CMR INSTITUTE OF TECHNOLOGY

USN | | | | | | | | | |

Internal Assessment Test 1 – Sep. 2020

Sub:	Storage Area Networks							Sub Code:	18MCA554
Date:	18/09/2020	Duration:	90 min's	Max Marks:	50	Sem & Sec:	5 A & B	Branch:	MCA

Note : Answer any full FIVE Questions

	PART I	MARKS	OBE CO	OBE RBT
1	What do you understand by data center? Name and explain core elements of data center. What are characteristics of data center? **Sol:** **Data Center** : It is a facility that contains storage , compute network, and other IT resources to provide centralized data – processing capabilities. **CORE ELEMENTS OF A DATA CENTER** Five core elements are essential for the functionality of a data center: ➢ **Application**: A computer program that provides the logic for computing operations. ➢ **Database management system (DBMS):** Provides a structured way to store data in logically organized tables that are interrelated. ➢ **Host or compute:** A computing platform (hardware, firmware, and software) that runs applications and databases. ➢ **Network:** A data path that facilitates communication among various networked devices. ➢ **Storage:** A device that stores data persistently for subsequent use ➢ The figure below shows an example of an online order transaction system that involves the five core elements of a data center and illustrates their functionality in a business process. **KEY CHARACTERISTICS OF A DATA CENTER** ➢ Uninterrupted operation of data centers is critical to the survival and success of a business. ➢ Organizations must have a reliable infrastructure that ensures that data is accessible at all times. ➢ The following are the key characteristics of Data Center ➢ **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.	[10]	CO1	L2

	➢ **Security**: Data centers must establish policies, procedures, and core element integration to prevent unauthorized access to information. ➢ **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. ➢ **Performance**: All the elements of the data center should provide optimal performance based on the required service levels. ➢ **Data integrity**: Data integrity refers to mechanisms, such as error correction codes or parity bits, which ensures that data is stored and retrieved exactly as it was received. ➢ **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. ➢ **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.  **Fig : K**ey characteristics of a Data Center			
2	Explain different RAID techniques by drawing neat diagram. Sol : **RAID TECHNIQUES** Different RAID techniques are 1. **Striping** 2. **Mirroring** **1. Parity 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, 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**.	[10]	CO1	L1

- The set of aligned strips that spans across all the disks within the RAID set is called a **stripe**.
- The below figure shows physical and logical representations of a striped RAID set.

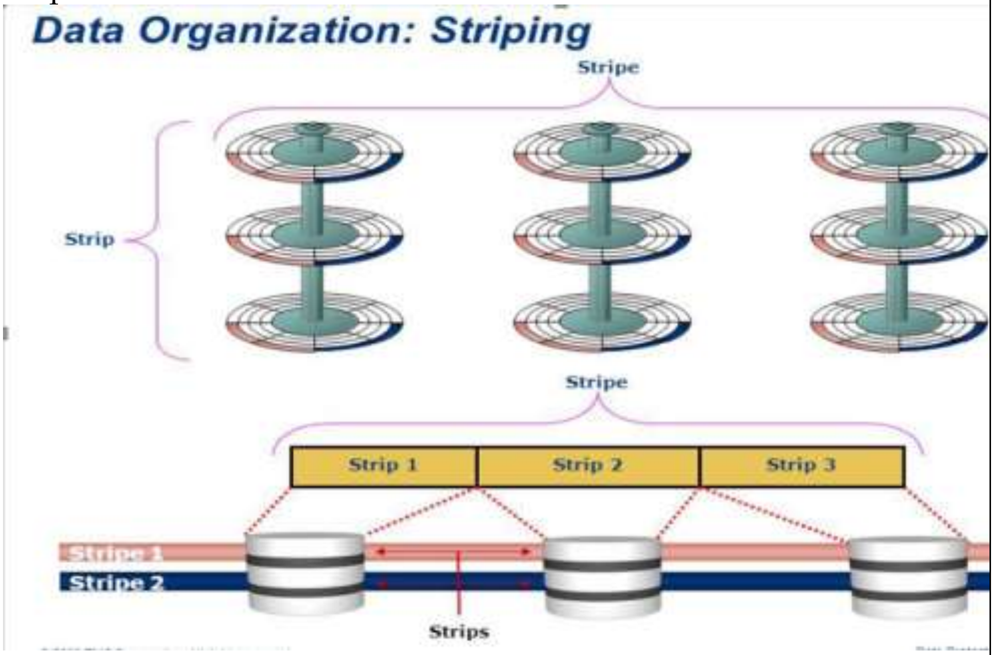

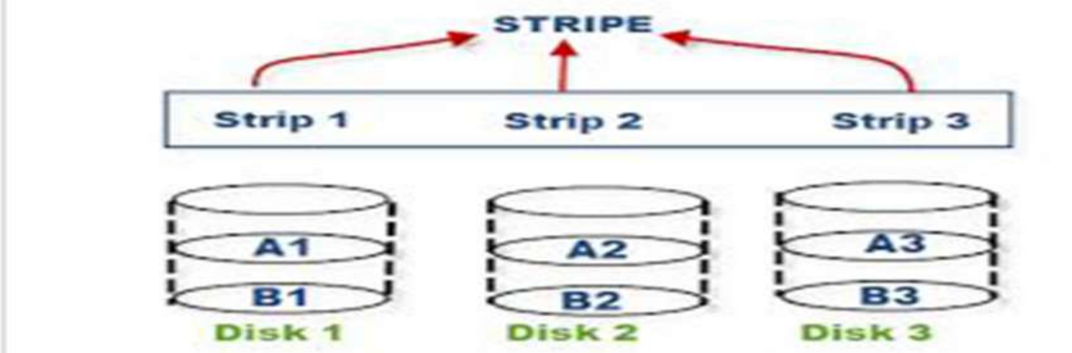

- Strip size (also called stripe depth) describes the number of blocks in a strip and is the maximum amount of data that can be written to or read from a single disk in the set, assuming that the accessed data starts at the beginning of the strip.
- All strips in a stripe have the same number of blocks.
- Having a smaller strip size means that data is broken into smaller pieces while spread across the disks.
- Stripe size is a multiple of strip size by the number of data disks in the RAID set. For example, in a five disk striped RAID set with a strip size of 64 KB, the stripe size is 320 KB(64KB X 5). Stripe width refers to the number of data strips in a stripe.
- Striped RAID does not provide any data protection unless parity or mirroring is used.

2. 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 as shown in the figure below and the controller continues to service the host's data requests from the surviving disk of a mirrored pair.

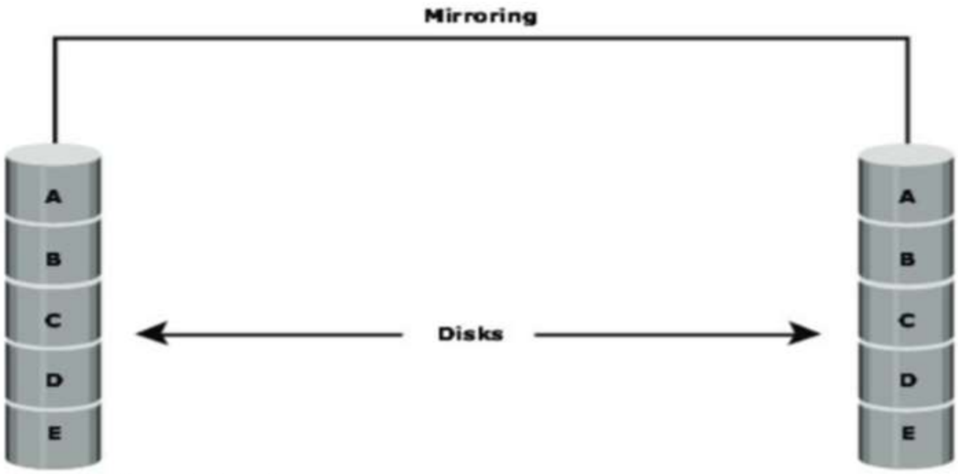

Figure 3-3: Mirrored disks in an array

- 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 protection 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.

- Mirroring involves duplication of data — the amount of storage capacity needed is twice the amount of data being stored. Therefore, mirroring is considered expensive and is preferred for mission-critical applications that cannot afford the risk of any data loss.
- Mirroring improves read performance because read requests can be serviced by both disks.
- However, write performance is slightly lower than that in a single disk because each write request manifests as two writes on the disk drives.
- Mirroring does not deliver the same levels of write performance as a striped RAID.

3.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 recreation 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. Parity information can be stored on separate, dedicated disk drives or distributed across all the drives in a RAID set.
- Figure below shows a parity RAID set.
- The first three disks, labeled “Data Disks,” contain the data.
- The fourth disk, labeled “Parity Disk,” stores the parity information, which, in this case, is the sum of the elements in each row.
- Now, if one of the data disks fails, the missing value can be calculated by subtracting the sum of the rest of the elements from the parity value.

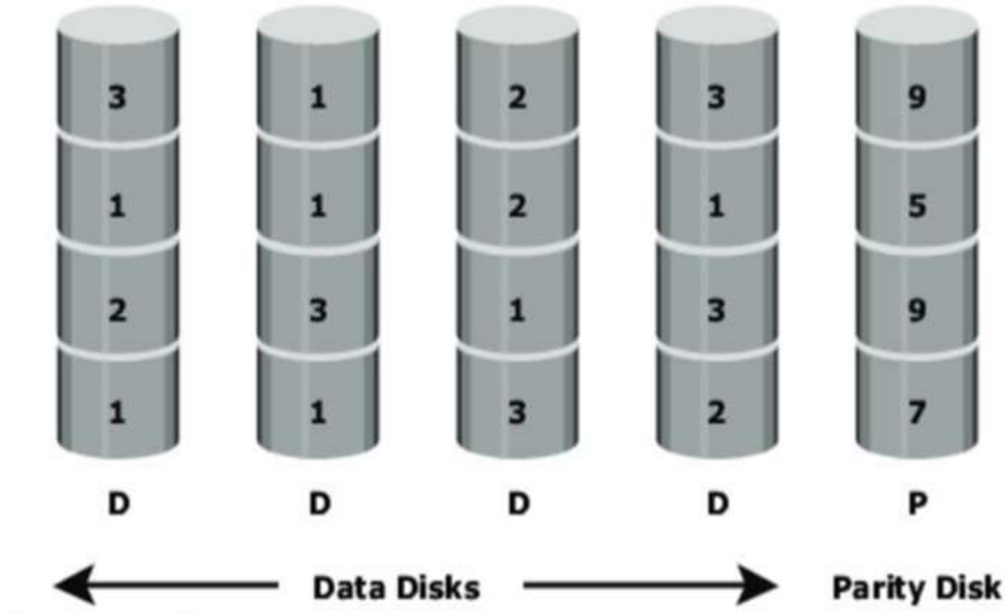

Figure 3-4: Parity RAID

Disadvantages of using parity:

- Parity information is generated from data on the data disk. Therefore, parity is recalculated every time there is a change in data. This recalculation is time-consuming and affects the performance of the RAID array.
- For parity RAID, the stripe size calculation does not include the parity strip.
- For example in a five (4 + 1) disk parity RAID set with a strip size of 64 KB, the stripe size will be 256 KB (64 KB ¥ 4).

3	Compare the following: (a) RAID 1+0 and RAID 0+1 **Sol :** Most data centers require data redundancy and performance from their RAID arrays. **RAID 0+1** and **RAID 1+0** combine the **performance benefits of RAID 0** with the **redundancy benefits of RAID 1**. They use striping and mirroring techniques and combine their benefits. These types of RAID require an even number of disks, the minimum being four (see Figure 3-7). RAID 1+0 is also known as RAID 10 (Ten) or RAID 1/0. Similarly, RAID 0+1 is also known as RAID 01 or RAID 0/1.	[5]	CO1	L1

RAID 1+0 performs well for workloads that use small, random, write-intensive I/O. Some applications that benefit from RAID 1+0 include the following:

1. High transaction rate Online Transaction Processing (OLTP)
2. Large messaging installations
3. Database applications that require high I/O rate, random access, and high availability.

A common misconception is that RAID 1+0 and RAID 0+1 are the same. **RAID 1+0 is also called *striped mirror*.** The basic element of RAID 1+0 is a mirrored pair, which means that data is first mirrored and then both copies of data are striped across multiple HDDs in a RAID set. When replacing a failed drive, only the mirror is rebuilt, i.e. the disk array controller uses the surviving drive in the mirrored pair for data recovery and continuous operation. Data from the surviving disk is copied to the replacement disk.

RAID 0+1 is also called *mirrored stripe*. The basic element of RAID 0+1 is a stripe. This means that the process of striping data across HDDs is performed initially and then the entire stripe is mirrored. If one drive fails, then the entire stripe is faulted. A rebuild operation copies the entire stripe, copying data from each disk in the healthy stripe to an equivalent disk in the failed stripe.

Disadv: This causes increased and unnecessary I/O load on the surviving disks and makes the RAID set more vulnerable to a second disk failure.

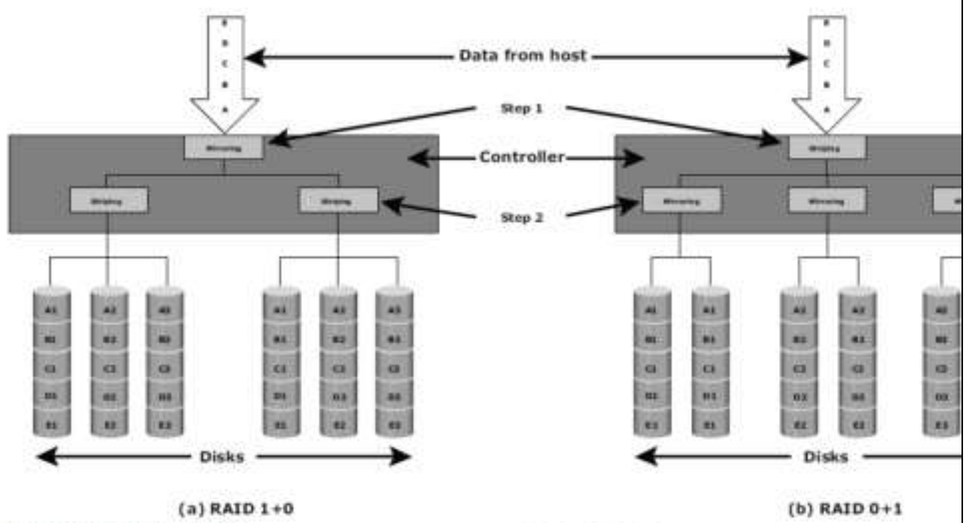

Figure 3-7: Nested RAID

b)RAID 5 and RAID 6 [5]

RAID 5

RAID 5 is a very versatile RAID implementation. It is similar to RAID 4 because it uses striping and the drives (strips) are independently accessible.

The difference between RAID 4 and RAID 5 is the **parity location**. In RAID 4, parity is written to a dedicated drive, creating a write bottleneck for the parity disk. In RAID 5, **parity is distributed** across all disks. The distribution of parity in RAID 5 overcomes the write bottleneck. Figure 3-9 illustrates the RAID 5 implementation.

RAID 5 is preferred for messaging, data mining, medium-performance media serving, and relational database management system (RDBMS) implementations in which database administrators (DBAs) optimize data access.

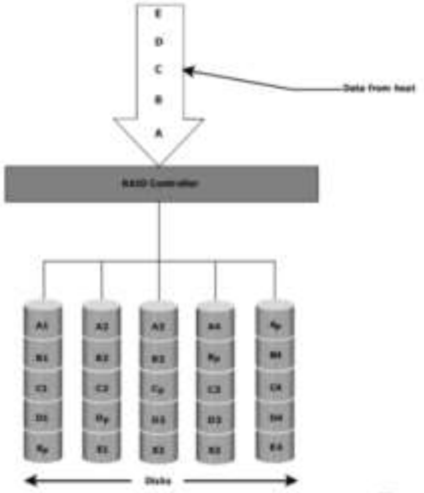

Figure 3-9: RAID 5

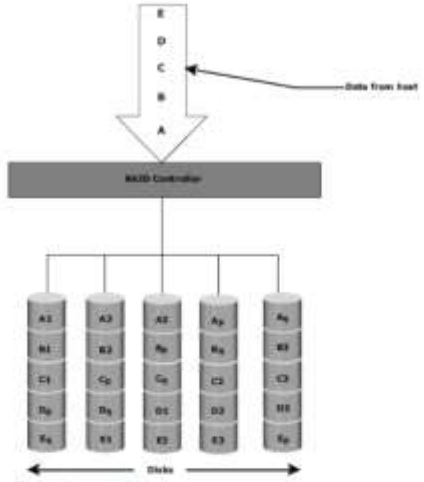

Figure 3-10: RAID 6

RAID 6

RAID 6 includes a second parity element to enable survival in the event of the failure of two disks in a RAID group (Figure 3-10). Therefore, a RAID 6 implementation requires at least four disks. RAID 6 distributes the parity across all the disks. The write penalty in RAID 6 is more than that in RAID 5; therefore, RAID 5 writes perform better than RAID 6. The rebuild operation in RAID 6 may take longer than that in RAID 5 due to the presence of two parity sets.

Comparison:

- RAID 6 is similar to RAID 5, except that it includes a second parity element to allow survival in the event of two disk failures.
  - The probability for this to happen increases and the number of drives in the array increases.
  - Calculates both horizontal parity (as in RAID 5) and diagonal parity.
  - Has more write penalty than in RAID 5.
  - Rebuild operation may take longer than on RAID 5.

4 Explain the components of an Intelligent Storage System (ISS) in detail with a diagram. [10] CO2 L2

Sol : An intelligent storage system consists of four key components:

i) front end
ii) cache
iii) back end
iv) physical disks.

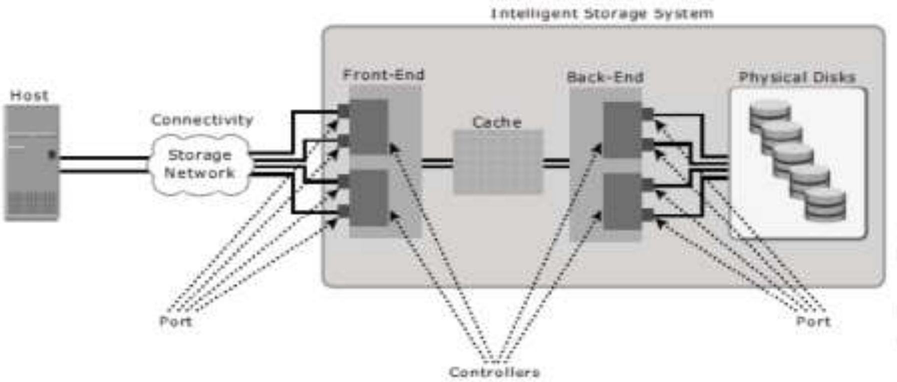

Figure 4-1: Components of an intelligent storage system

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 the 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 cache.

i) Front End

The front end provides the interface between the storage system and the host. It consists of two components: **front-end ports** and **front-end controllers**.

1. Front-end ports: enable hosts to connect to the intelligent storage system. Each front-end port has processing logic that executes the appropriate transport protocol, such as SCSI, Fibre Channel, or iSCSI, for storage connections. Redundant ports are provided on the front end for high availability.

2. Front-end controllers: route data to and from cache via the internal data bus.

ii) Cache

Cache is an important component that enhances the I/O performance in an intelligent storage system. Cache is semiconductor memory where data is placed temporarily to reduce the time required to service I/O requests from the host.

Cache improves storage system performance by isolating hosts from the mechanical delays associated with physical disks, which are the slowest components of an intelligent storage system. Accessing data from cache takes less than a millisecond. Write data is placed in cache and then written to disk. After the data is securely placed in cache, the host is acknowledged immediately.

Structure of Cache: Cache is organized into pages or slots, which is the smallest unit of cache allocation. The size of a cache page is configured according to the application I/O size. Cache consists of the *data store* and *tag RAM.* The data store holds the data while tag RAM tracks the location of the data in the data store

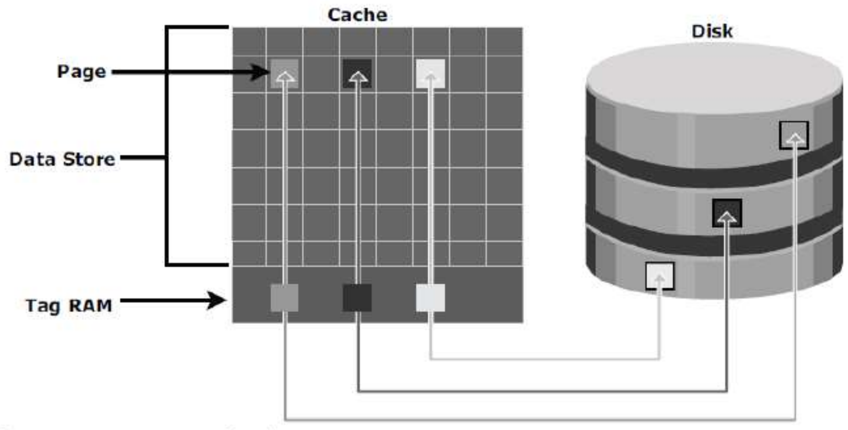

Figure 4-3: Structure of cache

Entries in tag RAM indicate where data is found in cache and where the data belongs on the disk. Tag RAM includes a *dirty bit* flag, which indicates whether the data in cache has been committed to the disk or not. It also contains time-based information, such as the time of last access, which is used to identify cached information that has not been accessed for a long period and may be freed up.

Read Operation with Cache: When a host issues a read request, the front-end controller accesses 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 (see Figure 4-4[a]). This provides a fast response time to the host (about a millisecond).

If the requested data is not found in cache, it is called a ***read cache miss*** *or* ***read miss*** and the data must be read from the disk (see Figure 4-4[b]).

The back-end controller accesses the appropriate disk and retrieves the requested data. Data is then placed in cache and is finally sent to the host through the front-end controller. Cache misses increase I/O response time.

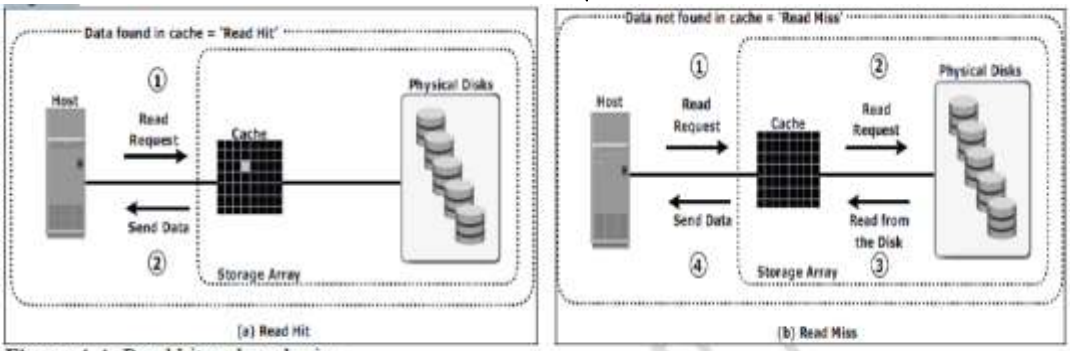

Figure 4-4: Read hit and read miss

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 less time (from the host's perspective) than it would take to write directly to disk. A write operation with cache is implemented in the following ways:

1. 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, as the write operations are isolated from the mechanical delays of the disk. However, uncommitted data is at risk of loss in the event of cache failures.

2. 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 write response time is longer because of the disk operations.

iii) Back End

The *back end* provides an interface between cache and the physical disks. It consists of two components: **back-end ports** and **back-end controllers.** The back end controls data transfers between cache and the physical disks. From cache, data is sent to the back end and then routed to the destination disk. Physical disks are connected to ports on the back end. The back end controller communicates with the disks when performing reads and writes and also provides additional, but limited, temporary data storage. The algorithms implemented on back-end controllers provide error detection and correction, along with RAID functionality.

iv) Physical Disk

A physical disk stores data persistently. Disks are connected to the back-end with either SCSI or a Fibre Channel interface (discussed in subsequent chapters). An intelligent storage system enables the use of a mixture of SCSI or Fibre Channel drives and IDE/ATA drives.

	Logical Unit Number: Physical drives or groups of RAID protected drives can be logically split into volumes known as logical volumes, commonly referred to as ***Logical Unit Numbers* (LUNs)**. The use of LUNs improves disk utilization. For example, without the use of LUNs, a host requiring only 200 GB could be allocated an entire 1TB physical disk. Using LUNs, only the required 200 GB would be allocated to the host, allowing the remaining 800 GB to be allocated to other hosts.			
5	Discuss Physical and Logical components of the Host. Sol : Applications runs on hosts or Servers. Hosts can range from simple laptops to complex server clusters ***Physical Components*** A host has three key physical components: ■ Central processing unit (CPU) ■ Storage, such as internal memory and disk devices ■ Input/Output (I/O) devices The CPU consists of four main components: ■**Arithmetic Logic Unit (ALU):** This is the fundamental building block of the CPU. It performs arithmetical and logical operations such as addition, subtraction, and Boolean functions (AND, OR, and NOT). ■ **Control Unit:** A digital circuit that controls CPU operations and coordinates the functionality of the CPU. ■ **Register:** A collection of high-speed storage locations. The registers store intermediate data that is required by the CPU to execute an instruction and provide fast access because of their proximity to the ALU. CPUs typically have a small number of registers. ■ **Level 1 (L1) cache:** Found on modern day CPUs, it holds data and program instructions that are likely to be needed by the CPU in the near future. The L1 cache is slower than registers, but provides more storage space. ***Storage*** Memory and storage media are used to store data, either persistently or temporarily. Memory modules are implemented using semiconductor chips, whereas storage devices use either magnetic or optical media. There are two types of memory on a host: ■ **Random Access Memory (RAM):** This allows direct access to any memory location and can have data written into it or read from it. RAM is volatile; this type of memory requires a constant supply of power to maintain memory cell content. Data is erased when the system's power is turned off or interrupted. ■ **Read-Only Memory (ROM):** Non-volatile and only allows data to be read from it. ROM holds data for execution of internal routines, such as system startup. ***I/O Devices*** I/O devices enable sending and receiving data to and from a host. This communication may be one of the following types: ■**User to host communications:** Handled by basic I/O devices, such as the keyboard, mouse, and monitor. These devices enable users to enter data and view the results of operations. ■**Host to host communications:** Enabled using devices such as a Network Interface Card (NIC) or modem. ■**Host to storage device communications:** Handled by a *Host Bus Adaptor (HBA)*. HBA is an application-specific integrated circuit (ASIC) board that performs I/O interface functions between the host and the storage, relieving the CPU from additional I/O processing workload. i) ***Logical Components of the Host*** a) ***Application***	[10]	CO1	L2

- Applications is an Interface between user and the host . It has
- Three-tiered architecture - Application UI, computing logic and underlying databases
- Application data access can be classifies as:
  Block-level access: Data stored and retrieved in blocks, specifying the LBA
  File-level access: Data stored and retrieved by specifying the name and path of files

b) Operating system

- Resides between the applications and the hardware
- Controls the environment

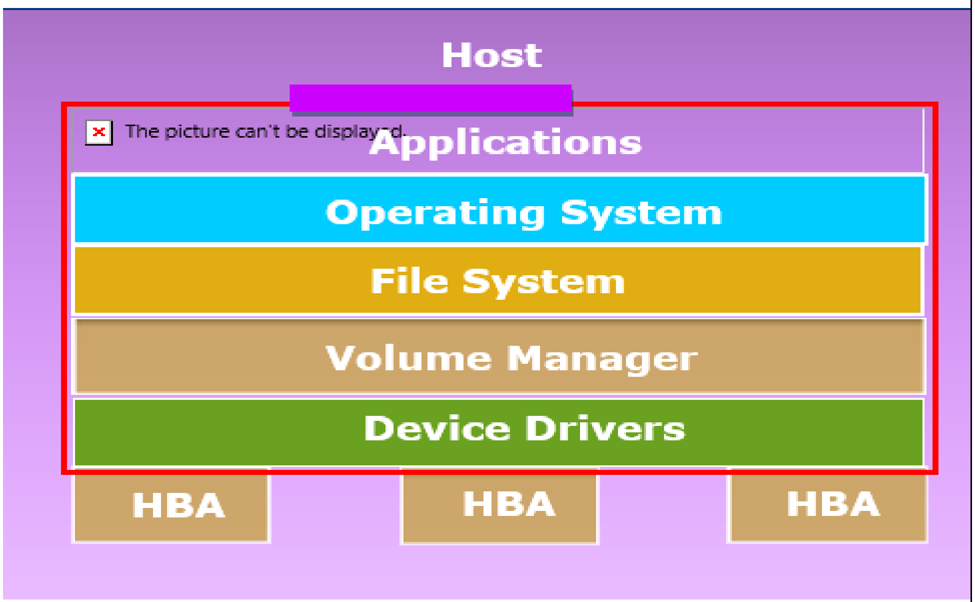

Logical data blocks are mapped to physical data blocks.
Logical Volume Managers
LVM Components are Physical Volumes and Volume Groups

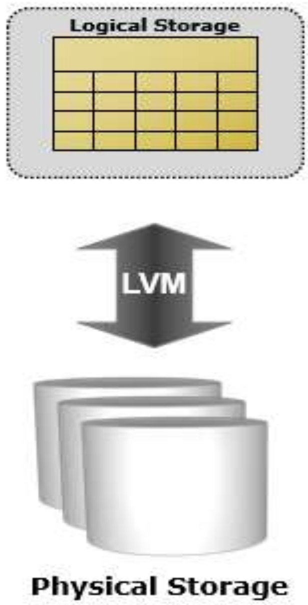

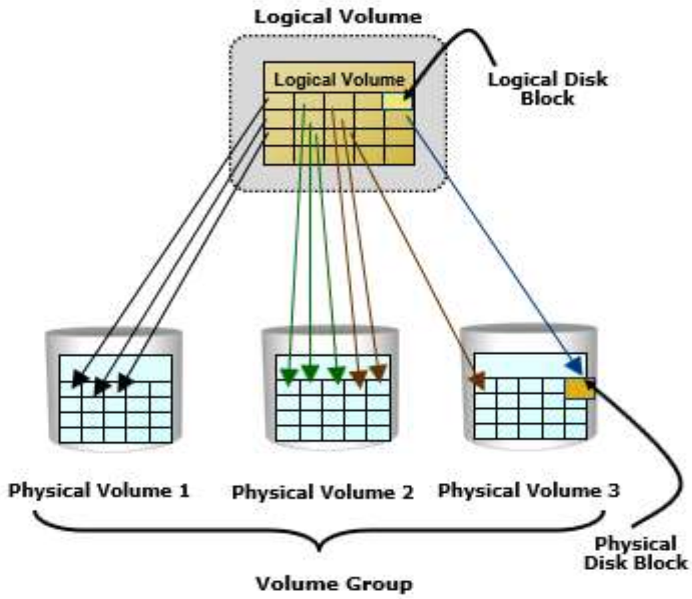

Volume Groups :

- One or more Physical Volumes form a Volume Group
- LVM manages Volume Groups as a single entity
- Physical Volumes can be added and removed from a Volume Group as necessary
- Physical Volumes are typically divided into contiguous equal-sized disk blocks
- A host will always have at least one disk group for the Operating System

Application and Operating System data maintained in separate volume groups

6	Briefly explain about the functionality of disk drive components with the help of neat and labeled diagram. A disk drive uses a rapidly moving arm to read and write data across a flat platter coated with magnetic particles. Data is transferred from the magnetic platter	[10]	CO1	L1

through the R/W head to the computer. Several platters are assembled together with the R/W head and controller, most commonly referred to as a *hard disk drive (HDD)*. Data can be recorded and erased on a magnetic disk any number of times.

Data on the disk is recorded on *tracks*, which are concentric rings on the platter around the spindle
The tracks are numbered, starting from zero, from the outer edge of the platter. Each track is divided into smaller units called *sectors*. A sector is the smallest, individually addressable unit of storage.

The key components of disk drive are :

- Platter
- spindle
- read/write head
- actuator arm assembly
- controller

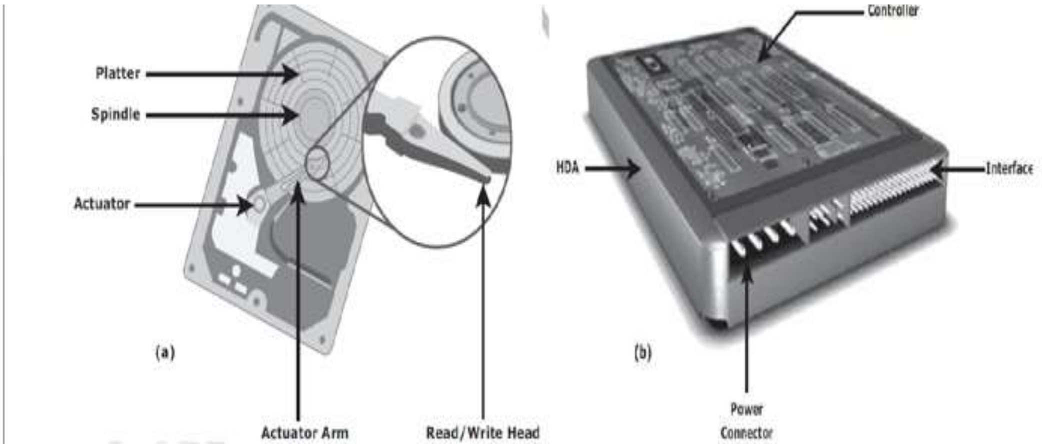

Figure 2-2: Disk Drive Components

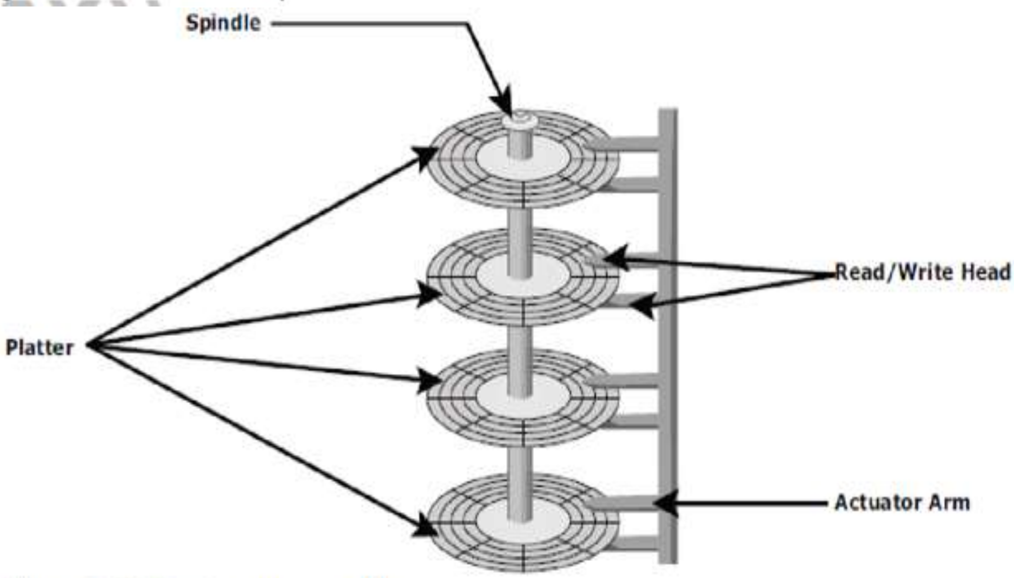

Figure 2-4: Actuator arm assembly

Platter :

- HDD consists of one or more flat circular disks called platters
- The data is recorded on these platters in binary (0s &1s).
- The set of rotating platters is sealed in a case, called a
- *Head Disk Assembly (HDA)*.
- The data is encoded by polarizing the magnetic area, of the disk surface.

	➢ Data can be written to or read from both surfaces of the platter **Spindle:** ➢ A spindle connects all the platters and 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). ➢ Disk drives have spindle speeds of 7,200 rpm, 10,000 rpm, or 15,000 rpm. ➢ platter - diameter of 3.5‖ (90 mm). **Read/Write head:** ➢ Drives have two R/W heads per platter to Read and write data one for each surface . ➢ when writing data- changes the magnetic polarization ➢ while reading data- detects magnetic polarization ➢ ensures that heads are moved to the landing zone before they touch the surface. ➢ If the R/W head accidentally touches the surface of the platter outside the landing zone, a ***head crash .*** **Actuator Arm Assembly:** ➢ The R/W heads are mounted on the actuator arm assembly ➢ To read or write data, R/W heads positions at the location on the platter ➢ The R/W heads for all platters are attached to one actuator arm assembly and move across the platters simultaneously. **Controller :** ➢ The *controller* is a printed circuit board, mounted at the bottom of a disk drive. ➢ It consists of a microprocessor, internal memory, circuitry, and firmware. ➢ The firmware controls power to the spindle motor and the speed of the motor. ➢ It also manages communication between the drive and the host.			
7	What do you know about Information Lifecycle Management? Discuss the information management lifecycle of an order processing system with the help of diagram. Discuss benefits of implementing ILM. **Information Lifecycle** The ***information lifecycle*** is the ―change in the value of information‖ over time. When data is first created, it often has the highest value and is used frequently. As data ages, it is accessed less frequently and is of less value to the organization. **For example**, in a sales order application, the value of the information changes from the time the order is placed until the time that the warranty becomes void (see Figure 1-7). The value of the information is highest when a company receives a new sales order and processes it to deliver the product. After order fulfillment, the customer or order data need not be available for real-time access. The company can transfer this data to less expensive secondary storage with lower accessibility and availability requirements unless or until a warranty claim or another event triggers its need.	[10]	CO2	L3

Information Lifecycle Management
Today's business requires data to be protected and available 24 × 7. Data centers can accomplish this with the optimal and appropriate use of storage infrastructure.

Information lifecycle management (ILM) is a proactive strategy that enables an IT organization to effectively manage the data throughout its lifecycle, based on predefined business policies. This allows an IT organization to optimize the storage infrastructure for maximum return on investment.

An ILM strategy should include the following characteristics:
1 Business-centric: It should be integrated with key processes, applications, and initiatives of the business to meet both current and future growth in information.
2 Centrally managed: All the information assets of a business should be under the purview of the ILM strategy.
3 Policy-based: The implementation of ILM should not be restricted to a few departments. ILM should be implemented as a policy and encompass all business applications, processes, and resources.
4 Heterogeneous: An ILM strategy should take into account all types of storage platforms and operating systems.
5 Optimized: Because the value of information varies, an ILM strategy should consider the different storage requirements and allocate storage resources based on the information's value to the business.

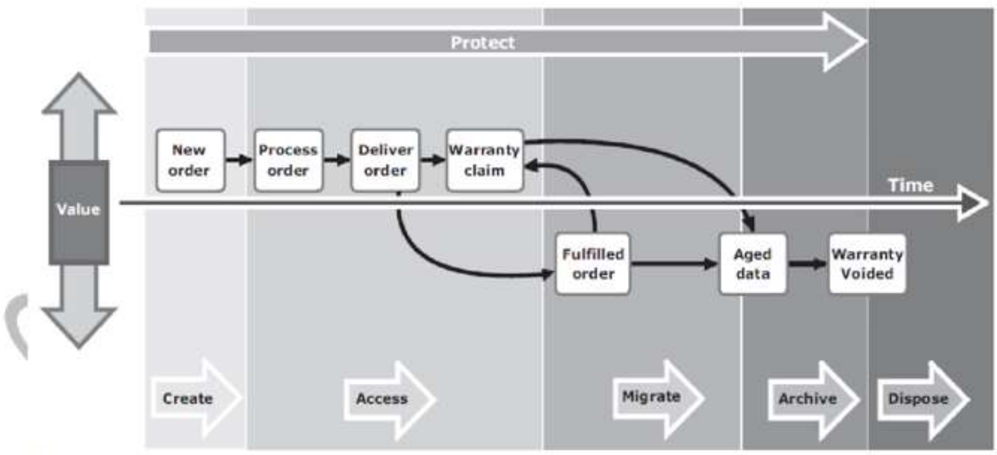

Figure 1-7: Changing value of sales order information

ILM Implementation
The process of developing an ILM strategy includes four activities—classifying, implementing, managing, and organizing:
1 Classifying data and applications on the basis of business rules and policies to enable differentiated treatment of information
2 Implementing policies by using information management tools, starting from the creation of data and ending with its disposal
3 Managing the environment by using integrated tools to reduce operational complexity
4 Organizing storage resources in tiers to align the resources with data classes, and storing information in the right type of infrastructure based on the information's current value.

Implementing ILM across an enterprise is an ongoing process. Figure 1-8 illustrates a three-step road map to enterprise-wide ILM.

Step 1 the goal is to implement a storage networking environment. Storage architectures offer varying levels of protection and performance and this acts as a foundation for future policy-based information management in Steps 2 and 3. The value of tiered storage platforms can be exploited by allocating appropriate storage resources to the applications based on the value of the information processed.

Step 2 takes ILM to the next level, with detailed application or data classification and linkage of the storage infrastructure to business policies. These classifications and the resultant policies can be automatically executed using tools for one or more applications, resulting in better management and optimal allocation of storage resources.

Step 3 of the implementation is to automate more of the applications or data classification and policy management activities in order to scale to a wider set of enterprise applications.

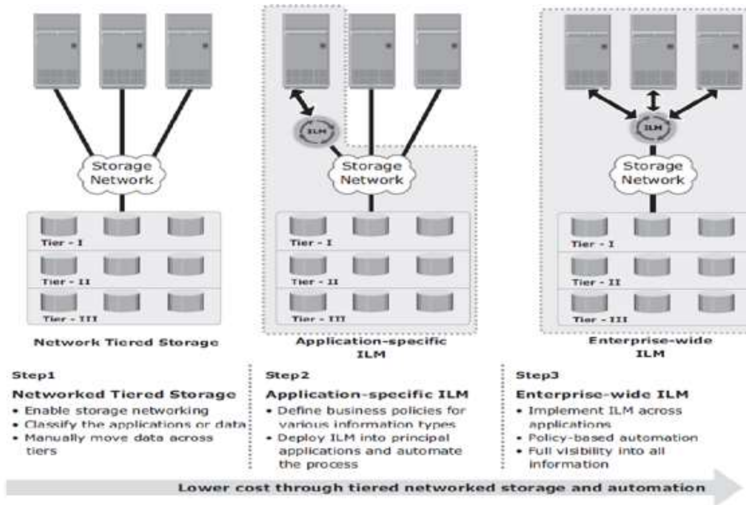

Figure 1-8: Implementation of ILM

ILM Benefits

Implementing an ILM strategy has the following key benefits that directly address the challenges of information management:

1 Improved utilization by using tiered storage platforms and increased visibility of all enterprise information.

2 Simplified management by integrating process steps and interfaces with individual tools and by increasing automation.

3 A wider range of options for backup, and recovery to balance the need for business continuity.

4 Maintaining compliance by knowing what data needs to be protected for what length of time.

5 Lower Total Cost of Ownership **(TCO)** by aligning the infrastructure and management costs with information value. As a result, resources are not wasted, and complexity is not introduced by managing low-value data at the expense of high-value data.

8	(a) Discuss measures of disk drive performance.	[6]	CO1	L1

Sol :

A disk drive is an electromechanical device that governs the overall performance of the storage system environment. The various factors that affect the performance of disk drives are discussed in this section.

Disk Service Time : Disk service time is the time taken by a disk to complete an I/O request. Components that contribute to service time on a disk drive are seek time, rotational latency, and data transfer rate.

1. **Seek Time**

The seek time (also called access time) describes the time taken to position the R/W heads across the platter with a radial movement (moving along the radius of the platter). In other words, it is the time taken to reposition and settle the arm and the head over the correct track. The lower the seek time, the faster the I/O operation. Disk vendors publish the following seek time specifications:

■Full Stroke: The time taken by the R/W head to move across the entire width of the disk, from the innermost track to the outermost track.

■Average: The average time taken by the R/W head to move from one random track to another, normally listed as the time for one-third of a full stroke.

■Track-to-Track: The time taken by the R/W head to move between adjacent tracks.

2. Rotational Latency

To access data, the actuator arm moves the R/W head over the platter to a particular track while the platter spins to position the requested sector under the R/W head. The time taken by the platter to rotate and position the data under the R/W head is called rotational latency. This latency depends on the rotation speed of the spindle and is measured in milliseconds.

3. Data Transfer Rate

The data transfer rate (also called transfer rate) refers to the average amount of data per unit time that the drive can deliver to the HBA. In a read operation, the data first moves from disk platters to R/W heads, and then it moves to the drive's internal buffer. Finally, data moves from the buffer through the interface to the host HBA. In a write operation, the data moves from the HBA to the internal buffer of the disk drive through the drive's interface. The data then moves from the buffer to the R/W heads. Finally, it moves from the R/W heads to the platters.

Internal transfer rate is the speed at which data moves from a single track of a platter's surface to internal buffer (cache) of the disk. Internal transfer rate takes into account factors such as the seek time.

External transfer rate is the rate at which data can be moved through the interface to the HBA. External transfer rate is generally the advertised speed of the interface, such as 133 MB/s for ATA.

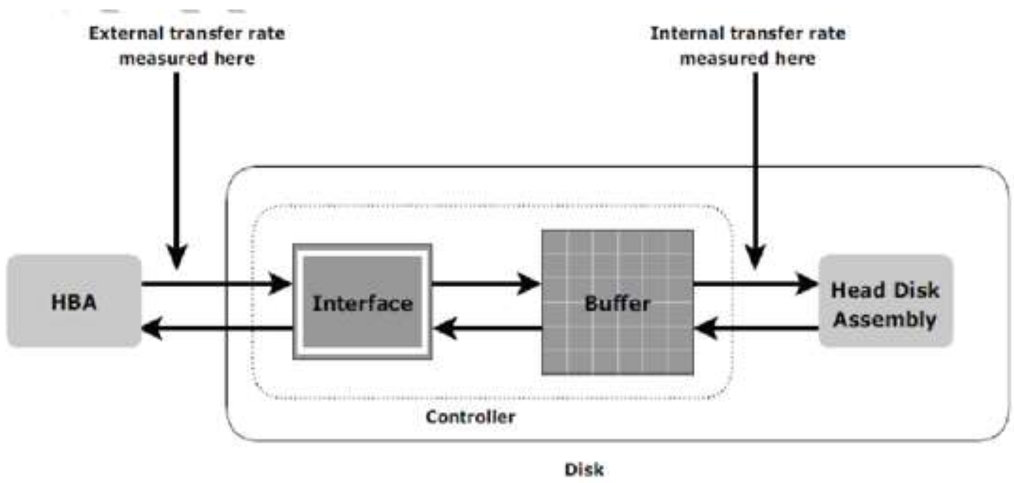

Figure 2-8: Data transfer rate

	b) Consider a disk I/O system in which an I/O request arrives at a rate of 100 I/Os per second. The service time, R_S, is 8 ms. Calculate utilization of I/O controller (U) and total time spent by a request in a queue. Solution: Arrival rate a= 100 I/O s =100/1000 I/O ms= 1/10 I/O ms Ra = 1/a = 10 ms R_S = 8 ms (given) Utilization (U)= R_S/ Ra= 8/10= 0.8% or 80% Response time(R)= R_S/(1-U) = 8/(1=0.8)= 40ms Total time spent by a request in a queue = UxR= 0.8x40= 32 ms	[4]	CO2	L2