CMR
INSTITUTE OF
TECHNOLOGY

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

Improvement Test

Sub:	STORAGE AREA NETWORKS							Code:	10IS765
Date	19/ 11 / 2016	Duration:	90 mins	Max Marks	50	Sem	VII A & B	Branch	ISE/CSE

Answer any five full Questions

		Marks	OBE	
			CO	RBT
1	SAN security architecture .Explain in detail	[10]	CO5	L2
2 (a)	How flow control works in FC?	[05]	CO3	L2
(b)	What are the considerations for performing backup from local replica?	[05]	CO6	L4
3 (a)	Explain the following terms i)MTTR ii) MTBF iii) RTO iv)RPO	[08]	CO4	L2
(b)	List out the Benefits of NAS	[02]	CO3	L1
4	Develop a checklist for auditing the security of a storage environment with SAN,NAS and iSCSI implementation .Explain how you will perform the audit .Assume that you discover at least 5 security loopholes during the audit process. List them and provide control mechanisms that should be implemented to eliminate them	[10]	CO5	L4
5	Describe Kerberos with necessary diagrams	[10]	CO5	L2
6.	Explain the components of NAS in detail	[10]	CO3	L4

Course Outcomes		PO1	PO2	PO3	PO4	PO5	PO6	PO7	PO8	PO9	PO10	PO11	PO12
CO1:	Analyze storage architectures, and the logical and physical components of storage infrastructure including storage subsystems	1	1	2	-	-	-	1	-	-	1	-	-
CO2:	Analyze RAID levels and components of an intelligent storage system	2	2	-	-	-	-	-	-	-	1	-	-
CO3:	Explain the storage networking technologies such as FC SAN, SCSI,iSCSI,FCIP&NAS	1	2	-	-	-	-	-	-	-	1	-	-
CO4:	Describe the architecture of backup/recovery and virtualization technologies	1	-	-	-	-	-	-	-	-	-	-	-
CO5:	Describe object storage and retrieval in CAS,Securing and Managing storage Infrastructure	1	2	-	-	-	-	-	-	-	-	-	-
CO6:	Differentiate between local and remote replication technologies	1	-	-	-	-	-	-	-	-	-	-	-

Cognitive level	KEYWORDS
L1	List, define, tell, describe, identify, show, label, collect, examine, tabulate, quote, name, who, when, where, etc.
L2	summarize, describe, interpret, contrast, predict, associate, distinguish, estimate, differentiate, discuss, extend
L3	Apply, demonstrate, calculate, complete, illustrate, show, solve, examine, modify, relate, change, classify, experiment, discover.
L4	Analyze, separate, order, explain, connect, classify, arrange, divide, compare, select, explain, infer.
L5	Assess, decide, rank, grade, test, measure, recommend, convince, select, judge, explain, discriminate, support, conclude, compare, summarize.

PO1 - *Engineering knowledge*; PO2 - *Problem analysis*; PO3 - *Design/development of solutions*; PO4 - *Conduct investigations of complex problems*; PO5 - *Modern tool usage*; PO6 - *The Engineer and society*; PO7- *Environment and sustainability*; PO8 – *Ethics*; PO9 - *Individual and team work*; PO10 - *Communication*; PO11 - *Project management and finance*; PO12 - *Life-long learning*

CMR
INSTITUTE OF
TECHNOLOGY

Improvement Test-SCHEMES AND SOLUTIONS

Sub:	**STORAGE AREA NETWORKS**							**Code:**	**10IS765**
Date	19/ 11 / 2016	**Duration:**	90 mins	**Max Marks**	50	**Sem**	**VII A & B**	**Branch**	**ISE/CSE**

Answer any five Questions

		Marks	OBE CO	OBE RBT
1	SAN security architecture .Explain in detail Storage networking environments are a potential target for unauthorized access, theft, and misuse because of the vastness and complexity of these environments. Therefore, security strategies are based on the *defense in depth* concept, which recommends multiple integrated layers of security. This ensures that the failure of one security control will not compromise the assets under protection. Figure 15-5 illustrates various levels (zones) of a storage networking environment that must be secured and the security measures that can be deployed. SANs not only suffer from certain risks and vulnerabilities that are unique, but also share common security problems associated with physical security and remote administrative access. In addition to implementing SAN-specific security measures, organizations must simultaneously leverage other security implementations in the enterprise. Table 15-2 provides a comprehensive list of protection strategies that must be implemented in various security zones. Note that some of the security mechanisms listed in Table 15-2 are not specific to SAN, but are commonly used data center techniques. For example, two-factor authentication is implemented widely; in a simple implementation it requires the use of a user name/password and an additional security component such as a smart card for authentication *Basic SAN Security Mechanisms* LUN masking and zoning, switch-wide and fabric-wide access control, RBAC, and logical partitioning of a fabric (Virtual SAN) are the most commonly used SAN security methods. *LUN Masking and Zoning* LUN masking and zoning are the basic SAN security mechanisms used to protect against unauthorized access to storage. LUN masking and zoning were detailed earlier in Chapter 4 and Chapter 6. Standard implementations of storage arrays mask the LUNs that are presented to a front-end storage port, based on the WWPNs of the source HBAs. A stronger variant of LUN masking may sometimes be offered whereby masking can be done on the basis of source FCIDs. Note that the FCID typically changes if the HBA is relocated across ports in the fabric. To avoid this problem, major switch vendors offer a mechanism to lock down the FCID of a given node port regardless of its location. *Hard zoning* or *port zoning* is the mechanism of choice in security-conscious environments. Unlike soft zoning or WWPN zoning, it actually filters frames to ensure that only authorized zone members can communicate. However, it lacks one significant advantage of WWPN zoning: The zoning configuration	[10]	CO5	L2

must change if the source or the target is relocated across ports in the fabric. There is a trade-off between ease of management and the security provided by WWPN zoning and port zoning.

Apart from zoning and LUN masking, additional security mechanisms such as port binding, port lockdown, port lockout, and persistent port disable can be implemented on switch ports. *Port binding* limits the number of devices that can attach to a particular switch port and allows only the corresponding switch port to connect to a node for fabric access. Port binding mitigates but does not eliminate WWPN spoofing. *Port lockdown* and *port lockout* restrict a switch port's type of initialization. Typical variants of port lockout ensure that the switch port cannot function as an E_Port and cannot be used to create an ISL, such as a rogue switch. Some variants ensure that the port role is restricted to only FL_Port, F_Port, E_Port, or a combination of these. *Persistent port disable* prevents a switch port from being enabled even after a switch reboot.

Switch-wide and Fabric-wide Access Control

As organizations grow their SANs locally or over longer distances there is a greater need to effectively manage SAN security. Network security can be configured on the FC switch by using *access control lists (ACLs)* and on the fabric by using fabric binding ACLs incorporate the device connection control and switch connection control policies. The device connection control policy specifies which HBAs and storage ports can be a part of the fabric, preventing unauthorized devices (identified by WWPNs) from accessing it. Similarly, the switch connection control policy specifies which switches are allowed to be part of the fabric, preventing unauthorized switches (identified by WWNs) from joining it.

Fabric binding prevents an unauthorized switch from joining any existing switch in the fabric. It ensures that authorized membership data exists on every switch and that any attempt to connect two switches by using an ISL causes the fabric to segment.

Role-based access control provides additional security to a SAN by preventing unauthorized management activity on the fabric for management operations. It enables the security administrator to assign roles to users that explicitly specify privileges or access rights after logging into the fabric. For example, the *zoneadmin* role is able to modify the zones on the fabric, whereas a basic user may only be able to view fabric-related information, such as port types and logged-in nodes.

Logical Partitioning of a Fabric: Virtual SAN

VSANs enable the creation of multiple logical SANs over a common physical SAN. They provide the capability to build larger consolidated fabrics and still maintain the required security and isolation between them. Figure 15-6 depicts logical partitioning in a VSAN.

Zoning should be done for each VSAN to secure the entire physical SAN. Each managed VSAN can have only one active zone set at a time. As depicted in the figure, VSAN 1 is the active zone set. The SAN administrator can create distinct VSANs other than VSAN 1 and populate each of them with switch ports. In the example, the switch ports are distributed over three VSANs: 1, 2, and 3—for the IT, Engineering, and HR divisions, respectively. A zone set is defined for each VSAN, providing connectivity for HBAs and storage ports logged into the VSAN. Therefore, each of the three divisions—Engineering, IT, and HR—has its own logical fabric. Although they share physical switching gear with other divisions, they can be managed individually as stand-alone fabrics.

VSANs minimize the impact of fabricwide disruptive events because management and control traffic on the SAN—which may include RSCNs, zone set activation events, and more—does not traverse VSAN boundaries. Therefore, VSANs are a cost-effective alternative for building isolated physical fabrics. They

contribute to information availability and security by isolating fabric events and providing a finer degree of authorization control within a single fabric.

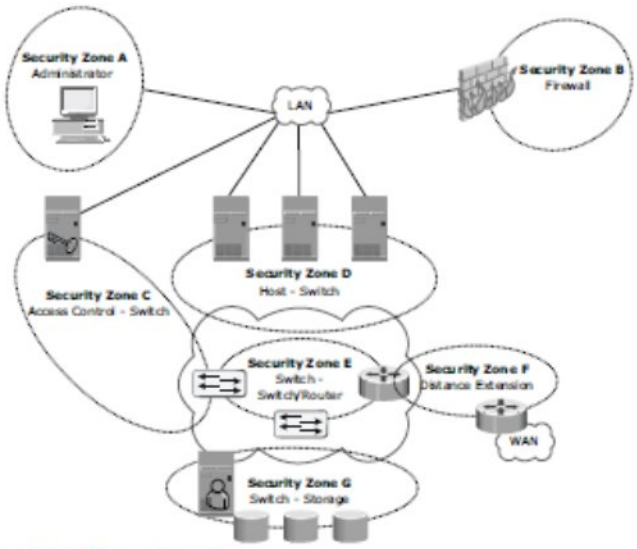

Figure 15-5: SAN security architecture

Table 15-2: Security Zones and Protection Strategies

SECURITY ZONES	PROTECTION STRATEGIES
Zone A (Authentication at the Management Console)	(a) Restrict management LAN access to authorized users (lock down MAC addresses) (b) Implement VPN tunneling for secure remote access to the management LAN (c) Use two-factor authentication for network access
Zone B (Firewall)	Block inappropriate or dangerous traffic by: (a) Filtering out addresses that should not be allowed on your LAN (b) Screening for allowable protocols—block well-known ports that are not in use
Zone C (Access Control Switch)	Authenticate users/administrators of FC switches using RADIUS (Remote Authentication Dial In User Service), DH-CHAP (Diffie-Hellman Challenge Handshake Authentication Protocol), etc.
Zone D (ACL and Zoning)	Restrict FC access to legitimate hosts by: (a) Implementing ACLs: Known HBAs can connect on specific switch ports only (b) Implementing a secure zoning method such as port zoning (also known as hard zoning)
Zone E (Switch to Switch/ Switch to Router)	Protect traffic on your fabric by: (a) Using E_Port authentication (b) Encrypting the traffic in transit (c) Implementing FC switch controls and port controls
Zone F (Distance Extension)	Implement encryption for in-flight data: (a) FCsec for long-distance FC extension (b) IPSec for SAN extension via FCIP
Zone G (Switch-Storage)	Protect the storage arrays on your SAN via: (a) WWPN-based LUN masking (b) S_ID locking: Masking based on source FCID (Fibre Channel ID/Address)

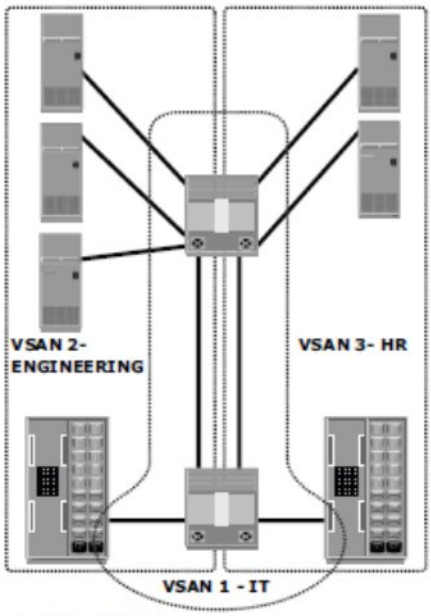

Figure 15-6: Securing SAN with VSAN

2 (a)	How flow control works in FC? Flow control defines the pace of the flow of data frames during data transmission. FC technology uses two flow-control mechanisms: buffer-to-buffer credit (BB_Credit) and end-to-end credit (EE_Credit). *BB_Credit* FC uses the *BB_Credit* mechanism for hardware-based flow control. BB_Credit controls the maximum number of frames that can be present over the link at any given point in time. In a switched fabric, BB_Credit management may take place between any two FC ports. The transmitting port maintains a count of free receiver buffers and continues to send frames if the count is greater than 0. The BB_Credit mechanism provides frame acknowledgment through the *Receiver Ready (R_RDY)* primitive. *EE_Credit* The function of end-to-end credit, known as EE_Credit, is similar to that of BB_ Credit. When an initiator and a target establish themselves as nodes communicating with each other, they exchange the EE_Credit parameters (part of Port Login). The EE_Credit mechanism affects the flow control for class 1 and class 2 traffic only	[05]	CO3	L2
(b)	What are the considerations for performing backup from local replica? `The replica should be consistent PIT copy of the source Replica should not be updated when the backup window is open` File systems buffer data in host memory to improve application response time. The buffered information is periodically written to disk. In UNIX operating systems, the *sync daemon* is the process that flushes the buffers to disk at set intervals. In some cases, the replica may be created in between the set intervals. Hence, the host memory buffers must be flushed to ensure data consistency on	[05]	CO6	L4

	the replica, prior to its creation. Figure 13-1 illustrates flushing of the buffer to its source, which is then replicated. If the host memory buffers are not flushed, data on the replica will not contain the information that was buffered in the host. If the file system is unmounted prior to the creation of the replica, the buffers would be automatically flushed and data would be consistent on the replica. If a mounted file system is replicated, some level of recovery such as *fsck* or *log replay* would be required on the replicated file system. When the file system replication process is completed, the replica file system can be mounted for operationaA database may be spread over numerous files, file systems, and devices. All of these must be replicated consistently to ensure that the replica is restorable and restartable. Replication can be performed with the database offline or online. If the database is offline, it is not available for I/O operations. Because no updates are occurring, the replica will be consistent. If the database is online, it is available for I/O operations. Transactions to the database will be updating data continuously. When a database is backed up while it is online, changes made to the database at this time must be applied to the backup copy to make it consistent. Performing an online backup requires additional procedures during backup and restore. Often these procedures can be scripted to automate the process, alleviating administrative work and minimizing human error. Most databases support some form of online or hot backups. There will be increased logging activity during the time when the database is in the hot backup mode. The sequence of operations in a hot backup mode is first to issue a database checkpoint to flush buffers to disk and place the database in hot backup mode. After taking a PIT copy, the database is taken out of hot backup mode. Logs collected are then applied to the replica to restore database consistently.l use. Creating a PIT copy for multiple devices happens quickly, but not instantaneously. It is possible that I/O transactions 3 and 4 were copied to the replica devices, but I/O transactions 1 and 2 were not copied. In this case, the data on the replica is inconsistent with the data on the source. If a restart were to be performed on the replica devices, I/O 4, which is available on the replica, might indicate that a particular transaction is complete, but all the data associated with the transaction will be unavailable on the replica, making the replica inconsistent. Another way to ensure consistency is to make sure that write I/O to all source devices is held for the duration of creating the replica. This creates a consistent image on the replica. Note that databases and applications can time out if the I/O is held for too long			
3 (a)	Explain the following terms i)MTTR ii) MTBF iii) RTO iv)RPO ▪ ▪ **Mean Time Between Failure (MTBF):** It is the average time available for a system or component to perform its normal operations between failures. ▪ ▪ **Mean Time To Repair (MTTR):** It is the average time required to repair a failed component. While calculating MTTR, it is assumed that the fault responsible for the failure is correctly identified and that the required spares and personnel are available. Note that a fault is a physical defect **Recovery-Point Objective (RPO):** This is the point in time to which systems and data must be recovered after an outage. It defines the amount of data loss that a business can endure. A large RPO signifies high tolerance to information loss in a business. Based on the RPO, organizations plan for	[08]	CO4	L2

	the minimum frequency with which a backup or replica must be made. For example, if the RPO is six hours, backups or replicas must be made at least once in 6 hours. Figure 11-2 shows various RPOs and their corresponding ideal recovery strategies. An organization can plan for an appropriate BC technology solution on the basis of the RPO it sets. For example: ▪▪ **RPO of 24 hours:** This ensures that backups are created on an offsite tape drive every midnight. The corresponding recovery strategy is to restore data from the set of last backup tapes. ▪▪ **RPO of 1 hour:** This ships database logs to the remote site every hour. The corresponding recovery strategy is to recover the database at the point of the last log shipment. ▪▪ **RPO of zero:** This mirrors mission-critical data synchronously to a remote site. ▪▪ **Recovery-Time Objective (RTO):** The time within which systems, applications, or functions must be recovered after an outage. It defines the amount of downtime that a business can endure and survive. Businesses can optimize disaster recovery plans after defining the RTO for a given data center or network. For example, if the RTO is two hours, then use a disk backup because it enables a faster restore than a tape backup. However, for an RTO of one week, tape backup will likely meet requirements. Some examples of RTOs and the recovery strategies to ensure data availability are listed below (refer to Figure 11-2): ▪▪ **RTO of 72 hours:** Restore from backup tapes at a cold site. ▪▪ **RTO of 12 hours:** Restore from tapes at a hot site. ▪▪ **RTO of 4 hours:** Use a data vault to a hot site. ▪▪ **RTO of 1 hour:** Cluster production servers with controller-based disk mirroring. ▪▪ **RTO of a few seconds:** Cluster production servers with bidirectional mirroring, enabling the applications to run at both sites simultaneously.			
(b)	List out the Benefits of 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.	[02]	CO3	L1

	▪ ▪ **Security:** Ensures security, user authentication, and file locking in conjunction with industry-standard security schemas.			
4	Develop a checklist for auditing the security of a storage environment with SAN,NAS and iSCSI implementation .Explain how you will perform the audit .Assume that you discover at least 5 security loopholes during the audit process. List them and provide control mechanisms that should be implemented to eliminate them	[10]	CO5	L4
5	Describe Kerberos with necessary diagrams(8 M FOR 8 points and 2 marks for diagram) Kerberos is a network authentication protocol. It is designed to provide strong authentication for client/server applications by using secret-key cryptography. It uses cryptography so that a client and server can prove their identity to each other across an insecure network connection. After the client and server have proven their identity, they can choose to encrypt all of their communications to ensure privacy and data integrity. In Kerberos, all authentications occur between clients and servers. The client gets a ticket for a service, and the server decrypts this ticket by using its secret key. Any entity, user, or host that gets a service ticket for a Kerberos service is called a *Kerberos client*. The term *Kerberos server* generally refers to the Key Distribution Center (KDC). The KDC implements the Authentication Service (AS) and the Ticket Granting Service (TGS). The KDC has a copy of every password associated with every principal, so it is absolutely vital that the KDC remain secure. In a NAS environment, Kerberos is primarily used when authenticating against a Microsoft Active Directory domain although it can be used to execute security functions in UNIX environments. The Kerberos authorization process shown in Figure 15-8 includes the following steps: 1. The user logs on to the workstation in the Active Directory domain (or forest) using an ID and a password. The client computer sends a request to the AS running on the KDC for a Kerberos ticket. The KDC verifies the user's login information from Active Directory. (Note that this step is not explicitly shown in Figure 15-8.) 2. The KDC responds with a TGT (TKT is a key used for identification and has limited validity period). It contains two parts, one decryptable by the client and the other by the KDC. 3. When the client requests a service from a server, it sends a request, consist of the previously generated TGT and the resource information, to the KDC. 4. The KDC checks the permissions in Active Directory and ensures that the user is authorized to use that service. 5. The KDC returns a service ticket to the client. This service ticket contains fields addressed to the client and to the server that is hosting the service. 6. The client then sends the service ticket to the server that houses the desired resources. 7. The server, in this case the NAS device, decrypts the server portion of the ticket and stores the information in a keytab file. As long as the client's Kerberos ticket is valid, this authorization process does not need to be repeated. The server automatically allows the client to access the appropriate resources. 8. A client/server session is now established. The server returns a session ID to the client, which is used to track client activity, such as file locking, as long as the session is active.	[10]	CO5	L2

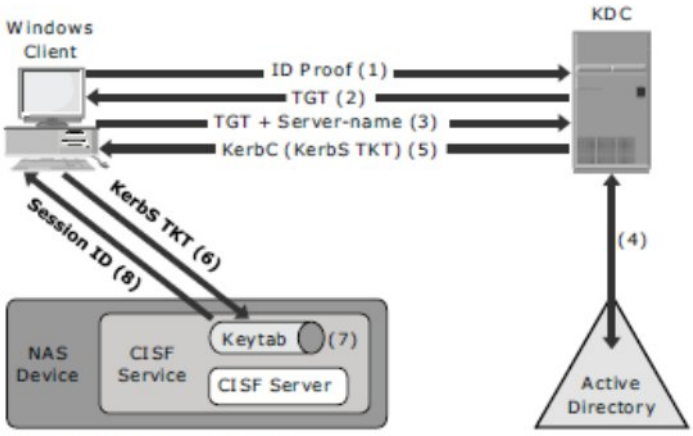

Figure 15-8: Kerberos authorization

6. Explain the components of NAS in detail [10] CO3 L4

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.

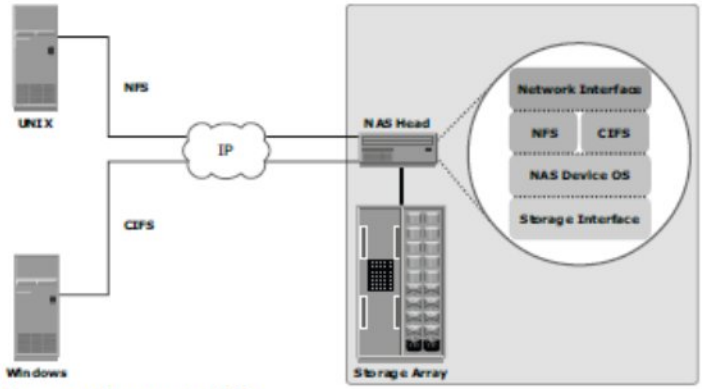

Figure 7-3: Components of NAS

NFS

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

- Searching files and directories
- Opening, reading, writing to, and closing a file
- Changing file attributes
- Modifying file links and directories

NFS uses the mount protocol to create a connection between the client and the

remote system to transfer data. NFS (NFSv3 and earlier) is a *stateless* protocol, which means that it does not maintain any kind of table to store information about open files and associated pointers. Therefore, each call provides a full set of arguments to access files on the server. These arguments include a file name and a location, a particular position to read or write, and the versions of NFS. Currently, three versions of NFS are in use:

- **NFS version 2 (NFSv2):** Uses UDP to provide a stateless network connection between a client and a server. Features such as locking are handled outside the protocol.
- **NFS version 3 (NFSv3):** The most commonly used version, it uses UDP or TCP, and is based on the stateless protocol design. It includes some new features, such as a 64-bit file size, asynchronous writes, and additional file attributes to reduce re-fetching.
- **NFS version 4 (NFSv4):** This version uses TCP and is based on a stateful protocol design. It offers enhanced security.

7.6.2 CIFS

CIFS is a client/server application protocol that enables client programs to make requests for files and services on remote computers over TCP/IP. It is a public, or open, variation of Server Message Block (SMB) protocol.

The CIFS protocol enables remote clients to gain access to files that are on a server. CIFS enables file sharing with other clients by using special locks. File names in CIFS are encoded using unicode characters. CIFS provides the following features to ensure data integrity:

It uses file and record locking to prevent users from overwriting the work of another user on a file or a record.

- It runs over TCP.
- It supports fault tolerance and can automatically restore connections and reopen files that were open prior to interruption. The fault tolerance features of CIFS depend on whether an application is written to take advantage of these features. Moreover, CIFS is a stateful protocol because the CIFS server maintains connection information regarding every connected client. In the event of a network failure or CIFS server failure, the client receives a disconnection notification. User disruption is minimized if the application has the embedded intelligence to restore the connection. However, if the embedded intelligence is missing, the user has to take steps to reestablish the CIFS connection.

Course Outcomes		PO1	PO2	PO3	PO4	PO5	PO6	PO7	PO8	PO9	PO10	PO11	PO12
CO1:	Analyze storage architectures, and the logical and physical components of storage infrastructure including storage subsystems	1	1	2	-	-	-	1	-	-	1	-	-
CO2:	Analyze RAID levels and components of an intelligent storage system	2	2	-	-	-	-	-	-	-	1	-	-
CO3:	Explain the storage networking technologies such as FC SAN, SCSI,iSCSI,FCIP&NAS	1	2	-	-	-	-	-	-	-	1	-	-
CO4:	Describe the architecture of backup/recovery and virtualization technologies	1	-	-	-	-	-	-	-	-	-	-	-
CO5:	Describe object storage and retrieval in CAS,Securing and Managing storage Infrastructure	1	2	-	-	-	-	-	-	-	-	-	-
CO6:	Differentiate between local and remote replication technologies	1	-	-	-	-	-	-	-	-	-	-	-

Cognitive level	KEYWORDS
L1	List, define, tell, describe, identify, show, label, collect, examine, tabulate, quote, name, who, when, where, etc.
L2	summarize, describe, interpret, contrast, predict, associate, distinguish, estimate, differentiate, discuss, extend
L3	Apply, demonstrate, calculate, complete, illustrate, show, solve, examine, modify, relate, change, classify, experiment, discover.
L4	Analyze, separate, order, explain, connect, classify, arrange, divide, compare, select, explain, infer.
L5	Assess, decide, rank, grade, test, measure, recommend, convince, select, judge, explain, discriminate, support, conclude, compare, summarize.

PO1 - *Engineering knowledge*; PO2 - *Problem analysis*; PO3 - *Design/development of solutions*; PO4 - *Conduct investigations of complex problems*; PO5 - *Modern tool usage*; PO6 - *The Engineer and society*; PO7- *Environment and sustainability*; PO8 – *Ethics*; PO9 - *Individual and team work*; PO10 - *Communication*; PO11 - *Project management and finance*; PO12 - *Life-long learning*