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

Internal Assessment Test 1 – May 2022

Sub:	SAN				Sub Code:	18CS822	Branch:	cse
Date:		Duration:	90 mins	Max Marks:	50	Sem/Sec:		OBE

Answer any FIVE FULL Questions

		MARKS	CO	RBT
1 (a)	Explain Evolution of Storage Architecture with neat diagram	[06]	L1	CO1

In conventional IT architectures, storage devices are normally only connected to a single server (Figure 1.1). To increase fault tolerance, storage devices are sometimes connected to two servers, with only one server actually able to use the storage device at any one time. In both cases, the storage device exists only in relation to the server to which it is connected. Other servers cannot directly access the data; they always have to go through the server that is

connected to the storage device. This conventional IT architecture is therefore called server-centric IT architecture. In this approach, servers and storage devices are generally connected

together by SCSI cables.

As mentioned above, in conventional server-centric IT architecture storage devices exist only in relation to the one or two servers to which they are connected. The failure of both of these computers would make it impossible to access this data. Most companies find this unacceptable: at least some of the company data (for example, patient files, websites) must be available around the clock.

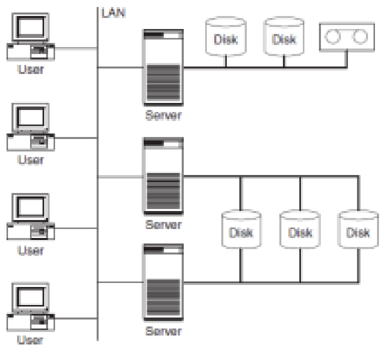

Figure 1.1 In a server-centric IT architecture storage devices exist only in relation to servers.

To overcome these challenges, storage evolved from server-centric to information-centric architecture (see Figure 1-4 [b]). In this architecture, storage devices

are managed centrally and independent of servers. These centrally-managed storage devices are shared with multiple servers. When a new server is deployed in the environment, storage is assigned from the same shared storage devices to that server. The capacity of shared storage can be increased dynamically by adding more storage devices without impacting information availability. In this architecture, information management is easier and cost-effective.

<table>
<tr>
<td></td>
<td>

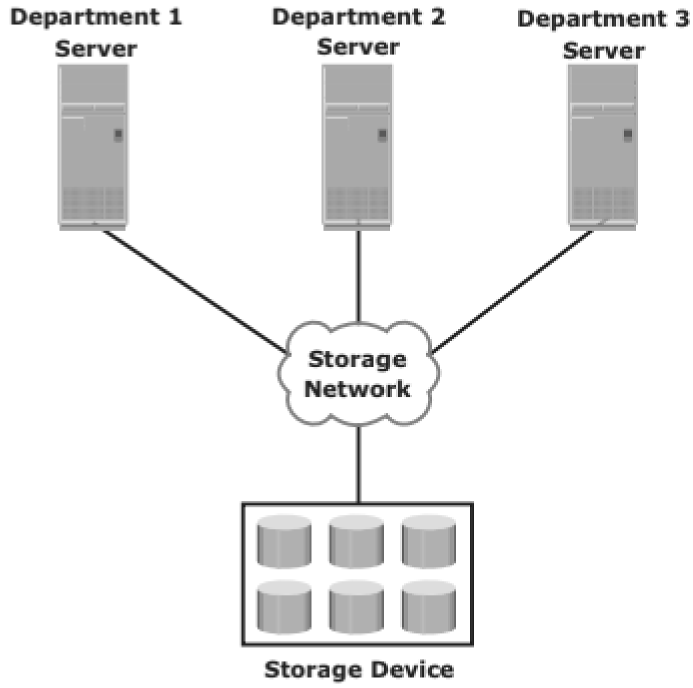

(b) Information-Centric Storage Architecture

</td>
<td></td>
<td></td>
<td></td>
</tr>
<tr>
<td>(b)</td>
<td>

Discuss Memory Virtualization -Diagram for Disk partitioning and concatenation\

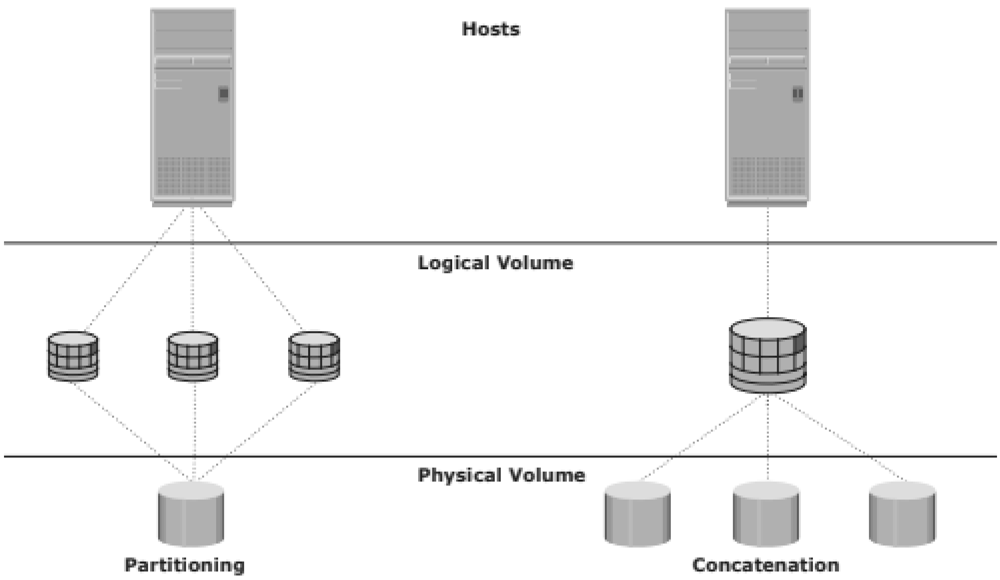

Figure 2-1: Disk partitioning and concatenation

Disk partitioning was introduced to improve the fl exibility and utilization of disk drives. In partitioning, a disk drive is divided into logical containers called logical volumes (LVs) (see Figure 2-1). For example, a large physical drive can be partitioned into multiple LVs to maintain data according to the fi le system and

application requirements. The partitions are created from groups of contigu- ous cylinders when the hard disk is initially set up on the host. The host's fi le

</td>
<td>[04]</td>
<td>L1</td>
<td>CO1</td>
</tr>
</table>

	system accesses the logical volumes without any knowledge of partitioning and physical structure of the disk. Concatenation is the process of grouping several physical drives and presenting them to the host as one big logical volume (see Figure 2-1). The LVM provides optimized storage access and simplifies storage resource management. It hides details about the physical disk and the location of data on the disk. It enables administrators to change the storage allocation even when The application is running.			
2 (a)	What is Protocol? Explain the popular interface Protocols used for host to storage communication. A protocol enables communication between the host and storage. Protocols are implemented using interface devices (or controllers) at both source and destination. IDE/ATA and Serial ATA IDE/ATA is a popular interface protocol standard used for connecting storage devices, such as disk drives and CD-ROM drives. This protocol supports parallel transmission and therefore is also known as Parallel ATA (PATA) or simply ATA. IDE/ATA has a variety of standards and names. The Ultra DMA/133 version of ATA supports a throughput of 133 MB per second. In a master-slave configuration, an ATA interface supports two storage devices per connector. However, if the performance of the drive is important, sharing a port between two devices is not recommended. SCSI and Serial SCSI SCSI has emerged as a preferred connectivity protocol in high-end computers. This protocol supports parallel transmission and offers improved performance, scalability, and compatibility compared to ATA. However, the high cost associated with SCSI limits its popularity among home or personal desktop users. Internet Protocol (IP) IP is a network protocol that has been traditionally used for host-to-host traffic. With the emergence of new technologies, an IP network has become a viable option for host-to-storage communication. IP offers several advantages in terms of cost and maturity and enables organizations to leverage their existing IP-based network. iSCSI and FCIP protocols are common examples that leverage IP for host-to-storage communication. Fibre Channel Fibre Channel is a widely used protocol for high-speed communication to the storage device. The Fibre Channel interface provides gigabit network speed. It provides a serial data transmission that operates over copper wire and optical fiber. The latest version of the FC interface (16FC) allows transmission of data up to 16 Gb/s.	[05]	L1	CO1
(b)	Discuss Volume Manager and Computer Virtualization in detail Volume Managers (LVMs) enabled dynamic extension of file system capacity and efficient storage management. The LVM is software that runs on the computer system and manages logical and physical storage. Compute virtualization is a technique for masking or abstracting the physical hardware from the operating system. It enables multiple operating systems to run concurrently on single or clustered physical machines. This technique enables creating portable virtual compute systems called virtual machines (VMs). Each VM runs an operating system and application instance in an isolated manner.	[05]	L1	CO1

	Compute virtualization is achieved by a virtualization layer that resides between the hardware and virtual machines. This layer is also called the hypervisor. The hypervisor provides hardware resources, such as CPU, memory, and network to all virtual machines. Within a physical server, a large number of virtual machines can be created depending on the hardware capabilities of the physical server.			
3 (a)	What is a Data Center? Discuss key Characteristics of Data Center with neat Diagram A data center is a facility that contains information storage and other physical information technology (IT) resources for computing, networking, and storing information. 1.Availability: A data center should ensure the availability of information when required. Unavailability of information could cost millions of dollars per hour to businesses, such as financial services, telecommunications, and e-commerce. 2.Security: Data centers must establish policies, procedures, and core element integration to prevent unauthorized access to information. 3.Scalability: Business growth often requires deploying more servers, new applications, and additional databases. Data center resources should scale based on requirements, without interrupting business operations. 4.Performance: All the elements of the data center should provide optimal performance based on the required service levels. 5 Data integrity: Data integrity refers to mechanisms, such as error correction codes or parity bits, which ensure that data is stored and retrieved exactly as it was received. 6 Capacity: Data center operations require adequate resources to store and process large amounts of data, efficiently. When capacity requirements increase, the data center must provide additional capacity without interrupting availability or with minimal disruption. Capacity may be managed by reallocating the existing resources or by adding new resources. 7 Manageability: A data center should provide easy and integrated management of all its elements. Manageability can be achieved through automation and reduction of human (manual) intervention in common tasks. 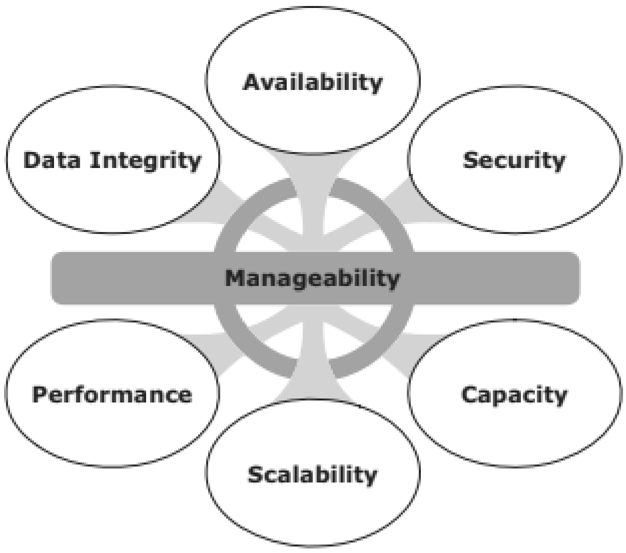 **Figure 1-6:** Key characteristics of a data center	[06]	L1	CO1
(b)	Explain key elements of Data center Core Elements of a Data Center	[04]	L1	CO1

	Five core elements are essential for the functionality of a data center: 1.Application: A computer program that provides the logic for computing operations 2.Database management system (DBMS): Provides a structured way to store data in logically organized tables that are interrelated 3.Host or compute: A computing platform (hardware, firmware, and software) that runs applications and databases 4.Network: A data path that facilitates communication among various networked devices 5. Storage: A device that stores data persistently for subsequent use			
4 (a)	**Explain RAID techniques — striping, mirroring, and parity with neat diagrams** **Striping:** Striping is a technique to spread data across multiple drives (more than one) to use the drives in parallel. All the read-write heads work simultaneously, allowing more data to be processed in a shorter time and increasing performance, com-compared to reading and writing from a single disk. Within each disk in a RAID set, a predefined number of contiguously addressable disk blocks are defined as a strip. The set of aligned strips that spans across all the disks within the RAID set are called a stripe. Figure 3-2 shows physical and logical representations of a striped RAID set. Strip size (also called stripe depth) describes the number of blocks in a strip and is the maximum amount of data that can be written to or read from a single disk in the set, assuming that the accessed data starts at the beginning of the strip. All strips in a stripe have the same number of blocks. Having a smaller strip size means that data is broken into smaller pieces while spread across the disks. **Mirroring:** Mirroring is a technique whereby the same data is stored on two different disk drives, yielding two copies of the data. If one disk drive failure occurs, the data is intact on the surviving disk drive (see Figure 3-3) and the controller continues to service the host's data requests from the surviving disk of a mirrored pair. When the failed disk is replaced with a new disk, the controller copies the data from the surviving disk of the mirrored pair. This activity is transparent to the host. In addition to providing complete data redundancy, mirroring enables fast recovery from disk failure. However, disk mirroring provides only data pro-tection and is not a substitute for data backup. Mirroring constantly captures changes in the data, whereas a backup captures point-in-time images of the data.	[05]	L1	CO1

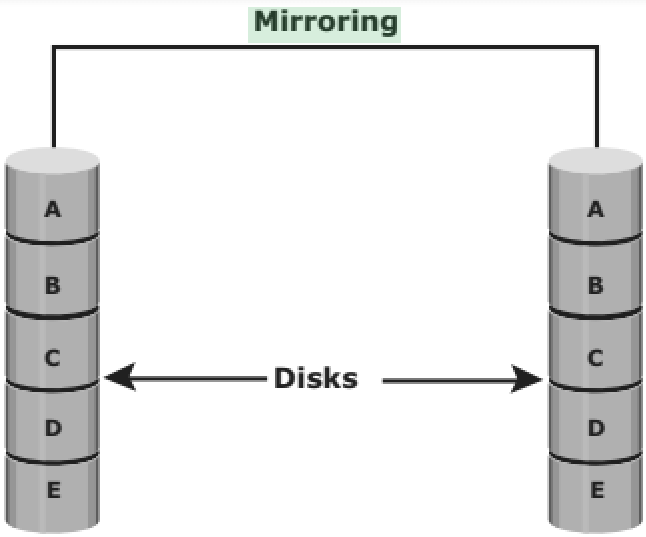

Figure 3-3: Mirrored disks in an array

Parity

Parity is a method to protect striped data from disk drive failure without the cost of mirroring. An additional disk drive is added to hold parity, a mathematical construct that allows re-creation of the missing data. Parity is a redundancy technique that ensures protection of data without maintaining a full set of duplicate data. Calculation of parity is a function of the RAID controller.

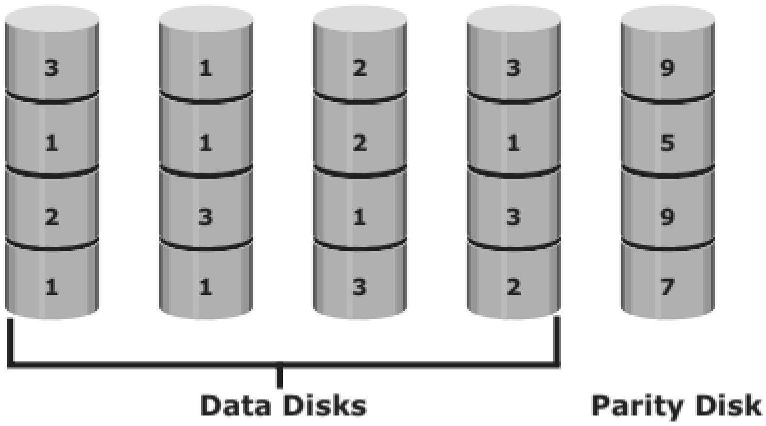

Figure 3-4: Parity RAID

	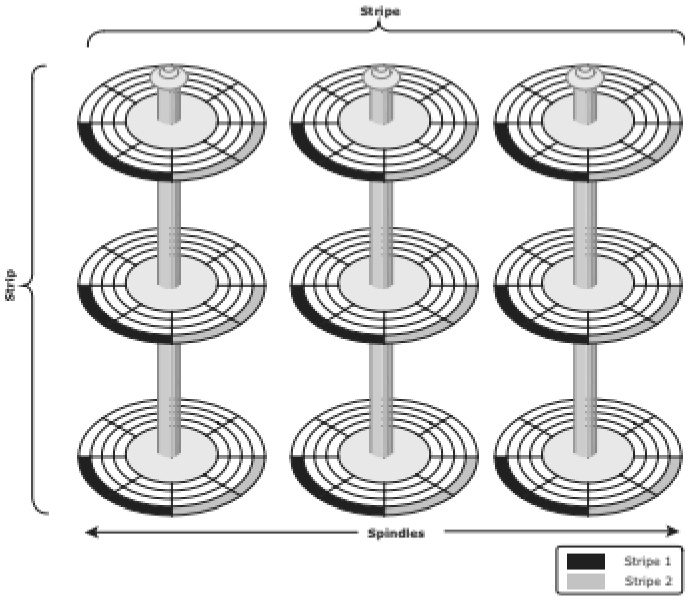 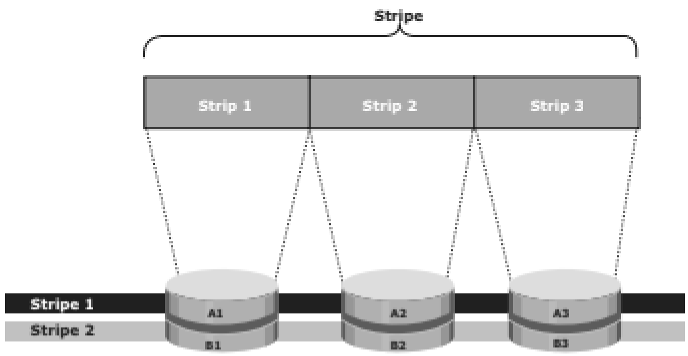 e 3-2: Striped RAID set			
(b)	Explain the RAID levels RAID 0, RAID 1 with neat diagrams **RAID 0** RAID 0 configuration uses data striping techniques, where data is striped across all the disks within a RAID set. Therefore it utilizes the full storage capacity of a RAID set. To read data, all the strips are put back together by the controller. Figure 3-5 shows RAID 0 in an array in which data is striped across five disks. When the number of drives in the RAID set increases, performance improves	[05]	L1	CO1

	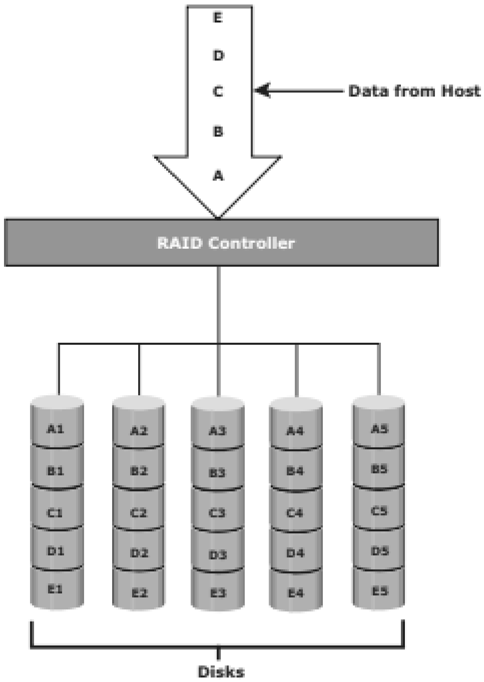 **Figure 3-5:** RAID 0 **RAID 1:** RAID 1 is based on the mirroring technique. In this RAID configuration, data is mirrored to provide fault tolerance (see Figure 3-6). A RAID 1 set consists of two disk drives and every write is written to both disks. The mirroring is transparent to the host. During disk failure, the impact on data recovery in RAID 1 is the least among all RAID implementations. This is because the RAID controller uses the mirror drive for data recovery. RAID 1 is suitable for applications that require high availability and cost is no constraint.			

<table>
<tr><td></td><td>

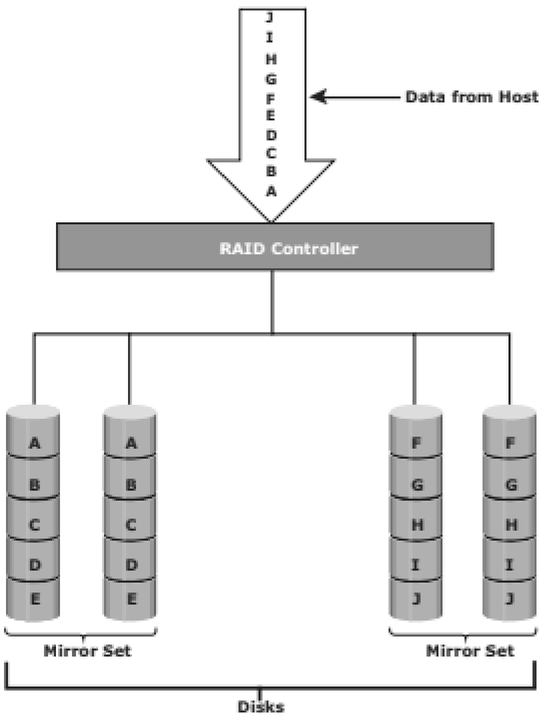

Figure 3-6: RAID 1

</td><td></td><td></td><td></td></tr>
<tr><td>5 (a)</td><td>

Explain with a diagram - Host Access to Data Internal and external DAS architecture with DAS Benefits and Limitations

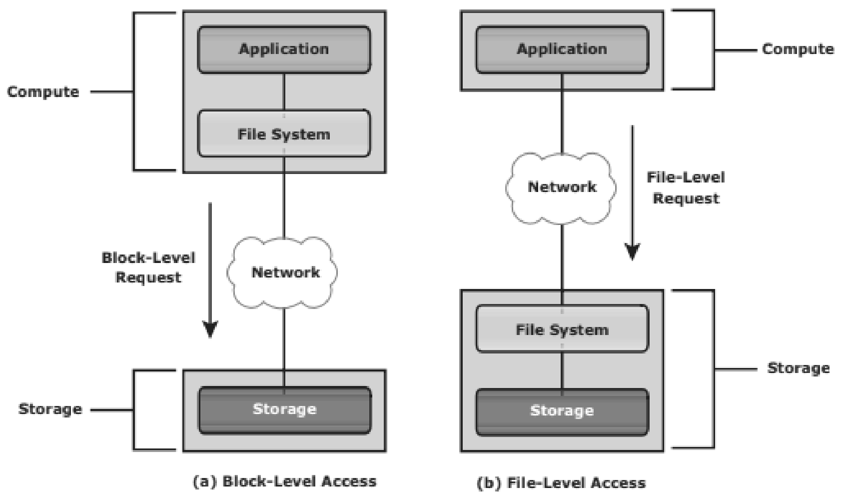

Figure 2-14: Host access to storage

the file system is created on a host, and data is accessed
on a network at the block level, as shown in Figure 2-14 (a). In this case, raw
disks or logical volumes are assigned to the host for creating the file system.
In internal DAS architectures, the storage device is internally connected to
the host by a serial or parallel bus (see Figure 2-15 [a]). The physical bus has

distance limitations and can be sustained only over a shorter distance for high-
speed connectivity. In addition, most internal buses can support only a limited

</td><td>[6]</td><td>L1</td><td>CO1</td></tr>
</table>

	number of devices, and they occupy a large amount of space inside the host, making maintenance of other components difficult. On the other hand, in external DAS architectures, the host connects directly to the external storage device, and data is accessed at the block level (see Figure 2-15 [b]). In most cases, communication between the host and the storage device takes place over a SCSI or FC protocol. Compared to internal DAS, an external DAS overcomes the distance and device count limitations and provides centralized management of storage devices. **DAS Benefits and Limitations** DAS requires a relatively lower initial investment than storage networking architectures. The DAS configuration is simple and can be deployed easily and rapidly. The setup is managed using host-based tools, such as the host OS, which makes storage management tasks easy for small environments. Because DAS has a simple architecture, it requires fewer management tasks and less hardware and software elements to set up and operate. However, DAS does not scale well. A storage array has a limited number of ports, which restricts the number of hosts that can directly connect to the storage. When capacities are reached, the service availability may be compromised. DAS does not make optimal use of resources due to its limited capability to share front-end ports. In DAS environments, unused resources cannot be easily re-allocated, resulting in islands of over-utilized and under-utilized storage pools.			
(b)	Discuss Utilization versus response time with a graph 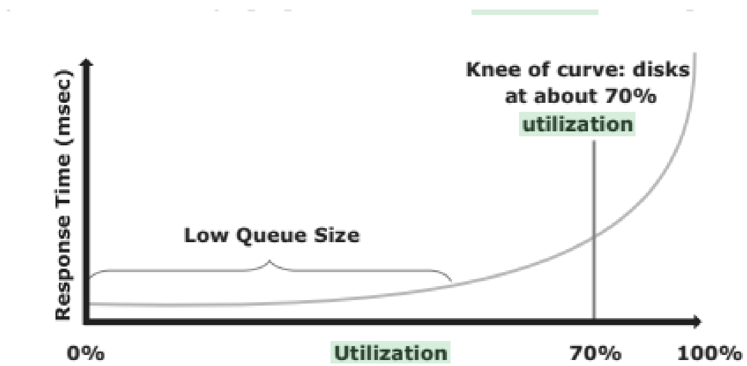 **Figure 2-13:** Utilization versus response time Utilization of a disk I/O controller has a significant impact on the I/O response time. To understand this impact, consider that a disk can be viewed as a black box consisting of two elements: n Queue: The location where an I/O request waits before it is processed by the I/O controller n Disk I/O Controller: Processes I/Os waiting in the queue one by one The I/O requests arrive at the controller at the rate generated by the application. This rate is also called the arrival rate. These requests are held in the I/O queue, and the I/O controller processes them one by one, as shown in Figure 2-12. The I/O arrival rate, the queue length, and the time taken by the I/O controller to process each request determines the I/O response time. If the controller is busy or heavily utilized, the queue size will be large and the response time will be high.	[2+2]	L1	CO1

6 (a)	Explain with neat diagram - Structure of Cache with its operations 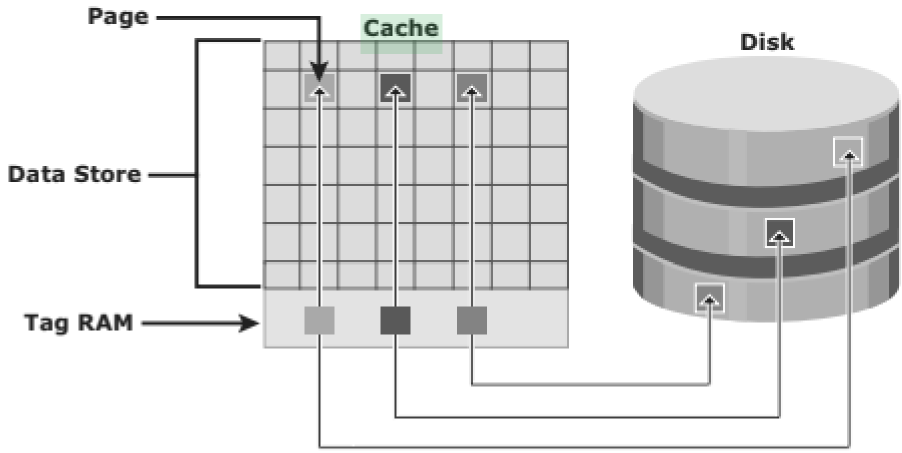 **Figure 4-2:** Structure of cache Read Operation with Cache When a host issues a read request, the storage controller reads the tag RAM to determine whether the required data is available in cache. If the requested data is found in the cache, it is called a read cache hit or read hit and data is sent directly to the host, without any disk operation A prefetch or read-ahead algorithm is used when read requests are sequential. In a sequential read request, a contiguous set of associated blocks is retrieved. Several other blocks that have not yet been requested by the host can be read from the disk and placed into cache in advance. When the host subsequently requests these blocks, the read operations will be read hits. This process signifi cantly improves the response time experienced by the host. The intelligent storage system offers fi xed and variable prefetch sizes. In fi xed prefetch, the intelligent storage system prefetches a fi xed amount of data. It is most suitable when host I/O sizes are uniform. In variable prefetch, the storage system prefetches an amount of data in multiples of the size of the host request. Maximum prefetch limits the number of data blocks that can be prefetched to prevent the disks from being rendered busy with prefetch at the expense of other I/Os.	[05]	L1	CO1
(b)	Explain with neat diagram - Components of an intelligent storage system An intelligent storage system consists of four key components: front end, cache, back end, and physical disks. Figure 4-1 illustrates these components and their interconnections. An I/O request received from the host at the front-end port is processed through cache and back end, to enable storage and retrieval of data from the physical disk. A read request can be serviced directly from cache if the requested data is found in the cache. In modern intelligent storage systems, front end, cache, and back end are typically integrated on a single board (referred to as a storage processor or storage controller).	[05]	L1	CO1

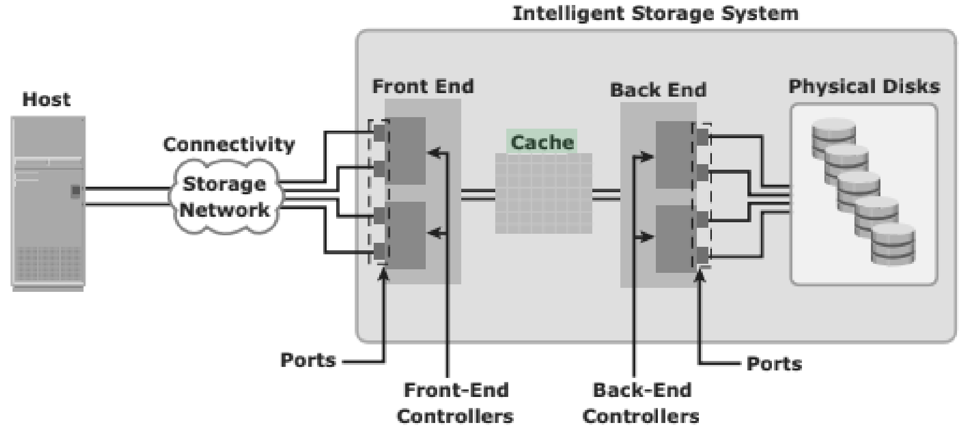

Figure 4-1: Components of an intelligent storage system

CO PO Mapping

Course Outcomes		Modules covered	PO1	PO2	PO3	PO4	PO5	PO6	PO7	PO8	PO9	PO10	PO11	PO12	PSO1	PSO2	PSO3	PSO4
CO1	Identify key challenges in managing information and analyze different storage networking technologies and virtualization	**1**	2	2	1	1	-	-	-	-	-	-	-	-	-	1	1	-
CO2	Explain components and the implementation of NAS	**2**	2	2	1	1	-	-	-	-	-	-	-	-	-	1	1	-

CO3	Describe CAS architecture and types of archives and forms of virtualization	**2,3**	2	2	1	1	-	-	-	-	-	-	-	-	-	1	1	-
CO4	Illustrate the storage infrastructure and management activities	**5**	2	2	1	1	-	-	-	-	-	-	-	-	-	2	1	-
CO5	Identify key challenges in managing information and analyze different storage networking technologies and virtualization	**1**	2	2	1	1	-	-	-	-	-	-	-	-	-	1	1	-

COGNITIVE LEVEL	REVISED BLOOMS TAXONOMY 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.

PROGRAM OUTCOMES (PO), PROGRAM SPECIFIC OUTCOMES (PSO)				CORRELATION LEVELS	
PO1	Engineering knowledge	PO7	Environment and sustainability	0	No Correlation
PO2	Problem analysis	PO8	Ethics	1	Slight/Low
PO3	Design/development of solutions	PO9	Individual and team work	2	Moderate/ Medium
PO4	Conduct investigations of complex problems	PO10	Communication	3	Substantial/ High
PO5	Modern tool usage	PO11	Project management and finance		
PO6	The Engineer and society	PO12	Life-long learning		
PSO1	Develop applications using different stacks of web and programming technologies				
PSO2	Design and develop secure, parallel, distributed, networked, and digital systems				
PSO3	Apply software engineering methods to design, develop, test and manage software systems.				
PSO4	Develop intelligent applications for business and industry				