(2½ hours)

Total Marks: 60

N. B.: (1) **All** questions are **compulsory**.
(2) Make **suitable assumptions** wherever necessary and **state the assumptions** made.
(3) Answers to the **same question** must be **written together**.
(4) Numbers to the **right** indicate **marks**.
(5) Draw **neat labeled diagrams** wherever **necessary**.
(6) Use of **Non-programmable** calculators is **allowed**.

1. Attempt *any two* of the following: **30**

1. Which of the following statements about various hard disks is wrong?
   a. SATA Disks support faster transfer rates and have support for hot swapping
   b. USB Hard disks store data on flash memory only
   c. ATA hard disks cannot be connected externally to computer
   d. SCSI hardisks are more reliable
2. Which of the following is not true about JBOD?
   a. JBOD can combine hard disks of different sizes into a single unit without loss of any capacity
   b. If a drive in a JBOD set dies then it may be easier to recover the files on the other Drives
   c. JBOD supports data redundancy
   d. JBOD doesn't has any storage controller intelligence
3. Pick odd one out
   a. Mirroring
   b. Striping
   c. Error-correction
   d. Fault tolerance
4. Pick the false statement about RAID 0
   a. Provides no redundancy
   b. Provides data striping
   c. improves performance
   d. Provides fault tolerance
5. Which one of the following is not an advantage of RAID?
   a. Data Security
   b. Increased & integrated capacity
   c. Improved performance
   d. Effective data capacity increases
6. Which of the following is NOT the component of intelligent storage system?
   a. front end
   b. cache
   c. physical disks
   d. Host

7. In an intelligent storage system, front end provides the interface between _____ and _____.
   a. Cache, Physical Disk
   b. Cache, Host
   c. Storage System, Host
   d. Storage System, Physical Disks
8. On intelligent arrays, write data is first placed in _____ and then written to disk.
   a. Cache
   b. Front End
   c. Host
   d. Network
9. Cache is organized into _____, which is the smallest unit of cache allocation.
   a. Tracks
   b. Pages
   c. Sectors
   d. Segments
10. In _____, data is placed in cache, an acknowledgment is sent to the host immediately and later, data from several writes are committed (de-staged) to the disk.
   a. Write-back cache
   b. Write-through cache
   c. Write-front cache
   d. Write-away-Cache
11. A NAS solution is most appropriate for what type of data environment
   a. Secured Access
   b. Shared access
   c. Remote access
   d. Parallel access
12. Which three statements describe differences between Storage Area Network (SAN) and Network Attached Storage (NAS) solutions? Choose three.
   i. SAN is generally more expensive but provides higher performance
   ii. NAS uses TCP/IP for communication between hosts and the NAS server
   iii. NAS requires additional hardware on a host: a host bus adapter for connectivity
   iv. SAN uses proprietary protocols for communication between hosts and the SAN fabric
   a. i,ii, iii
   b. i, ii, iv
   c. ii, iii, iv
   d. i, ii, iii
13. Which topology is best suited for medium sized enterprise.
   a. NAS
   b. SAN
   c. DAS
   d. CAS

14. Identify a network file protocol in the below mentioned set.
   a. FC
   b. CIFS

c. SCSI
d. NAS

15. Following are some of the file sharing protocols
    a. Telnet
    b. NFS & CIFS
    c. FTP
    d. HTTP
16. A host accessing the production data from one or more LUNs on the storage array is called as_____
    a. Target
    b. production host
    c. Point-in-Time (PIT)
    d. continuous replica
17. _____enables restarting business operations using the replicas
    a. Restartability
    b. Recoverability
    c. production
    d. The PIT replica
18. Which one is NOT associated with Replication Terminology?
    a. Source
    b. Recovery-Point Objective
    c. Point-In-Time
    d. Target
19. A LUN on which the production data is replicated, is called as
    a. simply the target
    b. production host
    c. production target
    d. Point-in-Time (PIT)
20. Which is TRUE for Remote and Migration in Virtualized Environment?
    a. Replication process can be either synchronous or asynchronous
    b. Hypervisor-to- Hypervisor VM migration is NOT possible
    c. Array-to-Array VM migration is NOT possible
    d. Hypervisor-to-Array VM migration is NOT possible
21. Which one is not the basic information security framework is built to achieve security goals:
    a. Confidentiality
    b. Accountability service
    c. Integrity
    d. Striping
22. Ensuring _____ requires detection of and protection against unauthorized alteration or deletion of information
    a. Accountability service
    b. Confidentiality
    c. Availability
    d. integrity
23. Which one is not risk terms in Risk Triad.

a. Threats
b. Domain
c. assets
d. vulnerabilities

24. To protect _____, organizations must develop a set of parameters to ensure the availability of the resources to authorized users and trusted networks.
    a. vulnerabilities
    b. threats
    c. Domain
    d. assets
25. _____attacks are attempts to gain unauthorized access into the system.
    a. Active
    b. assets
    c. Passive
    d. Availability
26. _____ refers to the amount of time and effort required to exploit an attack vector
    a. Availability
    b. attack vector
    c. Work factor
    d. integrity
27. Which is one of technologies for Network-Based Local Replication?
    a. Target-Based Protection
    b. Post-Process Protection
    c. Continuous Data Protection
    d. File System Snapshot
28. File access protocols operate in which layer of the OSI model
    a. Application
    b. Session
    c. Transport
    d. Data link
29. _____ is the process of assigning storage resources to hosts based on capacity, availability, and performance requirements of applications running on the hosts.
    a. Storage provisioning
    b. Virtual provisioning
    c. Write-back cache
    d. Write-front cache
30. Pick the wrong statement about the hard disk?
    a. Hard disk has multiple platters and each platter has two read/write heads one on each side
    b. It's a non-volatile & random access storage device
    c. Hard disk can only have IDE or USB interface
    d. Data Transfer rate is over 80 MBPS

2. Attempt *any one* of the following: **08**

a. With the help of neat and clean diagram explain the Hard Disk components.

Platter

A typical HDD consists of one or more flat circular disks called *platters* (Figure 2-6). The data is recorded on these platters in binary codes (0s and 1s).The set of rotating platters is sealed in a case, called the *Head Disk Assembly*(HDA). A platter is a rigid, round disk coated with magnetic material on both surfaces (top and bottom). The data is encoded by polarizing the magnetic area, or domains, of the disk surface. Data can be written to or read from both surfaces of the platter. The number of platters and the storage capacity of eachplatter determine the total capacity of the drive.

Spindle

A spindle connects all the platters (refer to Figure 2-6) and is connected to a motor. The motor of the spindle rotates with a constant speed.

The disk platter spins at a speed of several thousands of revolutions per minute (rpm). Common spindle speeds are 5,400 rpm, 7,200 rpm, 10,000 rpm, and 15,000 rpm. The speed of the platter is increasing with improvements in technology, although the extent to which it can be improved is limited.

Read/Write Head

Read/Write (R/W) heads, as shown in Figure 2-7, read and write data from or to platters. Drives have two R/W heads per platter, one for each surface of the platter. The R/W head changes the magnetic polarization on the surface of the platter when writing data. While reading data, the head detects the magnetic polarization on the surface of the platter. During reads and writes, the R/W head senses the magnetic polarization and never touches the surface of the platter. When the spindle is rotating, there is a microscopic air gap maintained between the R/W heads and the platters, known as the *head fl ying height*. This air gap is removed when the spindle stops rotating and the R/W head rests on a special area on the platter near the spindle. This area is called the *landing*

Actuator Arm Assembly

R/W heads are mounted on the *actuator arm assembly*, which positions the R/W head at the location on the platter where the data needs to be written or read (refer to Figure 2-7). The R/W heads for all platters on a drive are attached to one actuator arm assembly and move across the platters simultaneously

Drive Controller Board

The controller (refer to Figure 2-5 [b]) is a printed circuit board, mounted at the bottom of a disk drive. It consists of a microprocessor, internal memory, circuitry, and fi rmware. The fi rmware controls the power to the spindle motor and the speed of the motor. It also manages the communication between the drive and the host. In addition, it controls the R/W operations by moving the actuator arm and switching between different R/W heads, and performs the optimization of data access.

b. Write short noes on

- Zone bit recording

Zoned Bit Recording

Platters are made of concentric tracks; the outer tracks can hold more data than the inner tracks because the outer tracks are physically longer than the inner tracks. On older disk drives, the outer tracks had the same number of sectors as the inner tracks, so data density was low on the outer tracks. This was an

ineffi cient use of the available space, as shown in Figure 2-9 (a).
Zoned bit recording uses the disk effi ciently. As shown in Figure 2-9 (b), this mechanism groups tracks into zones based on their distance from the center of the disk. The zones are numbered, with the outermost zone being zone 0. An appropriate number of sectors per track are assigned to each zone, so a zone near the center of the platter has fewer sectors per track than a zone on the outer

- edge. However, tracks within a particular zone have the same number of sectors.

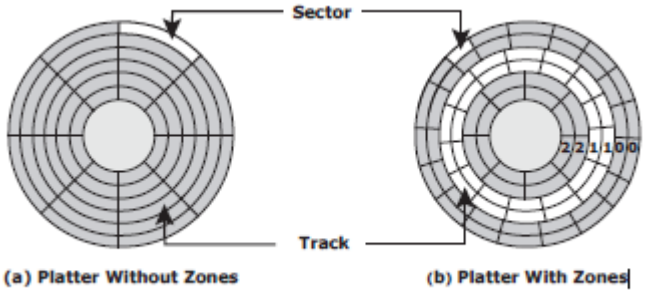

Figure 2-9: Zoned bit recording

-
- Logical block addressing

Earlier drives used physical addresses consisting of the *cylinder*, *head*, and *sector* (CHS) number to refer to specifi c locations on the disk, as shown in Figure 2-10 (a), and the host operating system had to be aware of the geometry of each disk used. *Logical block addressing* (LBA), as shown in Figure 2-10 (b), simplifi es addressing by using a linear address to access physical blocks of data. The disk controller translates LBA to a CHS address, and the host needs to know only the size of the disk drive in terms of the number of blocks. The logical blocks are mapped to physical sectors on a 1:1 basis.

In Figure 2-10 (b), the drive shows eight sectors per track, eight heads, and four cylinders. This means a total of 8 × 8 × 4 = 256 blocks, so the block number ranges from 0 to 255. Each block has its own unique address. Assuming that

- the sector holds 512 bytes, a 500 GB drive with a formatted capacity of 465.7 GB

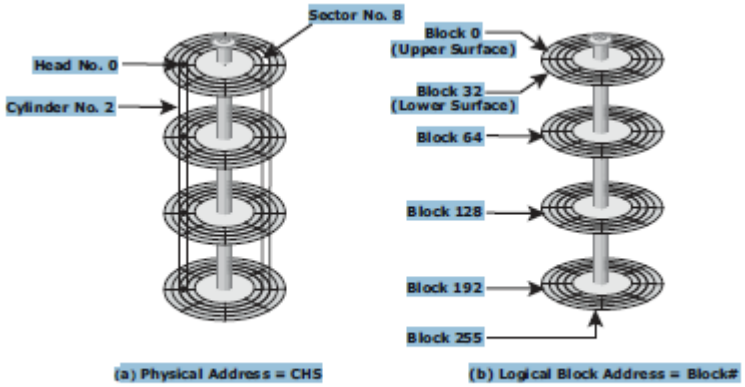

Figure 2-10: Physical address and logical block address

-

c. How nested RAID levels 0+1 and 1+0 differs from each other? Which is preferred in business continuity and why?

3. Attempt *any one* of the following: **08**

a. Write short noes on

- Write Operation with Cache

Write Operation with Cache

Write operations with cache provide performance advantages over writing directly to disks. When an I/O is written to cache and acknowledged, it is completed in far less time (from the host's perspective) than it would take to write directly to disk. Sequential writes also offer opportunities for optimization because many smaller writes can be coalesced for larger transfers to disk drives with the use of cache.

A write operation with cache is implemented in the following ways:

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

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

Cache can be bypassed under certain conditions, such as large size write I/O. In this implementation, if the size of an I/O request exceeds the predefi ned size, called *write aside size*, writes are sent to the disk directly to reduce the impact of large writes consuming a large cache space. This is particularly useful in an environment where cache resources are constrained and cache is required for

- small random I/O
- Cache Management

Cache Management

Cache is a fi nite and expensive resource that needs proper management. Even though modern intelligent storage systems come with a large amount of cache, when all cache pages are fi lled, some pages have to be freed up to accommodate new data and avoid performance degradation. Various cache management algorithms are implemented in intelligent storage systems to proactively maintain a set of free pages and a list of pages that can be potentially freed up whenever required. The most commonly used algorithms are discussed in the following list:

n **Least Recently Used (LRU):** An algorithm that continuously monitors data access in cache and identifi es the cache pages that have not been accessed for a long time. LRU either frees up these pages or marks them for reuse. This algorithm is based on the assumption that data that has not been accessed for a while will not be requested by the host. However, if a page contains write data that has not yet been committed to disk, the

data is fi rst written to disk before the page is reused.

n **Most Recently Used (MRU):** This algorithm is the opposite of LRU, where the pages that have been accessed most recently are freed up or marked for reuse. This algorithm is based on the assumption that recently accessed data may not be required for a while.

As cache fi lls, the storage system must take action to fl ush dirty pages (data written into the cache but not yet written to the disk) to manage space availability. *Flushing* is

the process that commits data from cache to the disk. On the basis of the I/O access rate and pattern, high and low levels called *watermarks* are set in cache to manage the fl ushing process. *High watermark* (*HWM*) is the cache utilization level at which the storage system starts high-speed fl ushing of cache data. *Low watermark* (*LWM*) is the point at which the storage system stops fl ushing data to the disks. The cache utilization level, as shown in Figure 4-4, drives the mode of fl ushing to be used:

n **Idle fl ushing:** Occurs continuously, at a modest rate, when the cache utilization level is between the high and low watermark.

n **High watermark fl ushing:** Activated when cache utilization hits the high watermark. The storage system dedicates some additional resources for fl ushing. This type of fl ushing has some impact on I/O processing.

n **Forced fl ushing:** Occurs in the event of a large I/O burst when cache reaches 100 percent of its capacity, which signifi cantly affects the I/O response time. In forced fl ushing, system fl ushes the cache on priority

- by allocating more resources

b. How Traditional Storage Provisioning and Virtual Storage provisioning differ from each other ?

Traditional Storage Provisioning

In traditional storage provisioning, physical disks are logically grouped together and a required RAID level is applied to form a set, called a RAID set. The number of drives in the RAID set and the RAID level determine the availability, capacity, and performance of the RAID set. It is highly recommend that the RAID set be created from drives of the same type, speed, and capacity to ensure maximum usable capacity, reliability, and consistency in performance. For example, if drives of different capacities are mixed in a RAID set, the capacity of the smallest drive is used from each disk in the set to make up the RAID set's overall capacity. The remaining capacity of the larger drives remains unused. Likewise, mixing higher revolutions per minute (RPM) drives with lower RPM drives lowers the overall performance of the RAID set.

RAID sets usually have a large capacity because they combine the total capacity of individual drives in the set. *Logical units* are created from the RAID sets by partitioning (seen as slices of the RAID set) the available capacity into smaller units. These units are then assigned to the host based on their storage requirements. Logical units are spread across all the physical disks that belong to that set. Each logical unit created from the RAID set is assigned a unique ID, called a *logical unit number* (LUN). LUNs hide the organization and composition of the RAID set from the hosts. LUNs created by traditional storage provisioning methods are also referred to as *thick LUNs* to distinguish them from the LUNs

created by virtual provisioning methods.

Figure 4-5 shows a RAID set consisting of fi ve disks that have been sliced, or partitioned, into two LUNs: LUN 0 and LUN 1. These LUNs are then assigned to Host1 and Host 2 for their storage requirements.

When a LUN is confi gured and assigned to a non-virtualized host, a bus scan is required to identify the LUN. This LUN appears as a raw disk to the operating system. To make this disk usable, it is formatted with a fi le system and then the fi le system is mounted.

In a virtualized host environment, the LUN is assigned to the hypervisor, which recognizes it as a raw disk. This disk is confi gured with the hypervisor fi le system, and then virtual disks are created on it. *Virtual disks* are fi les on the hypervisor *Virtual provisioning* enables creating and presenting a LUN with more capacity than is physically allocated to it on the storage array. The LUN created using virtual provisioning is called a *thin LUN* to distinguish it from the traditional LUN.

Thin LUNs do not require physical storage to be completely allocated to them at the time they are created and presented to a host. Physical storage is allocated to the host "on-demand" from a shared pool of physical capacity.

A *shared pool* consists of physical disks. A shared pool in virtual provisioning is analogous to a RAID group, which is a collection of drives on which LUNs are created. Similar to a RAID group, a shared pool supports a single RAID protection level. However, unlike a RAID group, a shared pool might contain large numbers of drives. Shared pools can be homogeneous (containing a single drive type) or heterogeneous (containing mixed drive types, such as fl ash, FC, SAS, and SATA drives).

Virtual provisioning enables more effi cient allocation of storage to hosts. Virtual provisioning also enables oversubscription, where more capacity is presented to the hosts than is actually available on the storage array. Both shared pool and thin LUN can be expanded nondisruptively as the storage requirements of the hosts grow. Multiple shared pools can be created within a storage array, and a shared pool may be shared by multiple thin LUNs.

Comparison between Virtual and Traditional Storage Provisioning

Administrators typically allocate storage capacity based on anticipated storage requirements. This generally results in the over provisioning of storage capacity, which then leads to higher costs and lower capacity utilization. Administrators often over-provision storage to an application for various reasons, such as, to avoid frequent provisioning of storage if the LUN capacity is exhausted, and to reduce disruption to application availability. Over provisioning of storage often leads to additional storage acquisition and operational costs.

Virtual provisioning addresses these challenges. Virtual provisioning improves storage capacity utilization and simplifi es storage management. Figure 4-9 shows an example, comparing virtual provisioning with traditional storage provisioning.

c. Explain Fibre Channel Protocol Stack.

Fibre Channel Protocol Stack

It is easier to understand a communication protocol by viewing it as a structure of independent layers. FCP defi nes the communication protocol in fi ve 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 defi ned protocols. Figure 5-12 illustrates the Fibre Channel protocol stack.

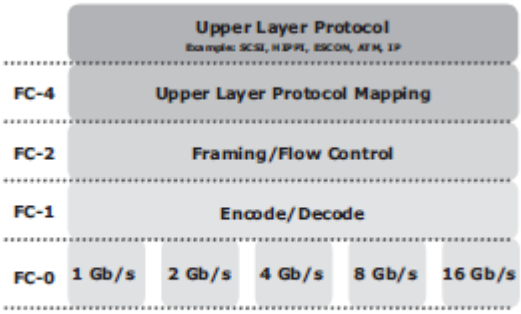

Figure 5-12: Fibre Channel protocol stack

FC-4 Layer

FC-4 is the uppermost layer in the FCP stack. This layer defi nes the application interfaces and the way *Upper Layer Protocols* (ULPs) are mapped to the lower FC layers. The FC standard defi nes several protocols that can operate on the FC-4 layer (see Figure 5-12). Some of the protocols include SCSI, High Performance Parallel Interface (HIPPI) Framing Protocol, Enterprise Storage Connectivity (ESCON), Asynchronous Transfer Mode (ATM), and IP.

FC-2 Layer

The FC-2 layer provides Fibre Channel addressing, structure, and organization of data (frames, sequences, and exchanges). It also defi nes fabric services, classes of service, fl ow control, and routing.

FC-1 Layer

The FC-1 layer defi nes how data is encoded prior to transmission and decoded upon receipt. 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 links with speeds of 10 Gbps and above use 64-bit to 66-bit encoding algorithms. The FC-1 layer also defi nes the transmission words, such as FC frame delimiters, which identify the start and end of a frame and primitive signals that indicate events at a transmitting port. In addition to these, the FC-1 layer performs link initialization and error recovery.

FC-0 Layer

FC-0 is the lowest layer in the FCP stack. This layer defi nes the physical interface, media, and transmission of bits. The FC-0 specifi cation includes cables, connectors, and optical and electrical parameters for a variety of data rates. The FC transmission can use both electrical and optical media.

4. Attempt *any one* of the following: **08**

a. Explain the components and benefits of NAS.

Components of NAS

A NAS device has two key components: NAS head and storage (see Figure 7-3). In some NAS implementations, the storage could be external to the NAS device and shared with other hosts. The NAS head includes the following components:
n CPU and memory
n One or more network interface cards (NICs), which provide connectivity to the client network. Examples of network protocols supported by NIC include Gigabit Ethernet, Fast Ethernet, ATM, and Fiber Distributed Data Interface (FDDI).
n An optimized operating system for managing the NAS functionality. It translates fi le-level requests into block-storage requests and further converts the data supplied at the block level to fi le data.
n NFS, CIFS, and other protocols for fi le sharing
n Industry-standard storage protocols and ports to connect and manage physical disk resources
The NAS environment includes clients accessing a NAS device over an IP network using fi le-sharing protocols.

b. Explain different Backup Granularity levels.

Backup Granularity

Backup granularity depends on business needs and the required RTO/RPO. Based on the granularity, backups can be categorized as full, incremental and cumulative (differential). Most organizations use a combination of these three backup types to meet their backup and recovery requirements. Figure 10-1 shows the different backup granularity levels.

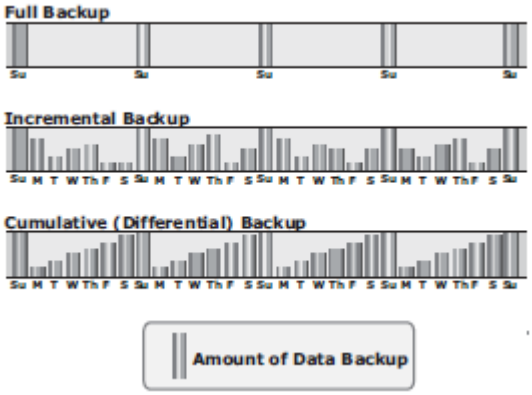

c. Write detail note on Content Addressed Storage .

Content-Addressed Storage

CAS is an object-based storage device designed for secure online storage and retrieval of fi xed content. CAS stores user data and its attributes as an object. The stored object is assigned a globally unique address, known as a *content address* (CA). This address is derived from the object's binary representation. CAS provides an optimized and centrally managed storage solution. Data access in CAS differs from other OSD devices. In CAS, the application server access the CAS device only via the CAS API running on the application server. However, the way CAS stores data is similar to the other OSD systems.
CAS provides all the features required for storing fi xed content. The key features of CAS are as follows:

n **Content authenticity:** It assures the genuineness of stored content. This is achieved by generating a unique content address for each object and validating the content address for stored objects at regular intervals. Content authenticity is assured because the address assigned to each object is as unique as a fi ngerprint. Every time an object is read, CAS uses a hashing algorithm to recalculate the object's content address as a validation step and compares the result to its original content address. If the object fails validation, CAS rebuilds the object using a mirror or parity protection scheme.

n **Content integrity:** It provides assurance that the stored content has not been altered. CAS uses a hashing algorithm for content authenticity and integrity. If the fi xed content is altered, CAS generates a new address for the altered content, rather than overwrite the original fi xed content.

n **Location independence:** CAS uses a unique content address, rather than directory path names or URLs, to retrieve data. This makes the physical location of the stored data irrelevant to the application that requests the data.

n **Single-instance storage (SIS):** CAS uses a unique content address to guarantee the storage of only a single instance of an object. When a new object is written, the CAS system is polled to see whether an object is already available with the same content address. If the object is available in the system, it is not stored; instead, only a pointer to that object is created.

n **Retention enforcement:** Protecting and retaining objects is a core requirement of an archive storage system. After an object is stored in the CAS system and the retention policy is defi ned, CAS does not make the object available for deletion until the policy expires.

n **Data protection:** CAS ensures that the content stored on the CAS system is available even if a disk or a node fails. CAS provides both local and remote protection to the data objects stored on it. In the local protection option, data objects are either mirrored or parity protected. In mirror protection, two copies of the data object are stored on two different nodes in the same cluster. This decreases the total available capacity by 50 percent. In parity protection, the data object is split in multiple parts and parity is generated from them. Each part of the data and its parity are stored on a different node. This method consumes less capacity to protect the stored data, but takes slightly longer to regenerate the data if corruption of data occurs. In the remote replication option, data objects are copied to a secondary CAS at the remote location. In this case, the objects remain accessible from the secondary CAS if the primary CAS system fails.

n **Fast record retrieval:** CAS stores all objects on disks, which provides faster access to the objects compared to tapes and optical discs.

n **Load balancing:** CAS distributes objects across multiple nodes to provide maximum throughput and availability.

n **Scalability:** CAS allows the addition of more nodes to the cluster without any interruption to data access and with minimum administrative overhead.

n **Event notifi cation:** CAS continuously monitors the state of the system

and raises an alert for any event that requires the administrator's attention. The event notifi cation is communicated to the administrator through SNMP, SMTP, or e-mail.

n **Self diagnosis and repair:** CAS automatically detects and repairs corrupted objects and alerts the administrator about the potential problem. CAS systems can be confi gured to alert remote support teams who can diagnose and repair the system remotely.

n **Audit trails:** CAS keeps track of management activities and any access or disposition of data. Audit trails are mandated by compliance requirements.

5. Attempt *any one* of the following: **08**

a. What are the uses of local replications?

Uses of Local Replicas

One or more local replicas of the source data may be created for various purposes, including the following:

n **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.

n **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.

n **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.

n **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.

n **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.

b. What are Remote Replication Technologies?

Remote Replication Technologies

Remote replication of data can be handled by the hosts or storage arrays. Other options include specialized network-based appliances to replicate data over the LAN or SAN. An advanced replication option such as three-site replication is discussed in section "12.3 Three-Site Replication."

Host-Based Remote Replication

Host-based remote replication uses the host resources to perform and manage the replication operation. There are two basic approaches to host-based remote replication: Logical volume manager (LVM) based replication and database replication via log shipping.

LVM-Based Remote Replication

LVM-based remote replication is performed and managed at the volume group level. Writes to the source volumes are transmitted to the remote host by the LVM. The LVM on the remote host receives the writes and commits them to the remote volume group.

Host-Based Log Shipping

Database replication via log shipping is a host-based replication technology supported by most databases. Transactions to the source database are captured in logs, which are periodically transmitted by the source host to the remote host (see Figure 12-6). The remote host receives the logs and applies them to the remote database.

c. Explain Three-Site Replication.

Three-Site Replication

In synchronous replication, the source and target sites are usually within a short distance. Therefore, if a regional disaster occurs, both the source and the target sites might become unavailable. This can lead to extended RPO and RTO because the last known good copy of data would need to come from another source, such as an offsite tape library.

A regional disaster will not affect the target site in asynchronous replication because the sites are typically several hundred or several thousand kilometers apart. If the source site fails, production can be shifted to the target site, but there is no further remote protection of data until the failure is resolved.

Three-site replication mitigates the risks identifi ed in two-site replication. In a three-site replication, data from the source site is replicated to two remote sites. Replication can be synchronous to one of the two sites, providing a near zero-RPO solution, and it can be asynchronous or disk buffered to the other remote site, providing a fi nite RPO. Three-site remote replication can be implemented as a cascade/multihop or a triangle/multitarget solution

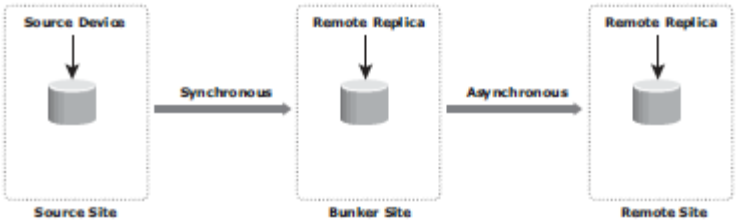

(a) Synchronous + Asynchronous

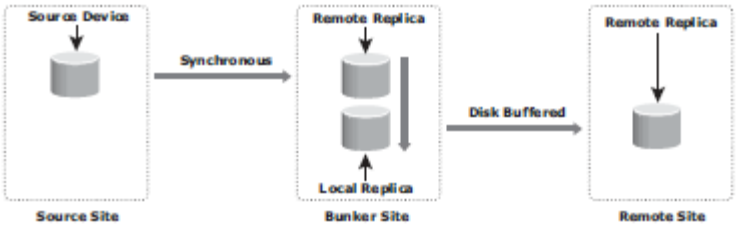

6.	**Attempt *any one* of the following:**	**08**
d.	Explain Information Security Framework. **Information Security Framework** The basic information security framework is built to achieve four security goals: confi dentiality, integrity, and availability (CIA), along with accountability. This framework incorporates all security standards, procedures, and controls, required to mitigate threats in the storage infrastructure environment. n **Confi dentiality:** Provides the required secrecy of information and ensures that only authorized users have access to data. This requires authentication of users who need to access information. Data in transit (data transmitted over cables) and data at rest (data residing on a primary storage, backup media, or in the archives) can be encrypted to maintain its confi dentiality. In addition to restricting unauthorized users from accessing information, confi dentiality also requires implementing traffi c fl ow protection measures as part of the security protocol. These protection measures generally include hiding source and destination addresses, frequency of data being sent, and amount of data sent. n **Integrity:** Ensures that the information is unaltered. Ensuring integrity requires detection of and protection against unauthorized alteration or deletion of information. Ensuring integrity stipulates measures such as error detection and correction for both data and systems. n **Availability:** This ensures that authorized users have reliable and timely access to systems, data, and applications residing on these systems. Availability requires protection against unauthorized deletion of data and denial of service (discussed in section “14.2.2 Threats”). Availability also implies that suffi cient resources are available to provide a service. n **Accountability service:** Refers to accounting for all the events and operations	

	that take place in the data center infrastructure. The accountability service maintains a log of events that can be audited or traced later for the purpose of security.	
e.	What are the basic SAN security mechanisms? ***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 are detailed in Chapter 4 and Chapter 5, respectively. The standard implementations of LUN masking on storage arrays mask the LUNs presented to a frontend 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 FC addresses. It offers a mechanism to lock down the FC address of a given node port to its WWN. *WWPN zoning* is the preferred choice in security-conscious environments. ***Securing Switch Ports*** 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 spoofi ng. *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 confi gured 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 specifi es which HBAs and storage ports can be a part of the fabric, preventing unauthorized devices from accessing it. Similarly, the switch connection control policy specifi es which switches are allowed to be part of the fabric, preventing unauthorized switches 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 any attempt to connect any switch in the fabric by using an ISL causes the fabric to segment. Role-based access control provides additional security to a SAN by preventing unauthorized 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 *zone admin* role can modify the zones on the fabric, whereas a basic user may view only fabricrelated 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 14-6 depicts logical partitioning in a VSAN. The SAN administrator can create distinct VSANs by populating each of them with switch ports. In the example, the switch ports are distributed over two VSANs: 10 and 20 — for the Engineering and HR divisions, respectively. Although they share physical switching gear with other divisions, they can be managed individually as standalone fabrics. 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. VSANs minimize the impact of fabricwide disruptive events because management and control traffi c 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 authorization control within a single fabri	
f.	Explain the different aspects of NAS Security.	