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

Internal Assessment Test II – Nov. 2017

Sub:	STORAGE AREA NETWORKS (SAN)					Sub Code:	10CS765	Branch:	ISE/CSE	
Date:	09.11.2017	Duration:	90 min's	Max Marks:	50	Sem / Sec:	VII		OBE	

Answer any FIVE FULL Questions

		MARKS	CO	RBT
1 (a)	Explain connectivity options of Fiber Channel architecture with relevant diagrams.	[10]	CO3	L3
2 (a)	Explain FC with respect to protocol stack, and zoning.	[10]	CO3	L3
3 (a)	Explain the architecture of CAS with neat diagram. List out the features of CAS.	[10]	CO4	L3
4 (a)	Explain object storage and retrieval in CAS with suitable diagrams	[10]	CO4	L3
5 (a)	What are the backup topologies? Explain with suitable diagrams.	[10]	CO5	L3
6 (a)	Draw and explain BC planning life cycle.	[10]	CO5	L3
7 (a)	Explain DAS, its types, advantages and disadvantages.	[10]	CO3	L3

1a.Explain connectivity options of Fiber channel architecture with relevant diagrams. (10 Marks)

Solution:

The FC architecture supports three basic interconnectivity options: point-to-point, arbitrated loop (FC-AL), and fabric connect.

Point to Point:

Point-to-point is the simplest FC configuration — two devices are connected directly to each other, as shown in Figure. This configuration provides a dedicated connection for data transmission between nodes. However, the point-to-point configuration offers limited connectivity, as only two devices can communicate with each other at a given time.

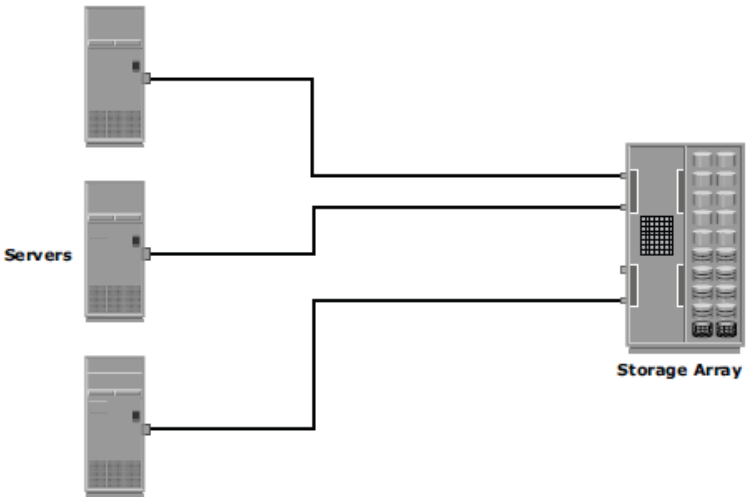

Fig: point-to-point connectivity

FCAL

In the FC-AL configuration, devices are attached to a shared loop, as shown in Figure. FC-AL has the characteristics of a token ring topology and a physical star topology. In FC-AL, each device contends with other devices to perform I/O operations. Devices on the loop must "arbitrate" to gain control of the loop. At any given time, only one device can perform I/O operations on the loop.

The FC-AL configuration has the following limitations in terms of scalability:
 FC-AL shares the bandwidth in the loop.
 Only one device can perform I/O operations at a time. Because each device in a loop has to wait for its turn to process an I/O request, the speed of data transmission is low in an FC-AL topology.
 FC-AL uses 8-bit addressing. It can support up to 127 devices on a loop.
 Adding or removing a device results in loop re-initialization, which can cause a momentary pause in loop traffic.

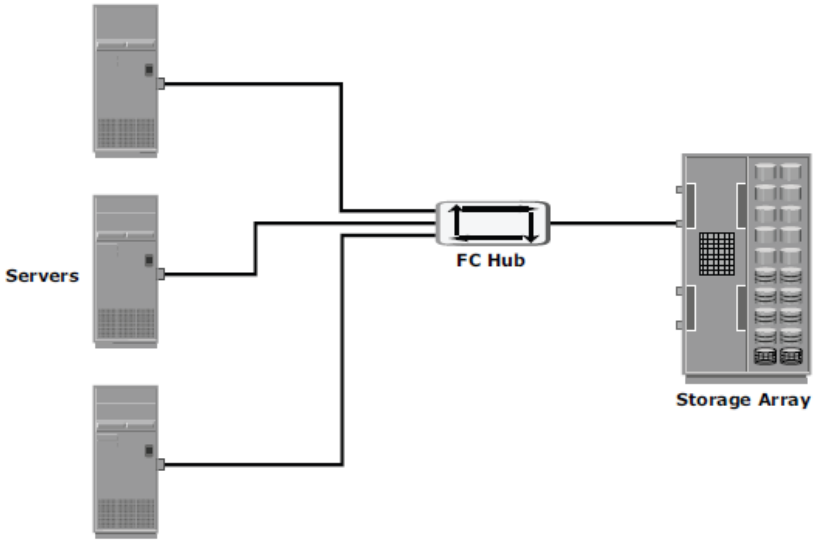

Fig: FCAL

FC SW

Fibre Channel switched fabric (FC-SW) network provides interconnected devices, dedicated bandwidth, and scalability. The addition or removal of a device in a switched fabric is minimally disruptive; it does not affect the ongoing traffic between other devices. FC-SW is also referred to as fabric connect. A fabric is a logical space in which all nodes communicate with one another in a network.

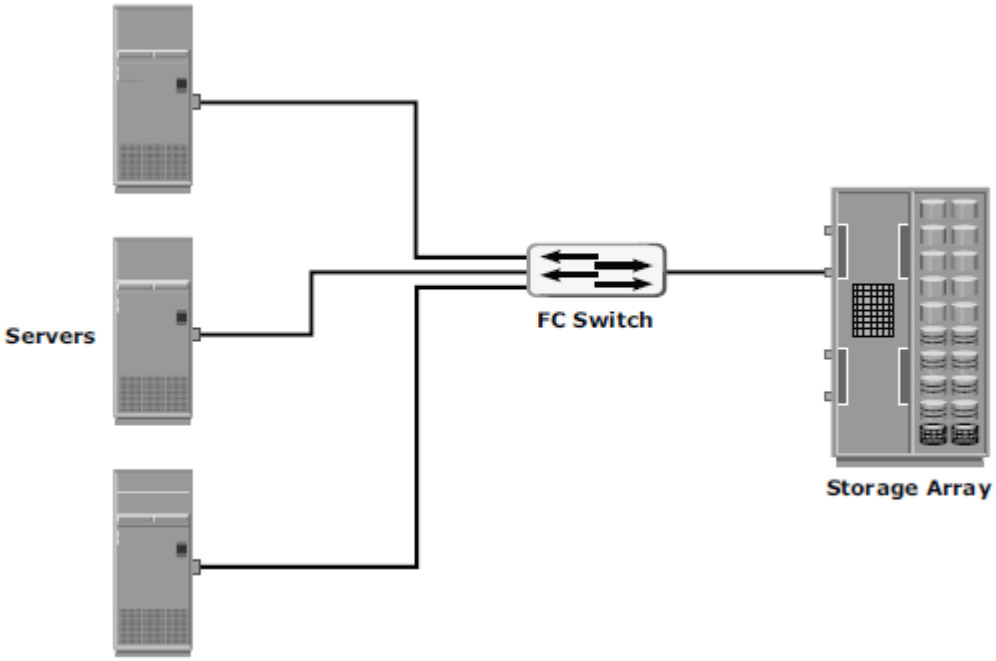

Fig: FC SW

2.a Explain FC with respect to protocol stack, and zoning. (10Marks)
Solution:

Fibre Channel Protocol Stack : It is easier to understand a communication protocol by viewing it as a structure of independent layers. FCP defines the communication protocol in five layers: FC-0 through FC-4 (except FC-3 layer, which is not implemented). In a layered communication model, the peer layers on each node talk to each other through defined protocols. Figure 6-13 illustrates the fibre channel protocol stack.

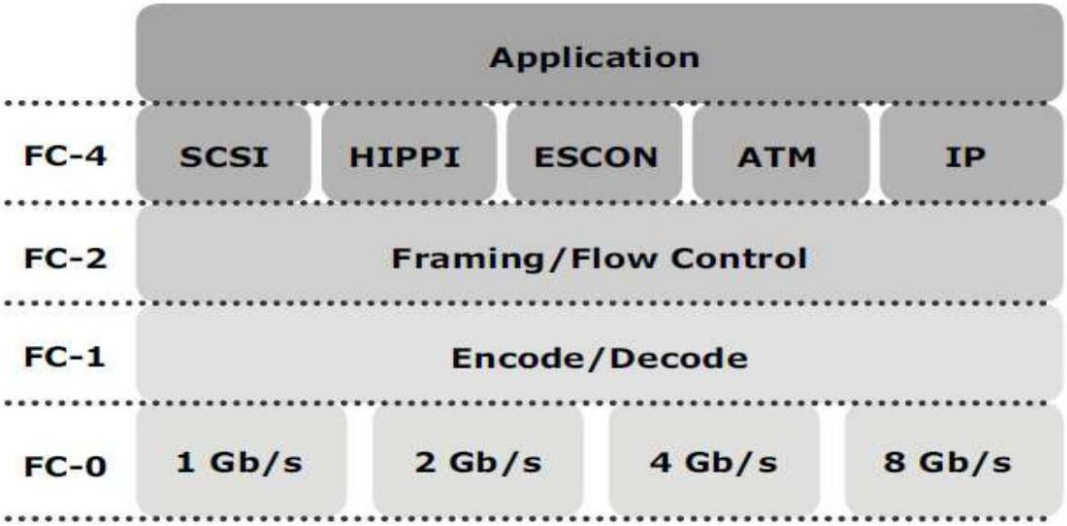

Figure 6-13: Fibre channel protocol stack

FC-4 Upper Layer Protocol: FC-4 is the uppermost layer in the FCP stack. This layer defines the application interfaces and the way Upper Layer Protocols (ULPs) are mapped to the lower FC layers. The FC standard defines several protocols that can operate on the FC-4 layer (see Figure 6-7). Some of the protocols include SCSI, HIPPI Framing Protocol, Enterprise Storage Connectivity (ESCON), ATM, and IP.

FC-2 Transport Layer*:* The FC-2 is the transport layer that contains the payload, addresses of the source and destination ports, and link control information. The FC-2 layer provides Fibre Channel addressing, structure, and organization of data (frames, sequences, and exchanges). It also defines fabric services, classes of service, flow control, and routing.

FC-1 Transmission Protocol: This layer defines the transmission protocol that includes serial encoding and decoding rules, special characters used, and error control. At the transmitter node, an 8-bit character is encoded into a 10-bit transmissions character. This character is then transmitted to the receiver node. At the receiver node, the 10-bit character is passed to the FC-1 layer, which decodes the 10-bit character into the original 8-bit character.

FC-0 Physical Interface*:* FC-0 is the lowest layer in the FCP stack. This layer defines the physical interface, media, and transmission of raw bits. The FC-0 specification includes cables, connectors, and optical and electrical parameters for a variety of data rates. The FC transmission can use both electrical and optical media.

Zoning: Zoning is an FC switch function that enables nodes within the fabric to be logically segmented into groups that can communicate with each other (see Figure 6-18). When a device (host or storage array) logs onto a fabric, it is registered with the name server. When a port logs onto the fabric, it goes through a device discovery process with other devices registered in the name server. The zoning function controls this process by allowing only the members in the same zone to establish these link-level services. Multiple zone sets may be defined in a fabric, but only one zone set can be active at a time. A zone set is a set of zones and a zone is a set of members. A member may be in multiple zones. Members, zones, and zone sets form the hierarchy defined in the zoning process (see Figure 6-19). *Members* are nodes within the SAN that can be included in a zone. *Zones* comprise a set of members that have access to one another. A port or a node can be a member of multiple zones. *Zone sets* comprise a group of zones that can be activated or deactivated as a single entity in a fabric. Only one

zone set per fabric can be active at a time. Zone sets are also referred to as *zone configurations*.

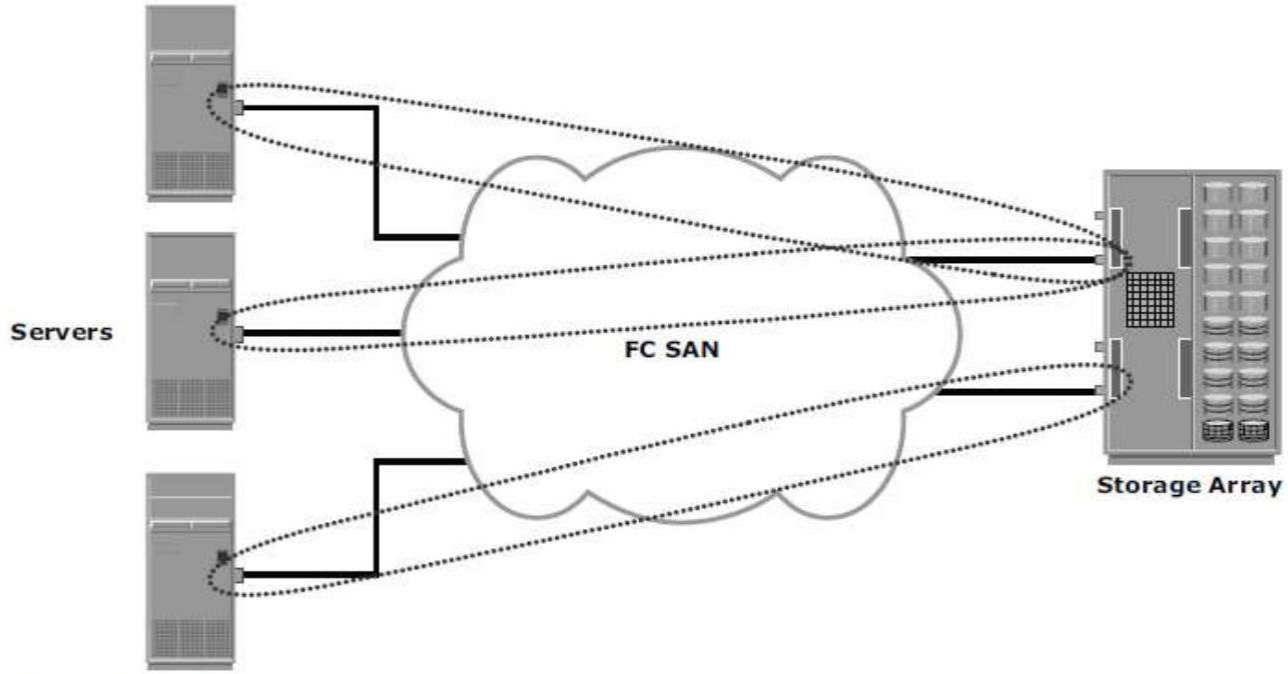

Figure 6-18: Zoning

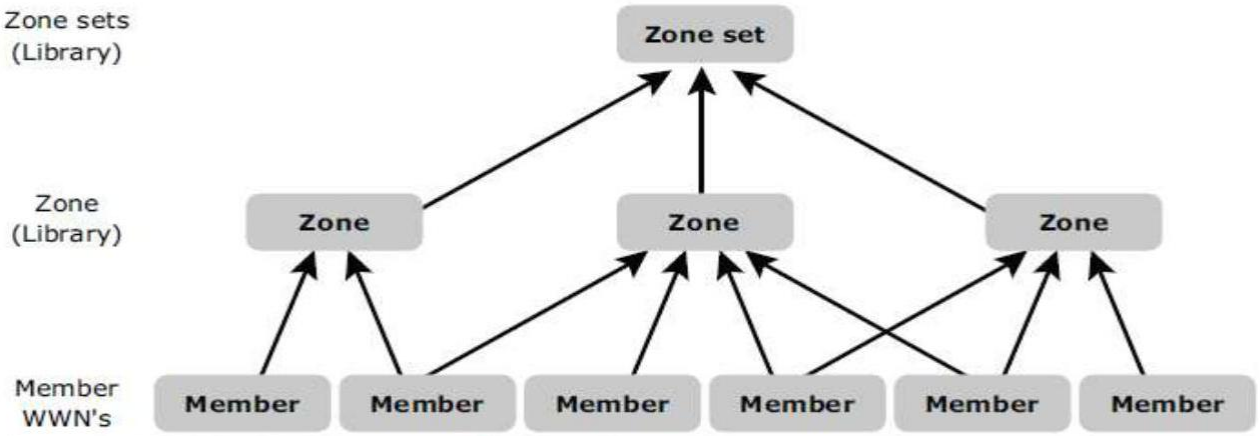

Figure 6-19: Members, zones, and zone sets

Types of Zoning : Zoning can be categorized into three types:

Port zoning: It uses the FC addresses of the physical ports to define zones. In port zoning, access to data is determined by the physical switch port to which a node is connected. The FC address is dynamically assigned when the port logs on to the fabric. Therefore, any change in the fabric configuration affects zoning. Port zoning is also called *hard zoning*. Although this method is secure, it requires updating of zoning configuration information in the event of fabric reconfiguration.

WWN zoning: It uses World Wide Names to define zones. WWN zoning is also referred to as *soft zoning*. A major advantage of WWN zoning is its flexibility. It allows the SAN to be recabled without reconfiguring the zone information. This is possible because the WWN is static to the node port.

Mixed zoning: It combines the qualities of both WWN zoning and port zoning. Using mixed zoning enables a specific port to be tied to the WWN of a node.

3.a Explain the architecture of CAS with neat diagram. List out the features of CAS. (10 Marks)

Solution:

CAS is an object-based system that has been purposely built for storing fixed content data. A client accesses the CAS-Based storage over a LAN through the server that runs the CAS API (application programming interface). The CAS API is responsible for performing functions that enable an application to store and retrieve the data. CAS architecture is a Redundant Array of Independent Nodes (RAIN). It contains storage nodes and access nodes networked as a cluster by using a private LAN that is internal to it. The internal LAN can be reconfigured automatically to detect the configuration changes such as the addition of storage or access nodes. Clients access the CAS on a separate LAN, which is used for interconnecting clients and servers to the CAS. The nodes are configured with low-cost, high-capacity ATA HDDs. These nodes run an operating system with special software that implements the features and functionality required in a CAS system. When the cluster is installed, the nodes are configured with a "role" defining the functionality they provide to the cluster. A node can be configured as a storage node, an access node, or a dual-role node. Storage nodes store and protect data objects. They are sometimes referred to as back-end nodes. Access nodes provide connectivity to application servers through the customer's LAN. They establish connectivity through a private LAN to the storage nodes in the cluster. The number of access nodes is determined by the amount of user required throughput from the cluster. If a node is configured solely as an "access node," its disk space cannot be used to store data objects. This configuration is generally found in older installations of CAS. Storage and retrieval requests are sent to the access node via the customer's LAN.

Dual-role nodes provide both storage and access node capabilities.

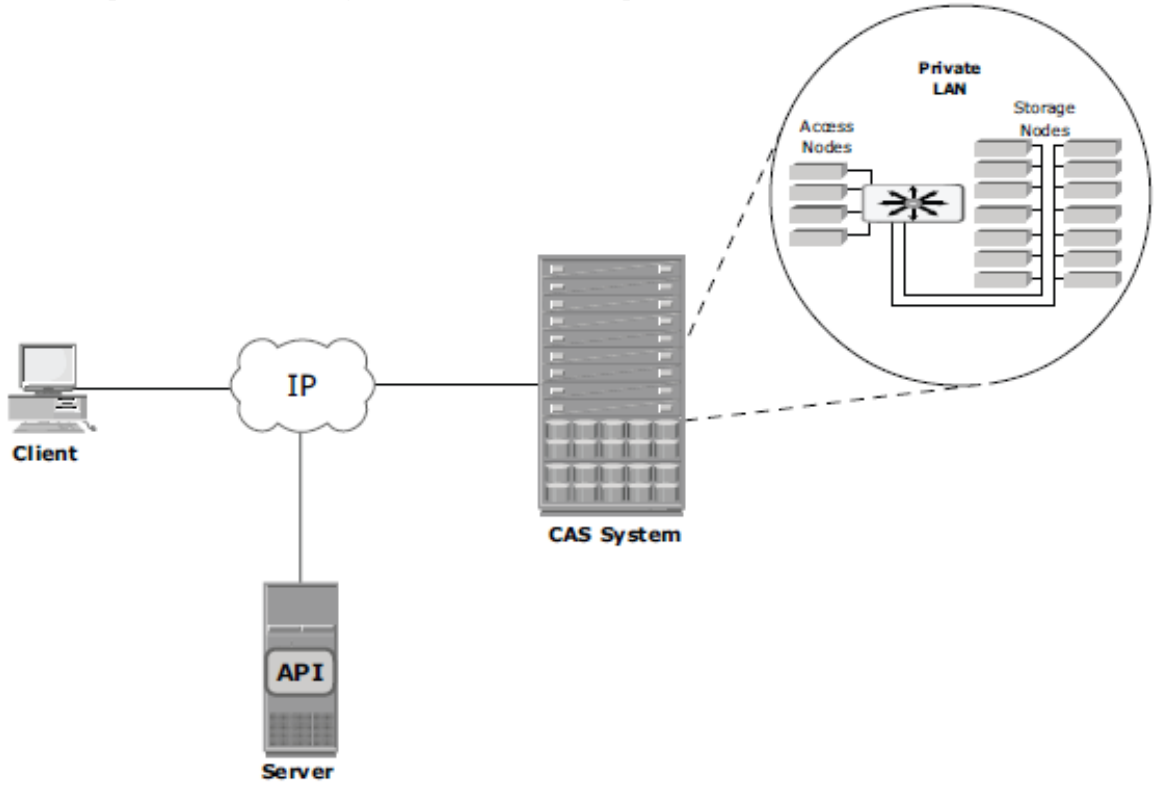

Fig: CAS Architecture

The following features are an essential part of any CAS solution:

- **Integrity checking:** It ensures that the content of the file matches the digital signature (hashed output or CA). The integrity checks can be done on every read or by using a background process. If problems are identified in any of the objects the nodes automatically repair or regenerate the object.
- **Data protection and node resilience:** This ensures that the content stored on the CAS system is available in the event of disk or node failure. Some CAS systems provide local replication or mirrors that copy a data object to another node in the same cluster.

- **Load balancing:** Distributes data objects on multiple nodes to provide maximum throughput, availability, and capacity utilization.
- S**calability:** Adding more nodes to the cluster without any interruption to data access and with minimum administrative overhead.
- **Self diagnosis and repair:** Automatically detects and repairs corrupted objects and alert the administrator of any potential problem. These failures can be at an object level or a node level. They are transparent to the users who access the archive. CAS systems can be configured to alert remote support teams who diagnose and make repairs remotely.
- **Report generation and event notification:** Provides on-demand reporting and event notification. A command-line interface (CLI) or a graphical user interface (GUI) can produce various types of reports. Any event notification can be communicated to the administrator through syslog, SNMP, SMTP, or e-mail.
- **Fault tolerance:** Ensures data availability if a component of the CAS system fails, through the use of redundant components and data protection schemes. If remote replication of CAS is implemented, failover to the remote CAS system occurs when the primary CAS system is unavailable.
- **Audit trails:** Enable documentation of management activity and any access and disposition of data. Audit trails are mandated by compliance requirements.

4a. Explain object storage and retrieval in CAS with suitable diagrams (10 Marks)

Solution:

The process of storing and retrieving objects in CAS is explained in Figures 9-3 and 9-4, respectively. This process requires an understanding of the following **CAS terminologies:**

1. **Application programming interface (API):** A high-level implementation of an interface that specifies the details of how clients can make service requests. The CAS API resides on the application server and is responsible for storing and retrieving the objects in a CAS system.
2. **Access profile:** Used by access applications to authenticate to a CAS cluster and by CAS clusters to authenticate themselves to each other for replication.
3. **Virtual pools:** Enable a single logical cluster to be broken up into multiple logical groupings of data.
4. **Binary large object (BLOB):** The actual data without the descriptive information (metadata). The distinct bit sequence of user data represents the actual content of a file and is independent of the name and physical location.
5. **Content address (CA):** An object's address, which is created by a hash algorithm run across the binary representation of the object. While generating a CA, the hash algorithm considers all aspects of the content, returning a unique content address to the user's application. A unique number is calculated from the sequence of bits that constitutes file content. If even a single character changes in the file, the resulting CA is different. A *hash output,* also called a *digest,* is a type of fingerprint for a variable-length data file. This output represents the file contents and is used to locate the file in a CAS system. The digest can be used to verify whether the data is authentic or has changed because of equipment failure or human intervention. When a user tries to retrieve or open a file, the server sends the CA to the CAS system with the appropriate function to read the file. The CAS system uses the CA to locate the file and passes it back to the application server.
6. **C-Clip:** A virtual package that contains data (BLOB) and its associated CDF. The *C-Clip ID* is the CA that the system returns to the client application. It is also referred as a *C-Clip handle* or *C-Clip reference*.
7. **C-Clip Descriptor File (CDF):** An XML file that the system creates while making a C-Clip. This file includes CAs for all referenced BLOBs and associated metadata. Metadata includes characteristics of CAS objects such as size, format, and expiration date.

Referring to Figure 9-3, the **data object storage process** in a CAS system is as follows:

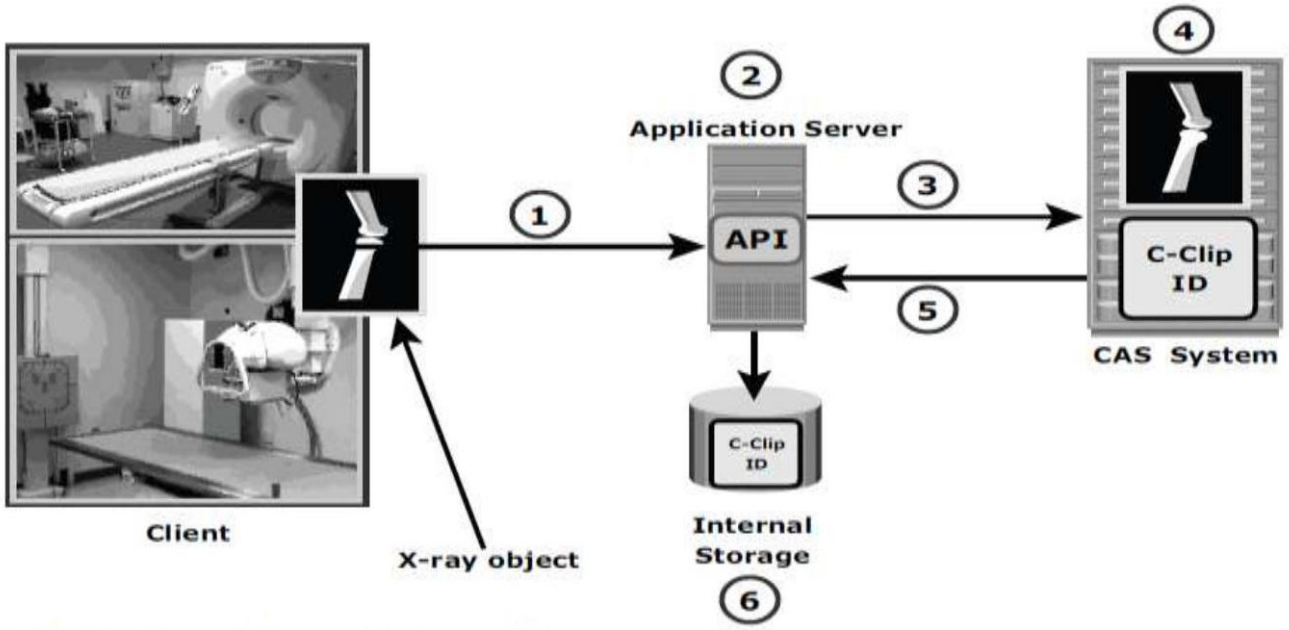

Figure 9-3: Storing data objects on CAS

End users present the data to be archived to the CAS API via an application. The application server may also interact directly with the source (e.g., an X-ray machine) that generated this fixed content.

1. The API separates the actual data (BLOB) from the metadata and the CA is calculated from the object's binary representation.
2. The content address and metadata of the object are then inserted into the C-Clip Descriptor File (CDF). The C-clip is then transferred to and stored on the CAS system.
3. The CAS system recalculates the object's CA as a validation step and stores the object. This is to ensure that the content of the object has not changed.
4. An acknowledgment is sent to the API after a mirrored copy of the CDF and a protected copy of the BLOB have been safely stored in the CAS system. After a data object is stored in the CAS system, the API is given a C-Clip ID and C-Clip ID is stored local to the application server.
5. Using the C-Clip ID, the application can read the data back from the CAS system.

Once an object is stored successfully, it is made available to end users for retrieval and use.

The process of **data retrieval** from CAS follows these steps:

1. The end user or an application requests an object.
2. The application queries the local table of C-Clip IDs stored in the local storage and locates the C-Clip ID for the requested object.
3. Using the API, a retrieval request is sent along with the C-Clip ID to the CAS system.
4. CAS system delivers the requested information to the application, which in turn delivers it to the end user.

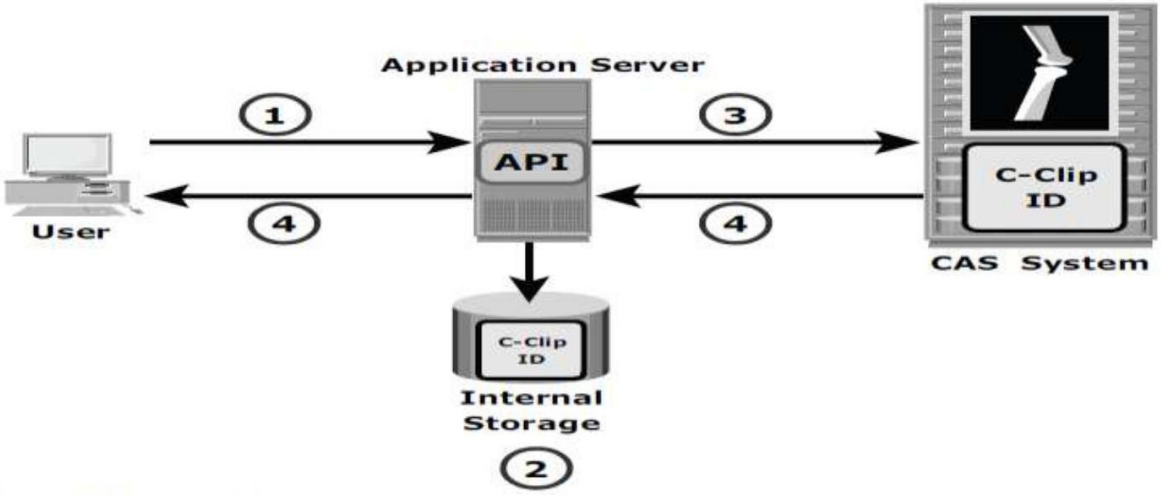

Figure 9-4: Data object retrieval from CAS system

5a. What are the backup topologies? Explain with suitable diagrams. (10 Marks)

Solution:

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.

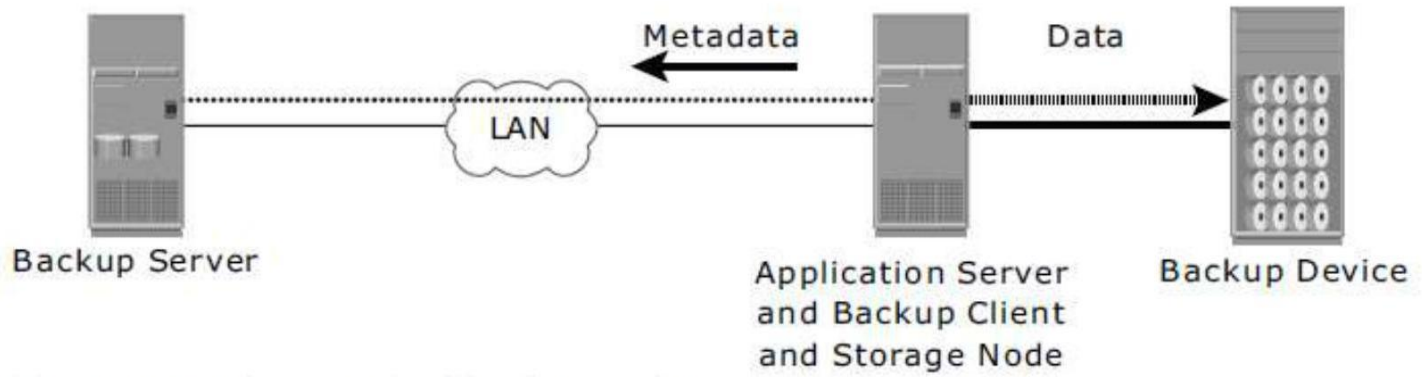

Figure 12-7: Direct-attached backup topology

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).

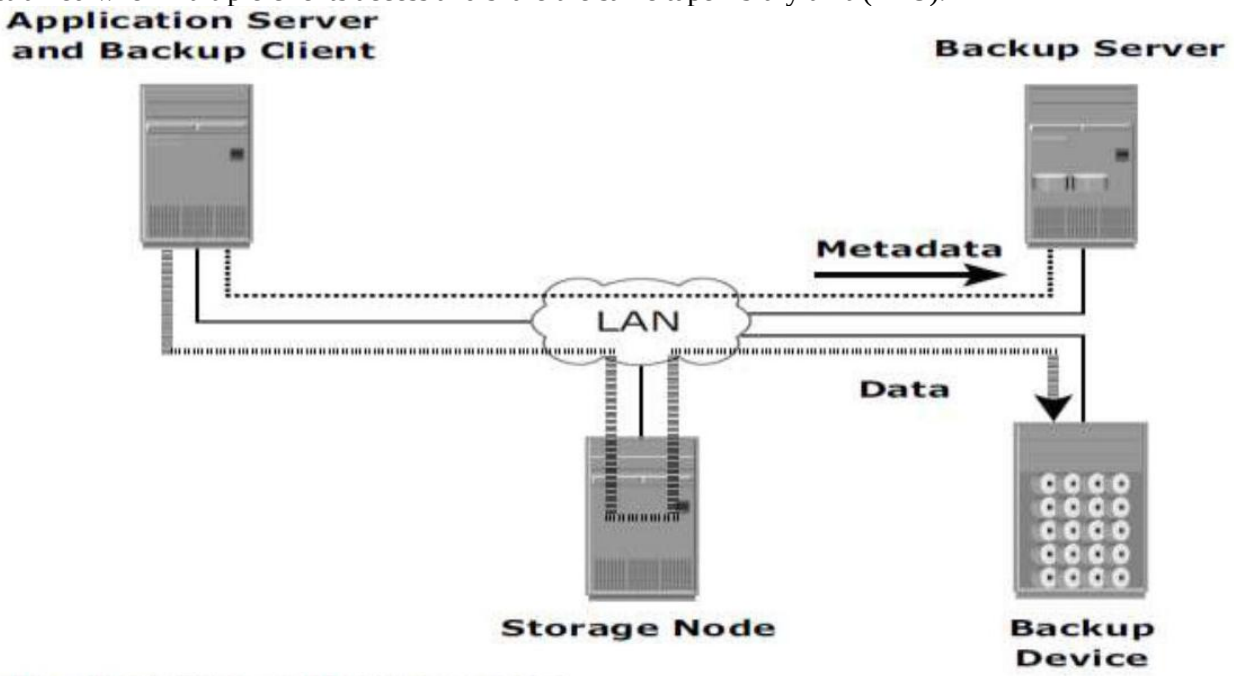

Figure 12-8: LAN-based backup topology

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.

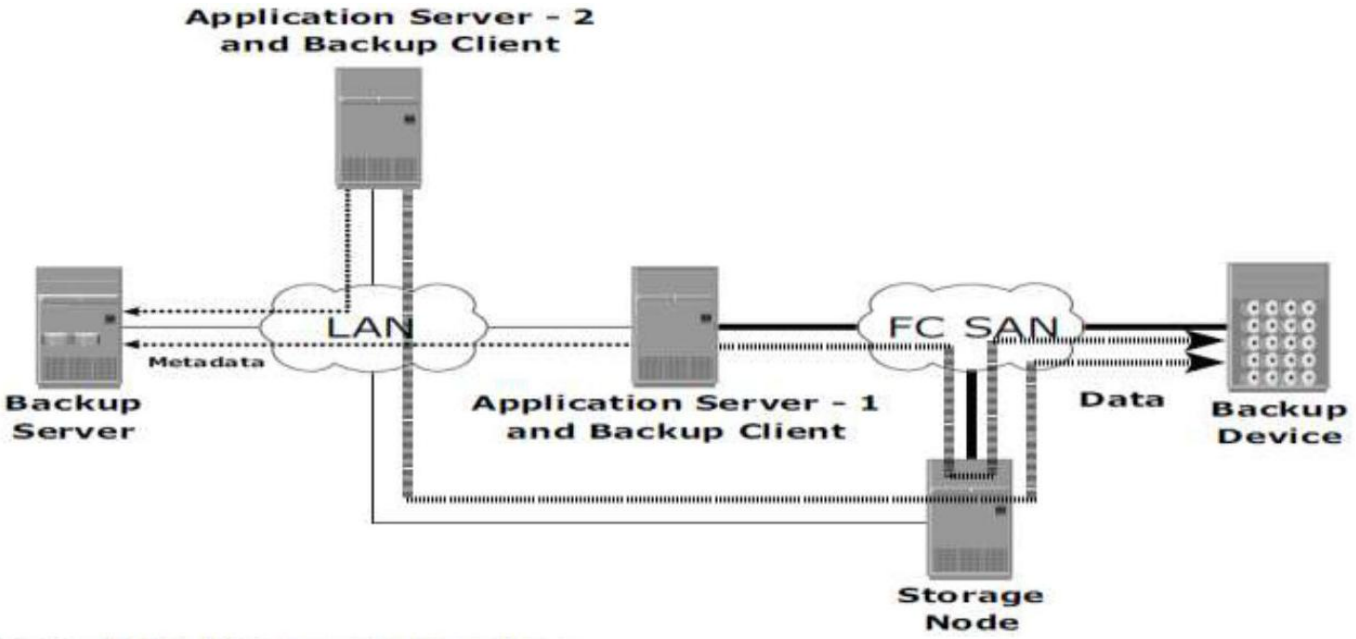

Figure 12-10: Mixed backup topology

6a. Draw and explain BC planning life cycle. (10 Marks)

Solution:

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

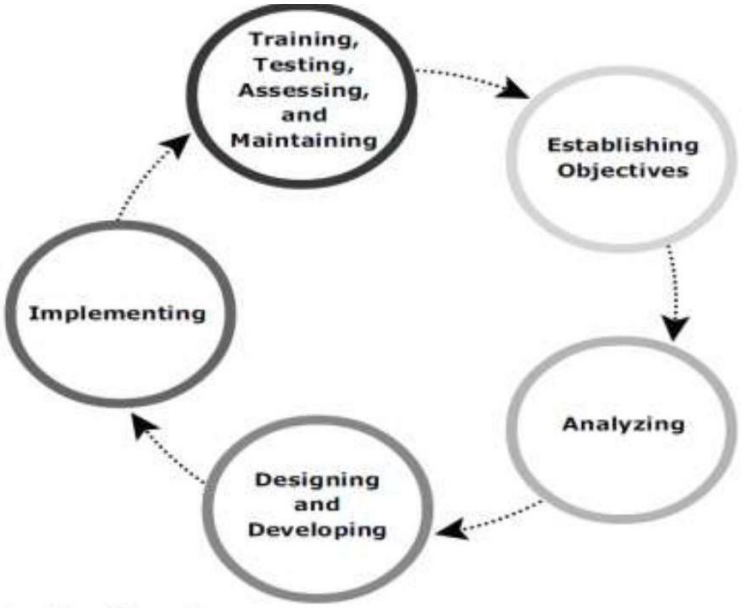

11-3: BC planning lifecycle

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

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.

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.

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.

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.

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.
- Assess the performance reports and identify limitations.
- Update the BC plans and recovery/restart procedures to reflect regular changes within the data center.

7a. Explain DAS, its types, advantages and disadvantages. (10 Marks)

Solution:

D*irect-Attached Storage (DAS)* is an architecture where storage connects directly to servers. Applications access data from DAS using block-level access protocols. The internal HDD of a host, tape libraries, and directly connected external HDD packs are some examples of DAS.

Types of DAS

DAS is classified as internal or external, based on the location of the storage device with respect to the host.

Internal DAS: In *internal DAS* architectures, the storage device is internally connected to the host by a serial or parallel bus. The physical bus has distance limitations and can only be sustained over a shorter distance for

high-speed connectivity. Supports limited number of devices, and they occupy a large amount of space inside the host, making maintenance of other components difficult.

External DAS: In *external DAS* architectures, the server connects directly to the external storage device (see Figure 5-1). In most cases, communication between the host and the storage device takes place over SCSI or FC protocol. Compared to internal DAS, an external DAS overcomes the distance and device count limitations and provides centralized management of storage devices.

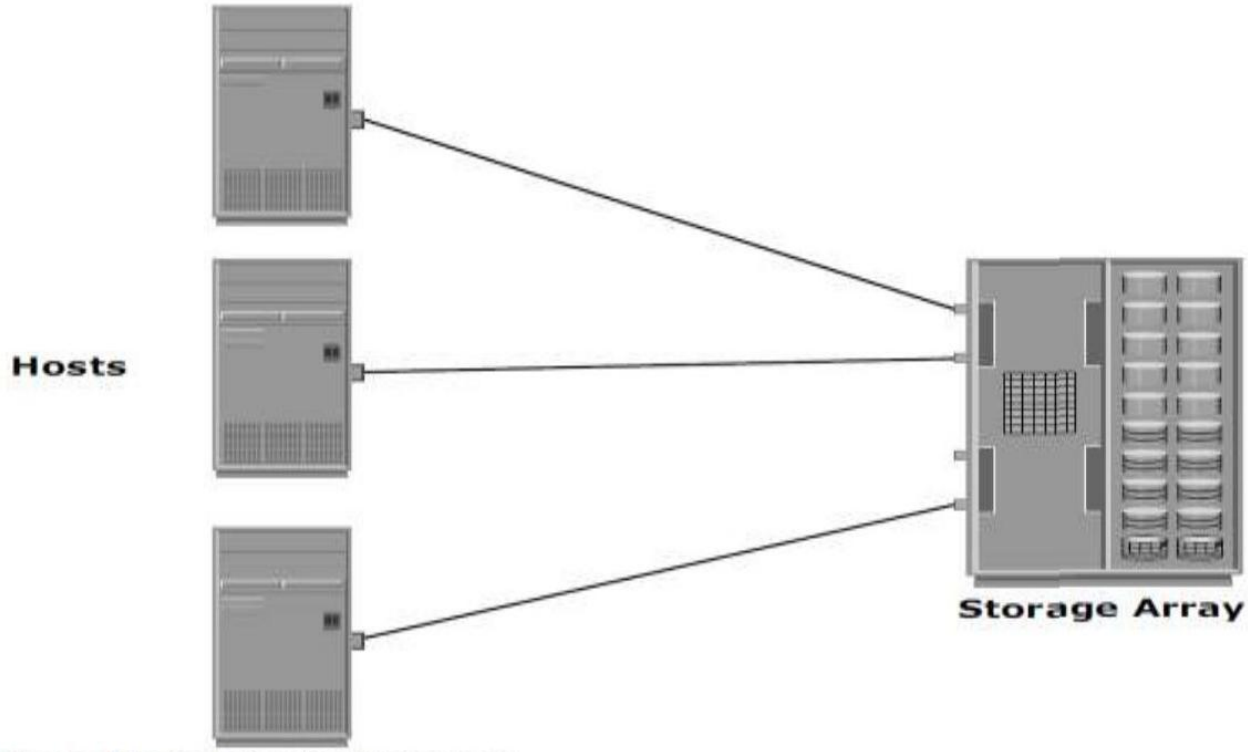

Figure 5-1: External DAS architecture

DAS Benefits and Limitations

Benefits

1. DAS requires a relatively lower initial investment than storage networking.
2. DAS configuration is simple and can be deployed easily and rapidly.
3. Setup is managed using host-based tools, such as the host OS, which makes storage management tasks easy for small and medium enterprises.
4. DAS is the simplest solution when compared to other storage networking models and requires fewer management tasks, and less hardware and software elements to set up and operate.

Limitations

1. DAS does not scale well.
2. A storage device has a limited number of ports, which restricts the number of hosts that can directly connect to the storage.
3. Limited bandwidth in DAS restricts the available I/O processing capability
4. When capacities are being reached, the service availability may be compromised, and this has a ripple effect on the performance of all hosts attached to that specific device or array.
5. DAS does not make optimal use of resources due to its limited ability to share front end ports.
6. In DAS environments, unused resources cannot be easily re-allocated, resulting in islands of over-utilized and under-utilized storage pools.