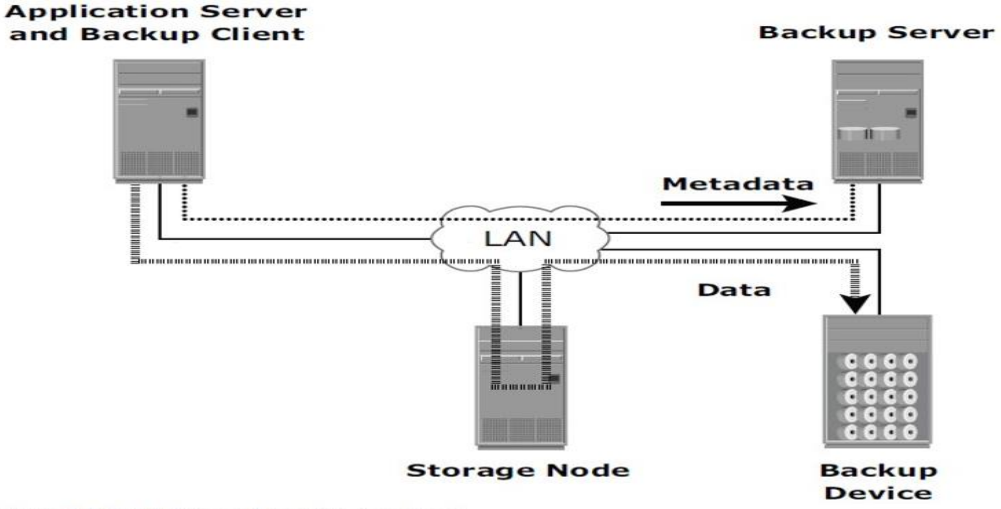

Figure 12-8: LAN-based backup topology

3. 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.

In this example, clients read the data from the mail servers in the SAN and write to the SAN attached backup device. The backup data traffic is restricted to the SAN, and backup metadata is transported over the LAN. However, the volume of metadata is insignificant when compared to production data. LAN performance is not degraded in this configuration.

4. The mixed topology uses both the LAN-based and SAN-based topologies, as shown in Figure 12-10. This topology might be implemented for several reasons, including cost, server location, reduction in administrative overhead, and performance considerations.

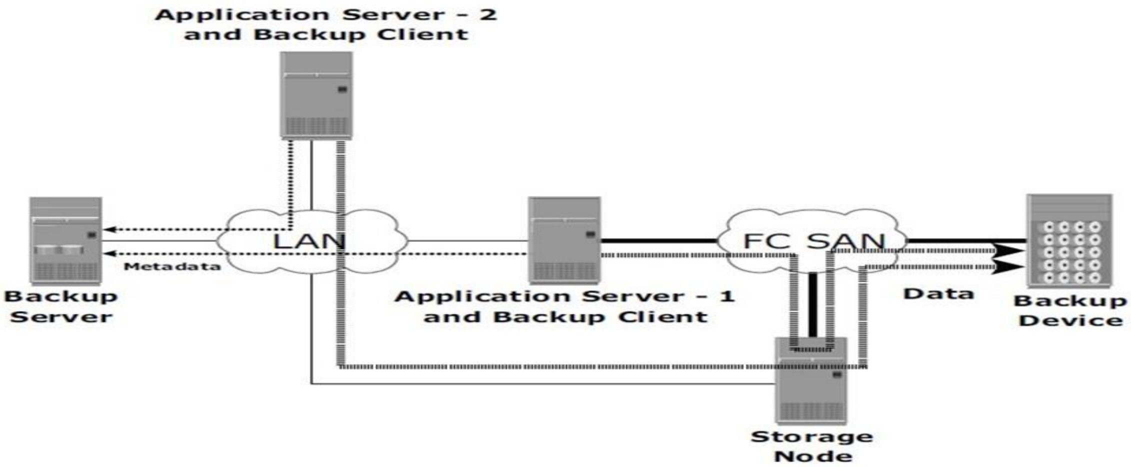

Figure 12-10: Mixed backup topology

10.b) Explain storage array based local Replication

Sol : Replication is the process of creating an exact copy of data. Creating one or more replicas of the production data is one of the ways to provide Business Continuity (BC). These replicas can be used for recovery and restart operations in the event of data loss. Local Replication technologies can be classified based on where the replication is performed.

i) Host based Local Replication

ii) Array based Local Replication

Storage Array–Based 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.

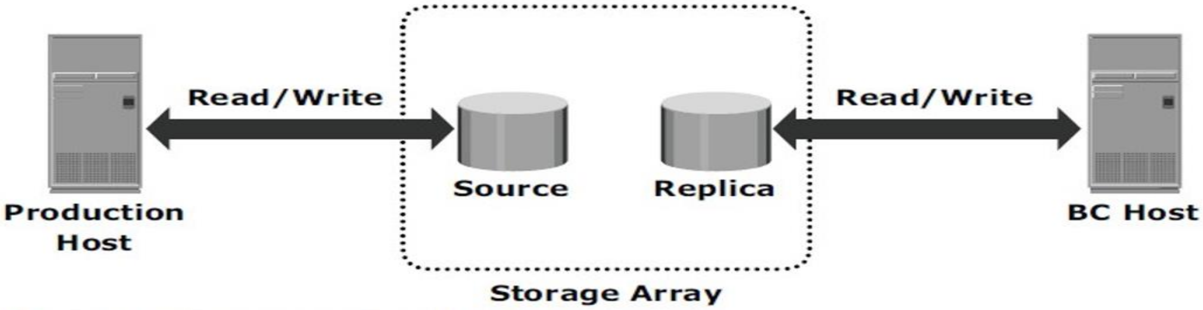

Figure 13-5: Storage array-based replication

Full-Volume Mirroring

In full-volume mirroring, the target is attached to the source and established as a mirror of the source (Figure 13-6 [a]). Existing data on the source is copied to the target. New updates to the source are also updated on the target. After all the data is copied and both the source and the target contain identical data, the target can be considered a mirror of the source. While the target is attached to the source and the synchronization is taking place, the target remains unavailable to any other host. However, the production host can access the source.

After synchronization is complete, the target can be detached from the source and is made available for BC operations. Figure 13-6 (b) shows full-volume mirroring when the target is detached from the source. Notice that both the source and the target can be accessed for read and write operations by the production hosts.

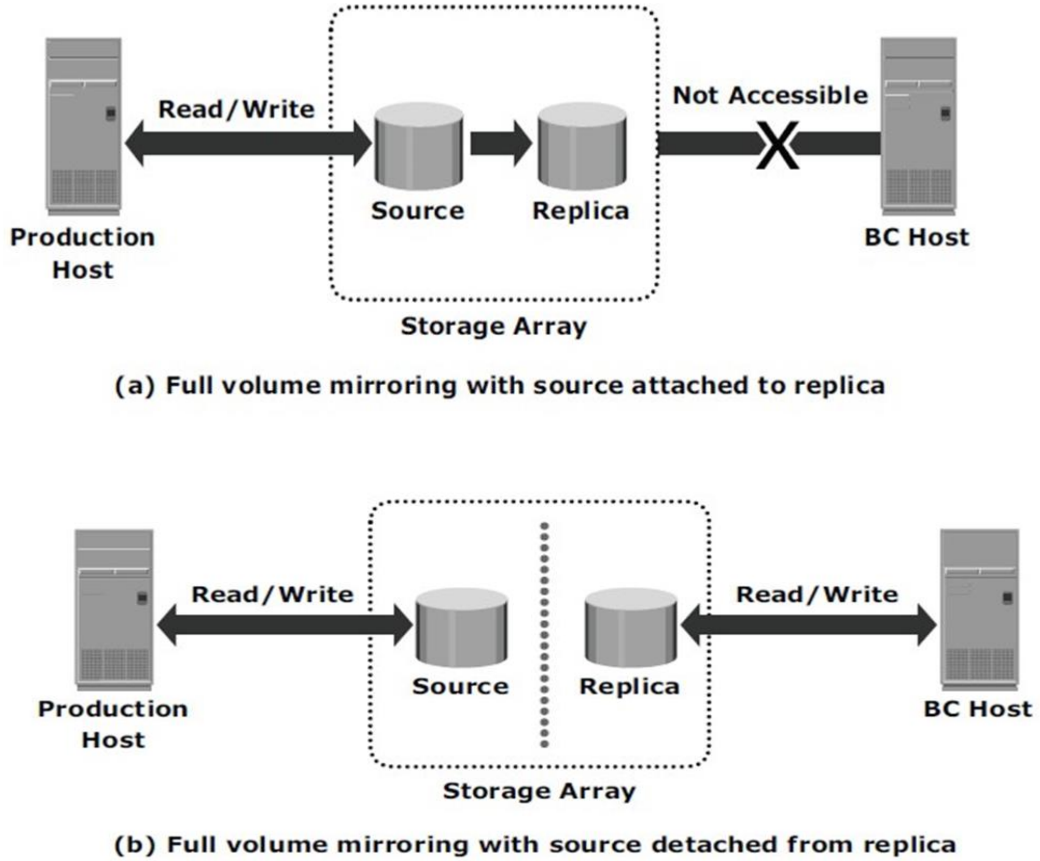

(a) Full volume mirroring with source attached to replica

(b) Full volume mirroring with source detached from replica

Figure 13-6: Full-volume mirroring

After the split from the source, the target becomes a PIT copy of the source. The point-in-time of a replica is determined by the time when the source is detached from the target. For example, if the time of detachment is 4:00 pm, the PIT for the target is 4:00 pm

After detachment, changes made to both source and replica can be tracked at some predefined granularity. This enables incremental resynchronization (source to target) or incremental restore (target to source). The granularity of the data change can range from 512 byte blocks to 64 KB blocks. Changes are typically tracked using bitmaps, with one bit assigned for each block. If any updates occur to a particular block, the whole block is marked as changed, regardless of the size of the actual update. However, for resynchronization (or restore), only the changed blocks have to be copied, eliminating the need for a full synchronization (or restore) operation. This method reduces the time required for these operations considerably.

In full-volume mirroring, the target is inaccessible for the duration of the synchronization process, until detachment from the source. For large databases, this can take a long time.

Pointer-Based, Full-Volume Replication

An alternative to full-volume mirroring is pointer-based full-volume replication. Like full-volume mirroring, this technology can provide full copies of the source data on the targets. Unlike full-volume mirroring, the target is made immediately available at the activation of the replication session. Hence, one need not wait for data synchronization to, and detachment of, the target in order to access it. The time of activation defines the PIT copy of source.

Pointer-based, full-volume replication can be activated in either Copy on First Access (CoFA) mode or Full Copy mode. In either case, at the time of activation, a protection bitmap is created for all data on the source devices. Pointers are initialized to map the (currently) empty data blocks on the target to the corresponding original data blocks on the source. The granularity can range from 512 byte blocks to 64 KB blocks or higher. Data is then copied from the source to the target, based on the mode of activation.

In CoFA, after the replication session is initiated, data is copied from the source to the target when the following occurs:

1. A write operation is issued to a specific address on the source for the first time (see Figure 13-7).
2. A read or write operation is issued to a specific address on the target for the first time (see Figure 13-8 and Figure 13-9).

When a write is issued to the source for the first time after session activation, original data at that address is copied to the target. After this operation, the new data is updated on the source. This ensures that original data at the point-in-time of activation is preserved on the target. This is illustrated in Figure 13-7.

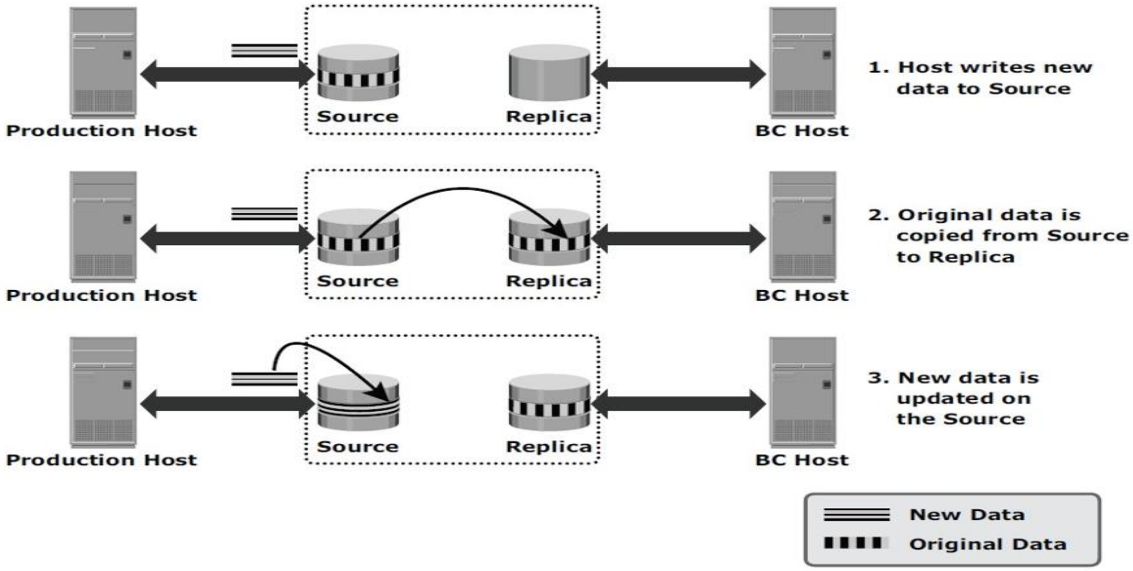

igure 13-7: Copy on first access (CoFA) — write to source

When a read is issued to the target for the first time after session activation, the original data is copied from the source to the target and is made available to the host. This is illustrated in Figure 13-8.

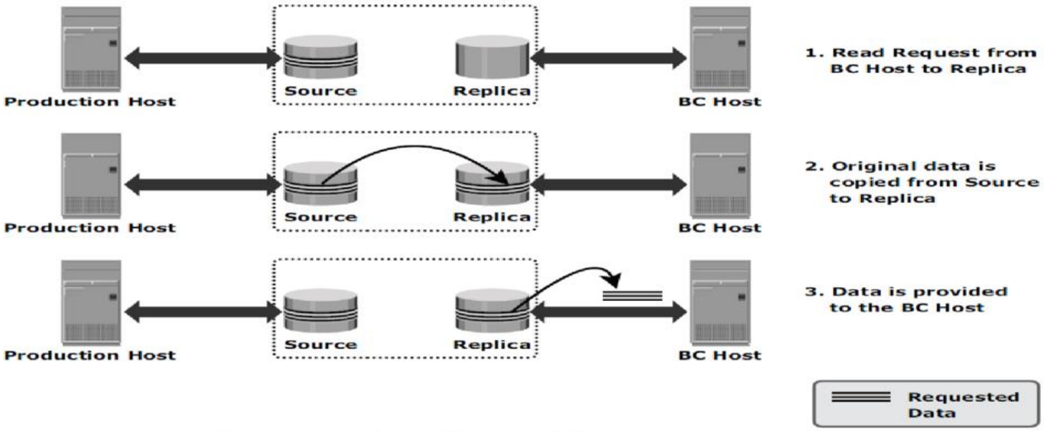
Production Host
Source
Replica
BC Host
1. Read Request from BC Host to Replica
Production Host
Source
Replica
BC Host
2. Original data is copied from Source to Replica
Production Host
Source
Replica
BC Host
3. Data is provided to the BC Host
Requested Data

Figure 13-8: Copy on first access (CoFA) — read from target

When a write is issued to the target for the first time after session activation, the original data is copied from the source to the target. After this, the new data is updated on the target. This is illustrated in Figure 13-9.

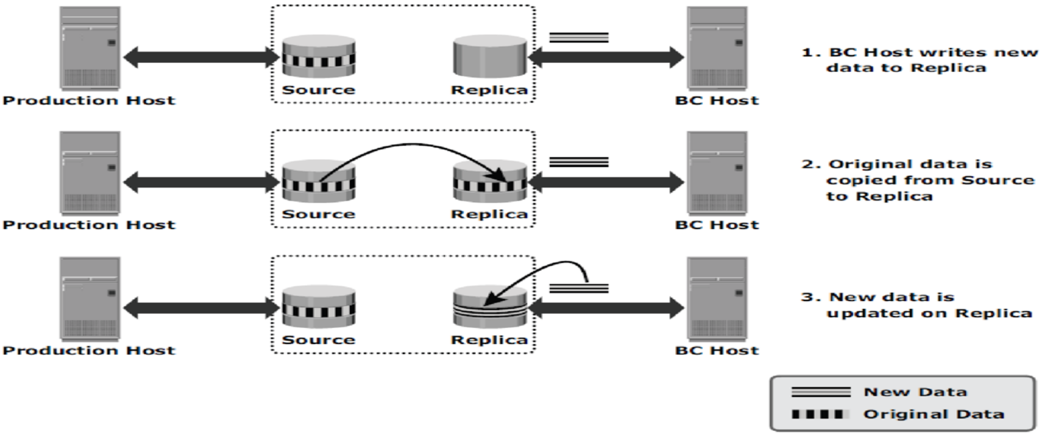

Figure 13-9: Copy on first access (CoFA) — write to target

In all cases, the protection bit for that block is reset to indicate that the original data has been copied over to the target. The pointer to the source data can now be discarded. Subsequent writes to the same data block on the source, and reads or writes to the same data blocks on the target, do not trigger a copy operation (and hence are termed Copy on First Access).

Pointer-Based Virtual Replication

In pointer-based virtual replication, at the time of session activation, the target contains pointers to the location of data on the source. The target does not contain data, at any time. Hence, the target is known as a virtual replica. Similar to pointer-based full-volume replication, a protection bitmap is created for all data on the source device, and the target is immediately accessible. Granularity can range from 512 byte blocks to 64 KB blocks or greater. When a write is issued to the source for the first time after session activation, original data at that address is copied to a predefined area in the array. This area is generally termed the save location. The pointer in the target is updated to point to this data address in the save location. After this, the new write is updated on the source. This process is illustrated in Figure 13-10.

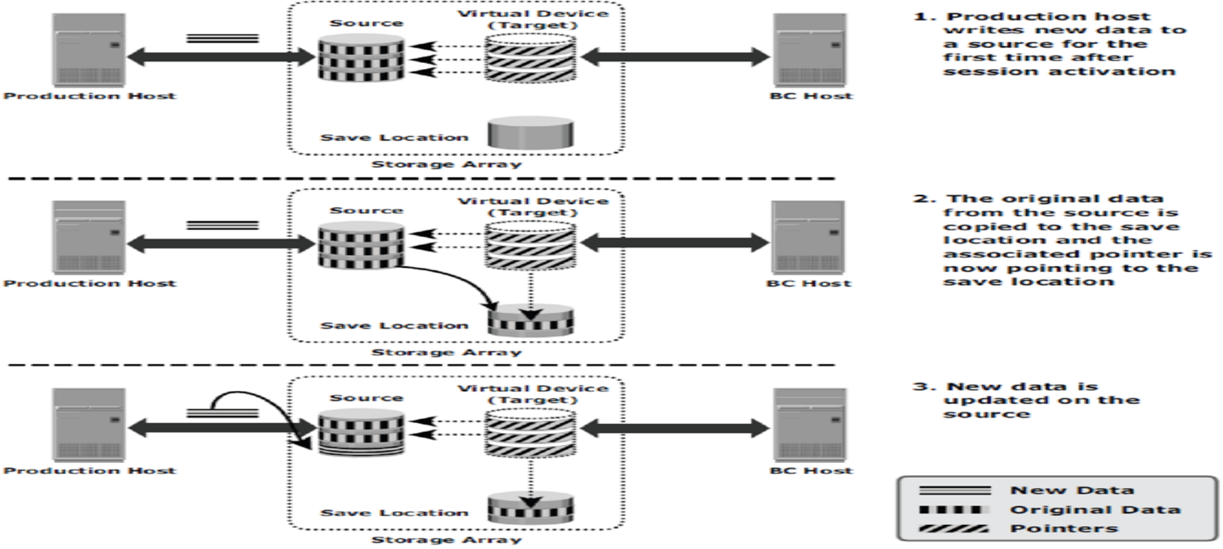

Figure 13-10: Pointer-based virtual replication — write to source

When a write is issued to the target for the first time after session activation, original data is copied from the source to the save location and similarly the pointer is updated to data in save location. Another copy of the original data is created in the save location before the new write is updated on the save location. This process is illustrated in Figure 13-11.

When reads are issued to the target, unchanged data blocks since session activation are read from the source. Original data blocks that have changed are read from the save location.

Pointer-based virtual replication uses CoFW technology. Subsequent writes to the same data block on the source or the target do not trigger a copy operation.

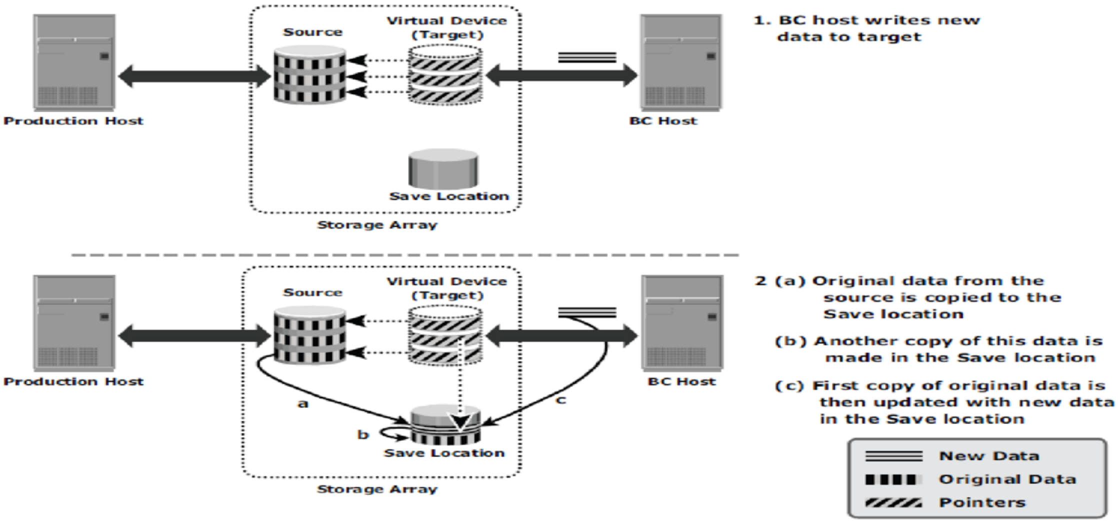
Source
Virtual Device (Target)
Production Host
BC Host
Save Location
Storage Array
1. BC host writes new data to target
Source
Virtual Device (Target)
Production Host
BC Host
a
b
c
Save Location
Storage Array
2 (a) Original data from the source is copied to the Save location
(b) Another copy of this data is made in the Save location
(c) First copy of original data is then updated with new data in the Save location
New Data
Original Data
Pointers

Figure 13-11: Pointer-based virtual replication — write to target

Data on the target is a combined view of unchanged data on the source and data on the save location. Unavailability of the source device invalidates the data on the target. As the target only contains pointers to data, the physical capacity required for the target is a fraction of the source device. The capacity required for the save location depends on the amount of expected data change.