Internal Assessment Test III – May 2023-Solution

Sub:	Storage Area Network					Sub Code:	18CS822	Branch:	ISE		
Date:	12-05-2023	Duration:	90 min's	Max Marks:	50	Sem/Sec:	VIII A& B			OBE	

	Answer any FIVE FULL Questions	MARKS	CO	RBT
1	Describe the various security concerns related to risk triad. **Solution:** **Definition:** Risk triad defines the risk in terms of threats, assets, and vulnerabilities. Risk arises when a threat agent (an attacker) seeks to access assets by exploiting an existing vulnerability. **Risk assessment** is the first step in determining 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. **Assets:** Information is one of the most important assets for any organization. Other assets include hardware, software, and the network infrastructure required to access this 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, the network infrastructure, and organizational policies. Several factors need to be considered when planning for asset security. **Security methods have two objectives.** 1. It has 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. 2. It has to make it very difficult for potential attackers to access and compromise the system. These methods should provide adequate protection against unauthorized access to resources, viruses, worms, Trojans and other malicious software programs. The effectiveness of a storage security methodology can be measured by two criteria. One, the cost of implementing the system should only be 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 protected data is worth. **Threats:** Threats are the potential attacks that can be carried out on an IT infrastructure. These attacks can be classified as active or passive.	10	CO6	L2

	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 modification attack, the unauthorized user attempts to modify information for malicious purposes. A modification attack can target data at rest or data in transit. These attacks pose a threat to data integrity. **Denial of Service (DoS) attacks** denies the use of resources to legitimate users. These attacks generally do not involve access to or modification of information on the computer system. Instead, they pose a threat to data availability. **Repudiation** is an attack against the accountability of the information. It attempts to provide false information by either impersonating someone or denying that an event or a transaction has taken place. **Vulnerabilities:** The paths that provide access to information are the most vulnerable to potential attacks. Each of these paths may contain various access points, each of which provides different levels of access to the storage resources. **An attack vector is a step** or a series of steps necessary to complete an attack. Work factor refers to the amount of time and effort required to exploit an attack vector. **Preventive controls avert** the vulnerabilities from being exploited and prevent an attack or reduce its impact. Corrective controls reduce the effect of an attack, while detective controls discover attacks and trigger preventive or corrective controls.			
2	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 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	10	CO6	L2

	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. 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.			

	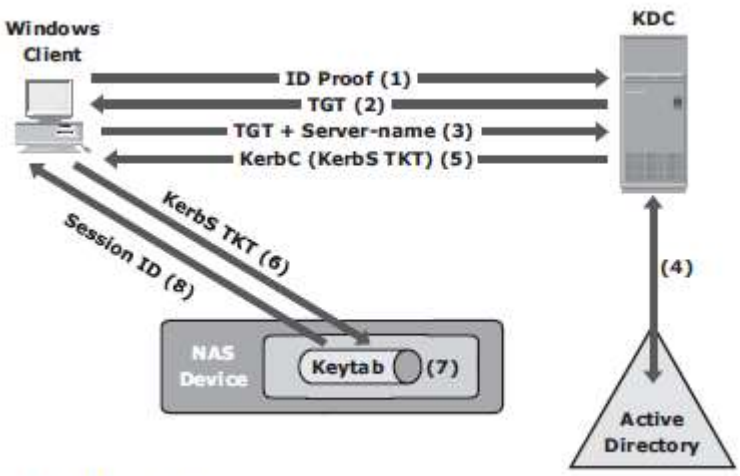 **Figure 14-8:** Kerberos authorization			
3	List and explain uses of local replicas and Replica consistency. **Solution:** One or more local replicas of the source data may be created for various purposes, including the following: **Alternative source for backup**: Under normal backup operations, data 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.	10	CO5	L2

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.

Replica Consistency:

Most file systems and databases buffer the data in the host before writing it to the disk. A consistent replica ensures that the data buffered in the host is captured on the disk when the replica is created. The data staged in the cache and not yet committed to the disk should be flushed before taking the replica. The storage array operating environment takes care of flushing its cache before the replication operation is initiated. Consistency ensures the usability of a replica and is a primary requirement for all the replication technologies.

11.3.1 Consistency of a Replicated File System

File systems buffer the data in the host memory to improve the application response time. The buffered data is periodically written to the disk. In UNIX operating systems, sync daemon is the process that flushes the buffers to the disk at set intervals. In some cases, the replica is created between the set intervals, which might result in the creation of an inconsistent replica. Therefore, host memory buffers must be flushed to ensure data consistency on the replica, prior to its creation. Figure 11-1 illustrates how the file system buffer is flushed to the source device before replication. If the host memory buffers are not flushed, the data on the replica will not contain the information that was buffered in the host. If the file system is unmounted before creating the replica, the buffers will be automatically flushed and the data will be consistent on the replica.

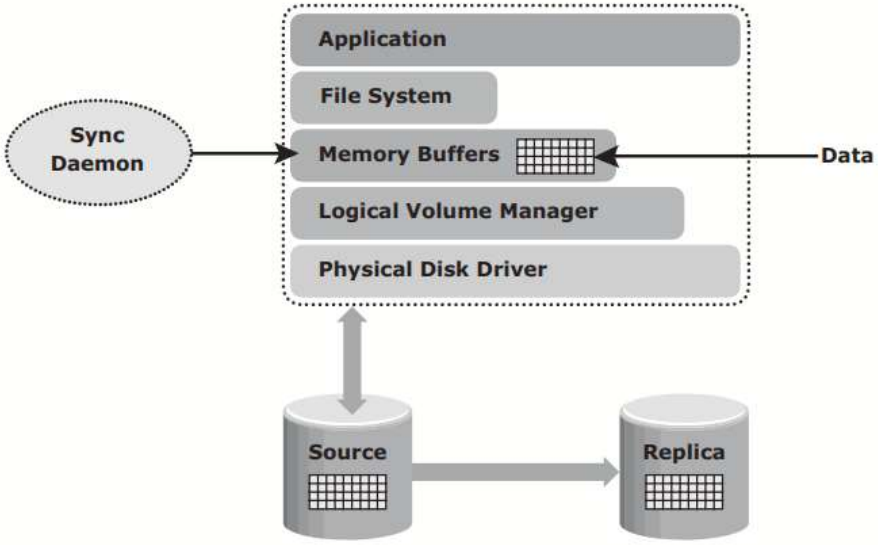

Figure 11-1: Flushing the file system buffer

If a mounted file system is replicated, some level of recovery, such as fsck or log replay, is required on the replicated file system. When the file system replication and check process are completed, the replica fi le system can be mounted for operational use.

11.3.2 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 is performed with the database offline or online. If the database is offline during the creation of the replica, it is not available for I/O operations. Because no updates occur on the source, the replica is consistent. If the database is online, it is available for I/O operations, and transactions to the database update the data continuously. When a database is replicated while it is online, changes made to the database at this time must be applied to the replica to make it consistent. A consistent replica of an online database is created by using the dependent write I/O principle or by holding I/Os momentarily to the source before creating the replica. A dependent write I/O principle is inherent in many applications and database management systems (DBMS) to ensure consistency. 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. For a transaction to be deemed complete, databases require a series of writes to have occurred in a particular order. These writes will be recorded on the various devices or fi le systems. Figure 11-2, illustrates the process of fl ushing the buffer from the host to the source; I/Os 1 to 4 must complete for the transaction to be considered complete. I/O 4 is dependent on I/O 3 and occurs 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).

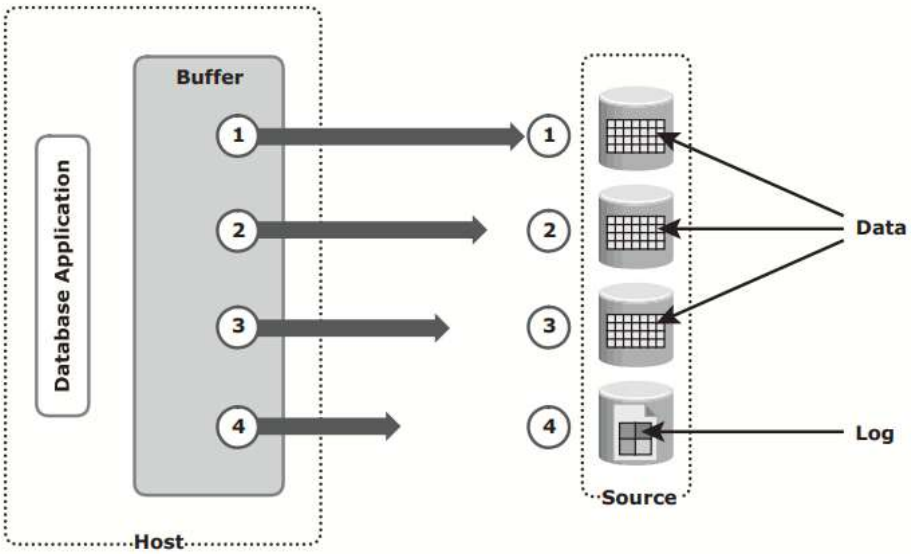

Figure 11-2: Dependent write consistency on sources

When the replica is created, all the writes to the source devices must be captured on the replica devices to ensure data consistency. Figure 11-3 illustrates the process of replication from the source to the replica. I/O transactions 1 to 4 must be carried out for the data to be consistent on the replica. 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. Figure 11-4 shows this situation.

	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. 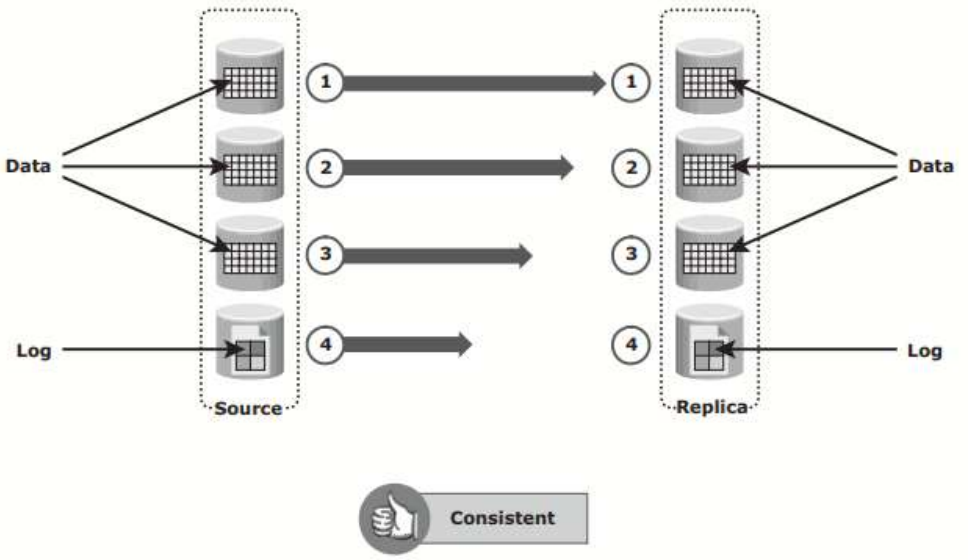 **Figure 11-3:** Dependent write consistency on replica 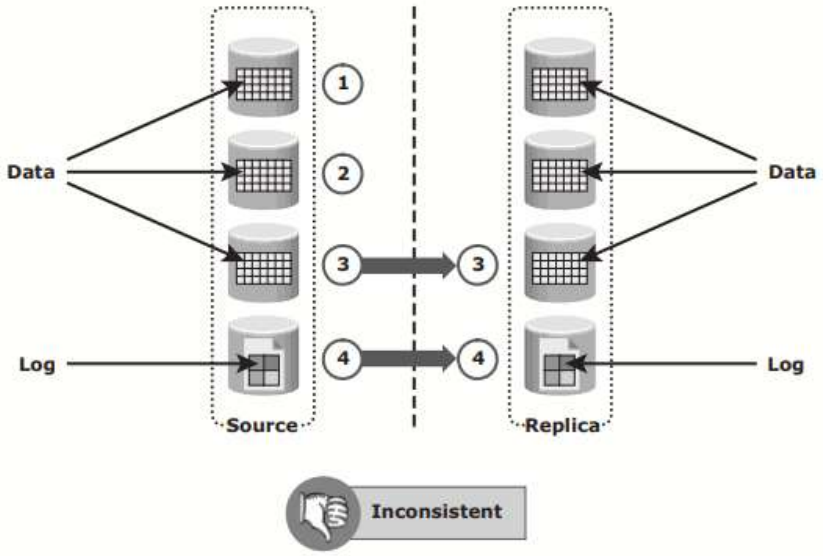 **Figure 11-4:** Inconsistent database replica			
4	Write the concept of Storage Array-Based Remote Replication and discuss. **Solution:** • In storage array-based local replication, the array-operating environment performs	10	CO5	L2

	the local replication process. The host resources, such as the 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 alternative host for other business operations. • In this replication, the required number of replica devices should be selected on the same array and then data should be replicated between the source-replica pairs. Figure 11-7 shows a storage array-based local replication, where the source and target are in the same array and accessed by different hosts. 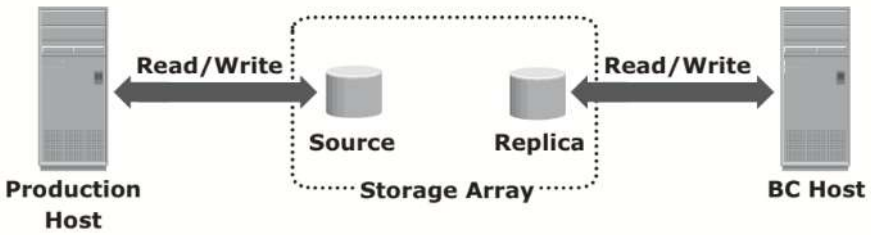 **Figure 11-7:** Storage array-based local replication • Storage array-based local replication is commonly implemented in three ways: full-volume mirroring, pointer-based full-volume replication, and pointer-based virtual replication. Replica devices are also referred as target devices, accessible by other hosts.			
5	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 countermeasures to reduce the possibility of occurrence	10	CO6	L4

	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.			
6	a. Illustrates backup and recovery operations with a neat diagram. **Solution:**	5+5	CO4	L2

Backup Operation

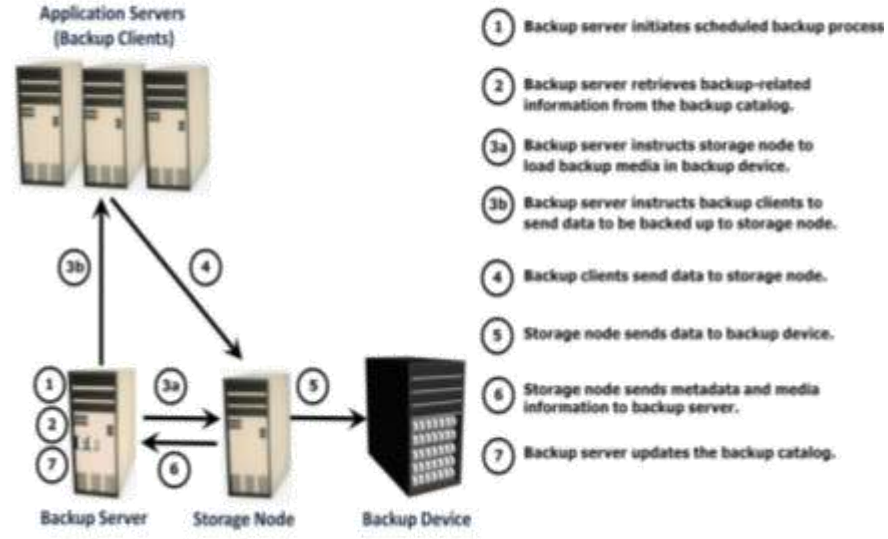

Recovery Operation

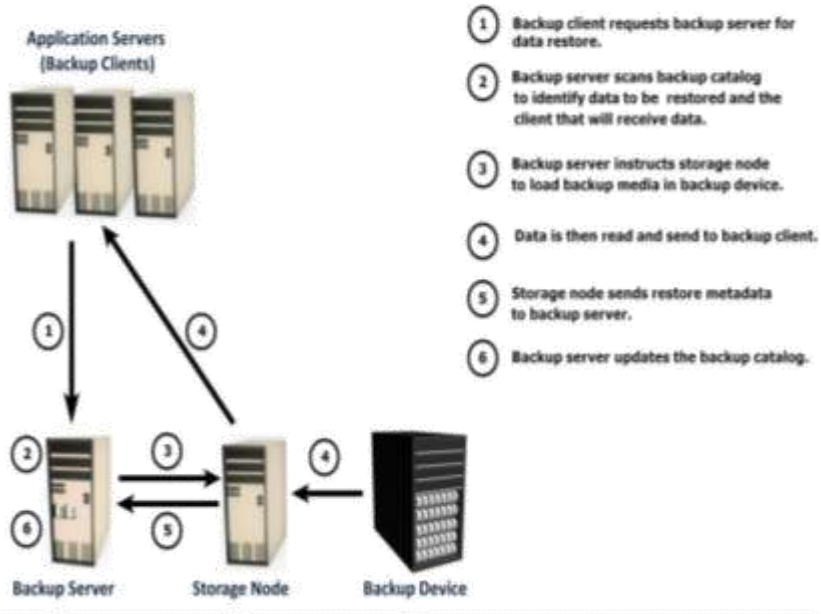

b. A system has three components and requires all three components to be operational during 8 a.m. through 5 p.m. business hours, Monday through Friday. Failure of component 2 occurs as follows:

Monday = 8 a.m. to 11 a.m.

Tuesday = No failure.

Wednesday = 4 p.m. to 7 p.m.

	Availability (%) = system uptime/(system uptime + system downtime) System downtime = 3 hours on Monday + 1 hour on Wednesday + 1 hour on Friday = 5 hours System uptime = total operational time – system downtime = 45 hours – 5 hours = 40 hours Availability (%) = 40/45 = 88.9%			

Faculty Signature CCI Signature HOD Signature