Internal Assessment Test III – Solution

Sub:	Storage Area Network					Sub Code:	18CS822	Branch:	ISE		
Date:	18-06-2022	Duration:	90 min's	Max Marks:	50	Sem/Sec:	VIII A, B & C			OBE	

	Answer any FIVE FULL Questions	MARKS	CO	RBT
1	Illustrates the concept of Kerberos with a neat diagram. **Solution:** 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 In secure 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 includes the following steps:** 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. 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.	10	CO3	L2

	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. 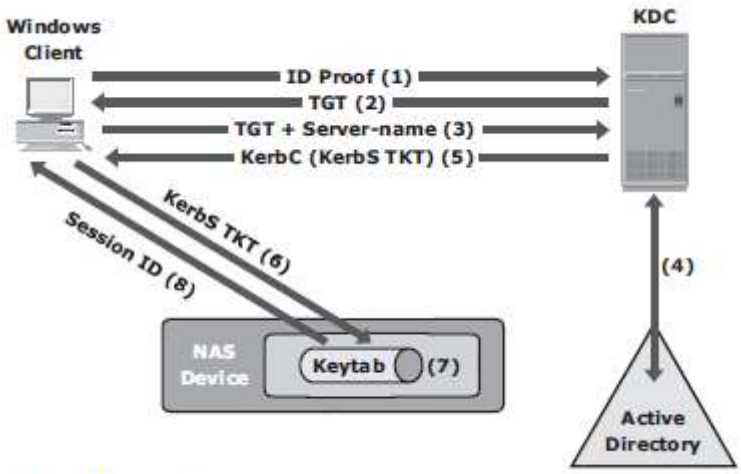 **Figure 14-8:** Kerberos authorization			
2	List and explain uses of local replicas and Replica consistency. **Alternate source for backup** • Production device burdened with simultaneously involved in production operations and servicing data for backup operations	10	CO3	L2

- Local replica contains exact point-in-time (PIT) copy of source data, can be used for backup

Fast recovery

If data loss / corrupt on the source, local replica can be used to recover the lost / corrupted data

Decision support activities

Reporting → Reduce I/O burden on the production device

Testing platform

If test is successful, the upgrade can be implemented on the production

Data Migration

From smaller capacity to larger capacity

Consistency of a Replicated database:

- 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 the replica, prior to its creation.

Figure 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.

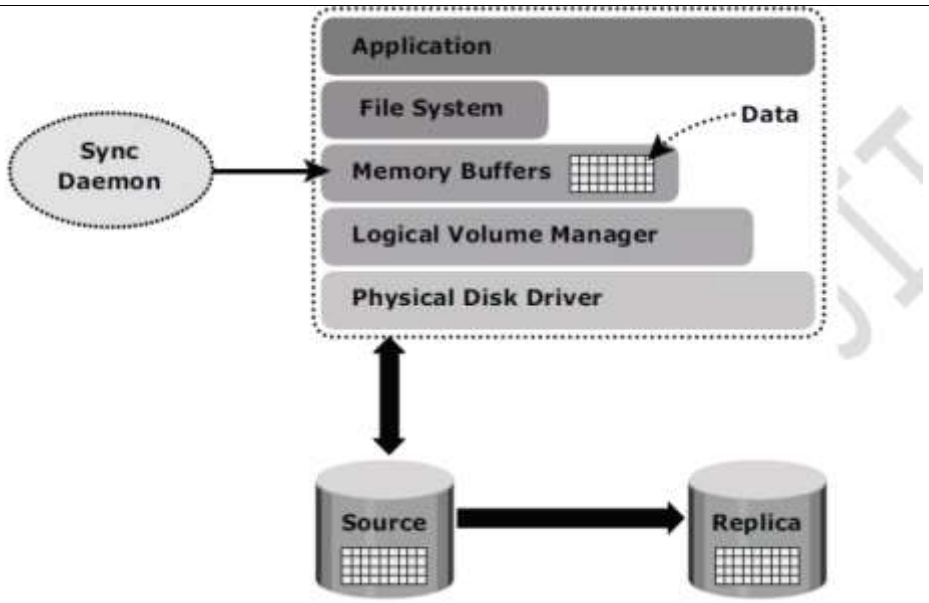

Fig. Flushing the file system buffer

Consistency of a Replicated Database

A 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.

An alternate approach exploits the dependent write I/O principle inherent in any database management system (DBMS).

According to this principle, a write I/O is not issued by an application until a prior related write I/O has completed. For example, a data write is dependent on the successful completion of the prior log write.

Dependent write consistency is required for protection against power outages, loss of local channel connectivity, or storage devices.

When the failure occurs a dependent write consistent image is created. A restart transforms the dependent write consistent image to a transactional consistent image — i.e., committed transactions are recovered, and in-flight transactions are discarded.

In order for a transaction to be deemed complete, databases require that a series of writes have to occur in a particular order.

These writes would be recorded on the various devices/file systems.

Figure illustrates the process of flushing the buffer from host to source; I/Os 1 to 4 must complete, in order for the transaction to be considered complete. I/O 4 is dependent on I/O 3 and will occur only if I/O 3 is complete. I/O 3 is dependent on I/O 2, which in turn depends on I/O 1. Each I/O completes only after completion of the previous I/O(s).

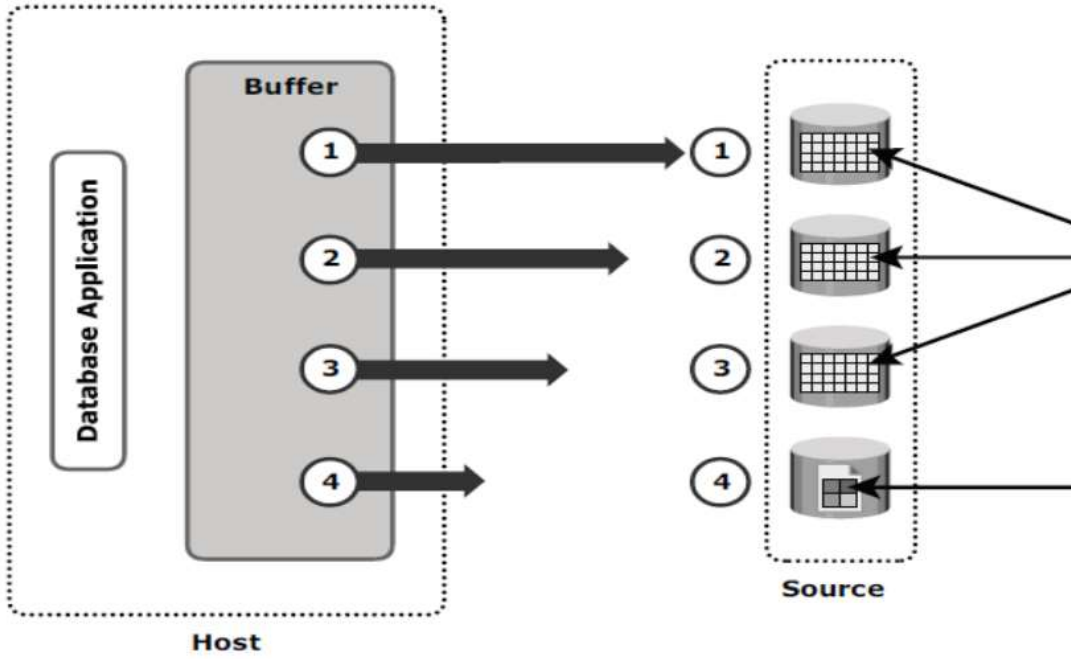

Fig. Dependent write consistency on sources

At the point in time when the replica is created, all the writes to the source devices must be captured on the replica devices to ensure data consistency. Figure, illustrates the process of replication from source to replica, I/O transactions 1 to 4 must be carried out in order for the data to be consistent on the replica.

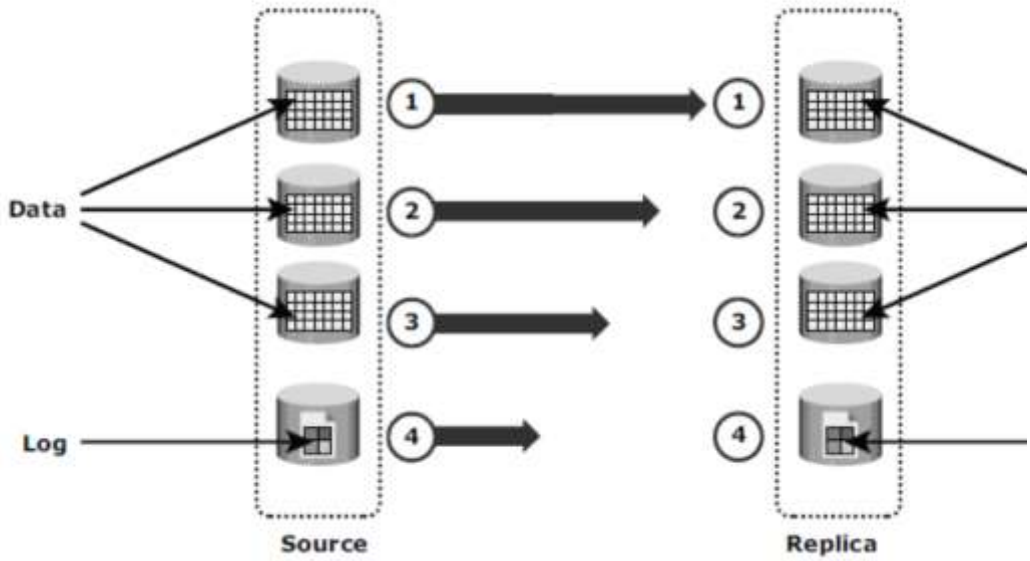

	Fig. Dependent write consistency on replica			
3	Illustrates backup and recovery operations with a neat diagram. **Solution:** **1.** 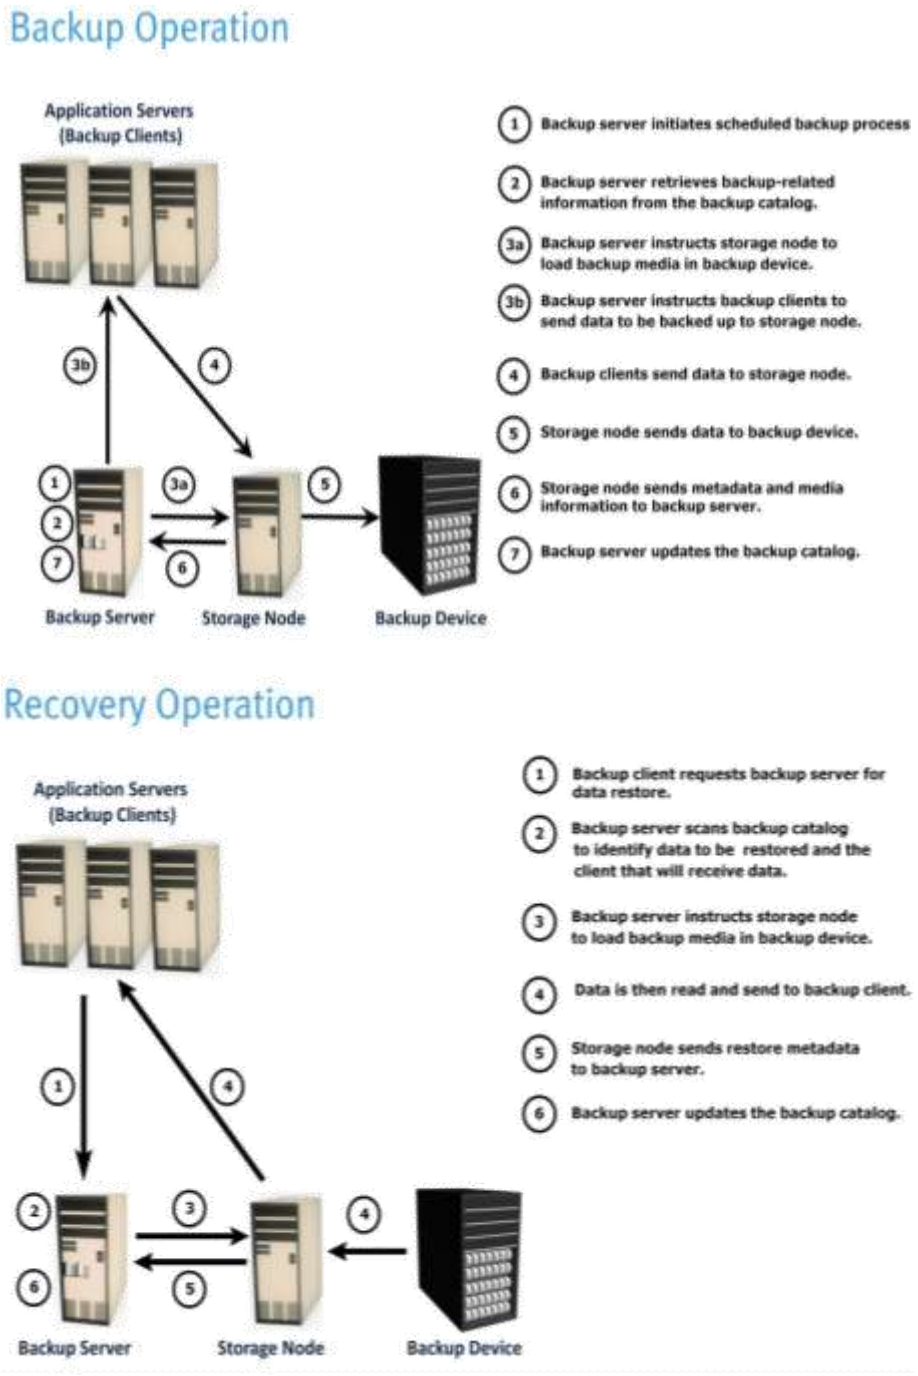	10	CO3	L2
4		10	CO3	L2

a) List and explain the replica terminologies.

Solution:

Source: A host accessing the production data from one or more LUNs on the storage array is called a production host, and these LUNs are known as source LUNs (devices/volumes), production LUNs, or simply the source.

Target: A LUN (or LUNs) on which the production data is replicated, is called the target LUN or simply the target or replica.
Point-in-Time (PIT) and continuous replica: Replicas can be either a PIT or a continuous copy. The PIT replica is an identical image of the source at some specifi c timestamp. For example, if a replica of a fi le system is created at 4:00 p.m. on Monday, this replica is the Monday 4:00 p.m. PIT copy. On the other hand, the continuous replica is in-sync with the production data at all times.
Recoverability and restartability: Recoverability enables restoration of data from the replicas to the source if data loss or corruption occurs.
Restartability enables restarting business operations using the replicas. The replica must be consistent with the source so that it is usable for both recovery and restart operations.

b) List and explain the uses of local Replicas

Alternative source for backup: Under normal backup operations, data is read from the production volumes (LUNs) and written to the backup device. This places an additional burden on the production infrastructure because production LUNs are simultaneously involved in production operations and servicing data for backup operations. The local replica contains an exact point-in-time (PIT) copy of the source data, and therefore can be used as a source to perform backup operations. This alleviates the backup I/O workload on the production volumes. Another benefi t of using local replicas for backup is that it reduces the backup window to zero.

Fast recovery: If data loss or data corruption occurs on the source, a local replica might be used to recover the lost or corrupted data. If a complete failure of the source occurs, some replication solutions enable a replica to be used to restore data onto a different set of source devices, or production can be restarted on the replica. In either case, this method provides faster recovery and minimal RTO compared to traditional recovery from tape backups. In many instances, business operations can be started using the source device before the data is completely copied from the replica.

Decision-support activities, such as reporting or data warehousing: Running the reports using the data on the replicas greatly reduces the I/O burden placed on the production device. Local replicas are also used for data-warehousing applications. The data-warehouse application may be populated by the data on the replica and thus avoid the impact on the

	production environment. **Testing platform:** Local replicas are also used for testing new applications or upgrades. For example, an organization may use the replica to test the production application upgrade; if the test is successful, the upgrade may be implemented on the production environment. **Data migration:** Another use for a local replica is data migration. Data migrations are performed for various reasons, such as migrating from a smaller capacity LUN to one of a larger capacity for newer versions of the application.			
5	Analyze LVM based local replication with a neat diagram. Discuss the advantages and Disadvantages. **Solution:** **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. In LVM-based replication, logical volume manager is responsible for creating and controlling the host-level logical volume. An LVM has three components: physical volumes (physical disk), volume groups, and logical volumes. A volume group is created by grouping together one or more physical volumes. Logical volumes are created within a given volume group. A volume group can have multiple logical volumes. In LVM-based replication, each logical partition in a logical volume is mapped to two physical partitions on two different physical volumes, as shown in Figure. An application write to a logical partition is written to the two physical partitions by the LVM device driver. This is also known as LVM mirroring. Mirrors can be split and the data contained therein can be independently accessed. LVM mirrors can be added or removed dynamically.	10	CO3	L3

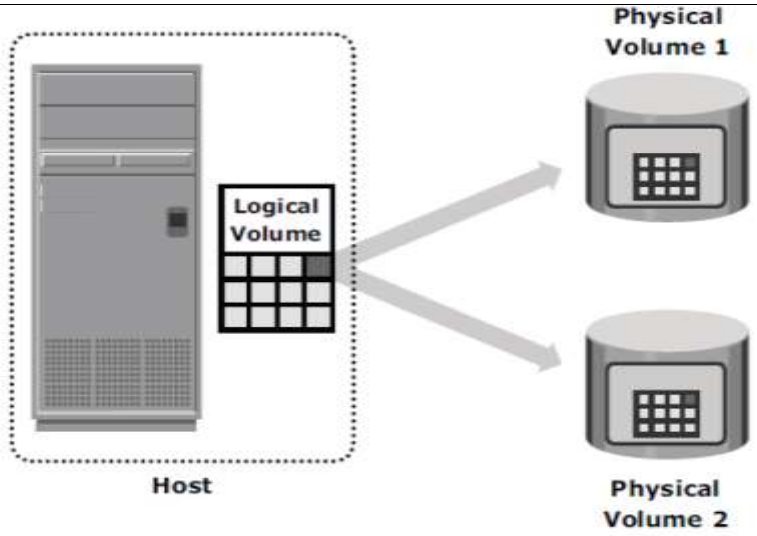

Fig: LVM Based Replication

Advantage:

The LVM-based replication technology is not dependent on a vendor-specific storage system. Typically,

LVM is part of the operating system and no additional license is required to deploy LVM mirroring.

Disadvantage:

As every write generated by an application translates into two writes on the disk, an additional burden is placed on the host CPU. This can degrade application performance.

Presenting an LVM-based local replica to a second host is usually not possible because the replica will still be part of the volume group, which is usually accessed by one host at any given time.

Tracking changes to the mirrors and performing incremental synchronization operations is also a challenge as all LVMs do not support incremental resynchronization. If the devices are already protected by some level of RAID on the array, then the additional protection provided by mirroring is unnecessary.

6	List and define the basic security elements can be applied to different aspects of SAN. **Solution:** **Risk triad** defines risk in terms of threats, assets, and vulnerabilities. Risk arises when a threat agent (an attacker) uses an existing vulnerability to compromise the security services of an asset, for example, if a sensitive document is transmitted without any protection (or encryption) over an insecure channel, an attacker might get unauthorized access to the document and may violate its confidentiality and integrity. This may, in turn, result in business loss for the organization. In this scenario potential business loss is the risk, which arises because an attacker uses vulnerability of the unprotected communication to access the document and tamper with it. To manage risks, organizations primarily focus on vulnerabilities because they cannot eliminate threat agents that appear in various forms and sources to its assets. Organizations can enforce	10	CO3	L2

countermeasures to reduce the possibility of occurrence of attacks and the severity of their impact.

Risk assessment is the first step to determine the extent of potential threats and risks in an IT infrastructure. The process assesses risk and helps to identify appropriate controls to mitigate or eliminate risks. To determine the probability of an adverse event occurring, threats to an IT system must be analyzed with the potential vulnerabilities and the existing security controls.

The severity of an adverse event is estimated by the impact that it may have on critical business activities. Based on this analysis, a relative value of criticality and sensitivity can be assigned to

IT assets and resources. For example, a particular IT system component may be assigned a high-criticality value if an attack on this particular component can cause a complete termination of mission-critical services.

Information is one of the most important ***assets*** for any organization. Other assets include hardware, software, and the network infrastructure required to access the information. To protect these assets, organizations must develop a set of parameters to ensure the availability of the resources to authorized users and trusted networks. These parameters apply to storage resources, network infrastructure, and organizational policies.

Security methods have two objectives. The first objective is to ensure that the network is easily accessible to authorized users. It should also be reliable and stable under disparate environmental conditions and volumes of usage. The second objective is to make it difficult for potential attackers to access and compromise the system.

The security methods should provide adequate protection against unauthorized access, viruses, worms, Trojans, and other malicious software programs. Security measures should also encrypt critical data and disable unused services to minimize the number of potential security gaps. The security method must ensure that updates to the operating system and other software are installed regularly. At the same time, it must provide adequate redundancy in the form of replication and mirroring of the production data to prevent catastrophic data loss if there is an unexpected data compromise. For the security system to function smoothly, all users are informed about the policies governing the use of the network.

The effectiveness of a storage security methodology can be measured by two criteria. One, the cost of implementing the system should be only a small fraction of the value of the protected data. Two, it should cost a potential attacker more, in terms of money and time, to compromise the system than the value of the protected data.

Threats are the potential attacks that can be carried out on an IT infrastructure. These attacks can be classified as active or passive. *Passive attacks* are attempts to gain unauthorized access into the system. They pose threats to confidentiality of information. *Active attacks* include data modification, denial of service (DoS), and repudiation attacks. They pose threats to data integrity and availability.

In a data modification attack, the unauthorized user attempts to modify information for malicious purposes. A modification attack can target the data at rest or the data in transit. These attacks pose a threat to data integrity.

Denial of service **(DoS)** attacks prevent legitimate users from accessing resources and services. . These attacks generally do not involve access to or modification of information on the computer system. Instead, they pose a threat to data availability. The intentional flooding of a network or website to prevent legitimate access to authorized users is one example of a DoS attack.

Repudiation is an attack against the accountability of information. It attempts to provide false information by either impersonating someone or denying that an event or a transaction has taken place. For example, a repudiation attack may involve performing an action and eliminating any evidence that could prove the identity of the user (attacker) who performed that action. Repudiation attacks include

	circumventing the logging of security events or tampering with the security log to conceal the identity of the attacker. **Vulnerabilities** The paths that provide access to information are the most vulnerable to potential attacks. Each of the paths may contain various access points, which provide different levels of access to the storage resources. It is important to implement adequate security controls at *all* the access points on an access path. Implementing security controls at each access point of every access path is known as *defense in depth*. Defense in depth recommends using multiple security measures to reduce the risk of security threats if one component of the protection is compromised**.** It is also known as a "layered approach measures for security at different levels, defense in depth gives additional time to detect and respond to an attack. This reduces the scope of a security breach.			