UNIT 2

Data Center Environment

This chapter provides an understanding of various logical components of hosts such as file systems, volume managers, and operating systems, and their role in the storage system environment.

2.1 Components of a Storage System Environment

Application

An application is a computer program that provides the logic for computing operations. Applications are

Infrastructure management application
Data protection application
e-mails
enterprise resource planning
DSS

DBMS

More commonly, a database management system (DBMS) provides a structured way to store data in logically organized tables that are interrelated. A DBMS optimizes the storage and retrieval of data.

Host

Users store and retrieve data through applications. The computers on which these applications run are referred to as hosts. Hosts can range from simple
Users store and retrieve data through applications. The computers on which these applications run are referred to as hosts.

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

CPU

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. Memory modules enable data access at a higher speed than the storage media. Generally, 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.Storage devices are less expensive than semiconductor memory.

Examples of storage devices are as follows:

■ Hard disk (magnetic)

■ CD-ROM or DVD-ROM (optical)

■ Floppy disk (magnetic)

■ Tape drive (magnetic)

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. HBAs also provide connectivity outlets known as ports to connect the host to the storage device. A host may have multiple HBAs.

Connectivity

Connectivity refers to the interconnection between hosts or between a host and any other peripheral devices, such as printers or storage devices. The discussion here focuses on the connectivity between the host and the storage device. The components of connectivity in a storage system environment can be classified as physical and logical. The physical components are the hardware elements that connect the host to storage and the logical components of connectivity are the protocols used for communication between the host and storage. The communication protocols are covered in Chapter 5.

Physical Components of Connectivity

The three physical components of connectivity between the host and storage are Bus, Port, and Cable .

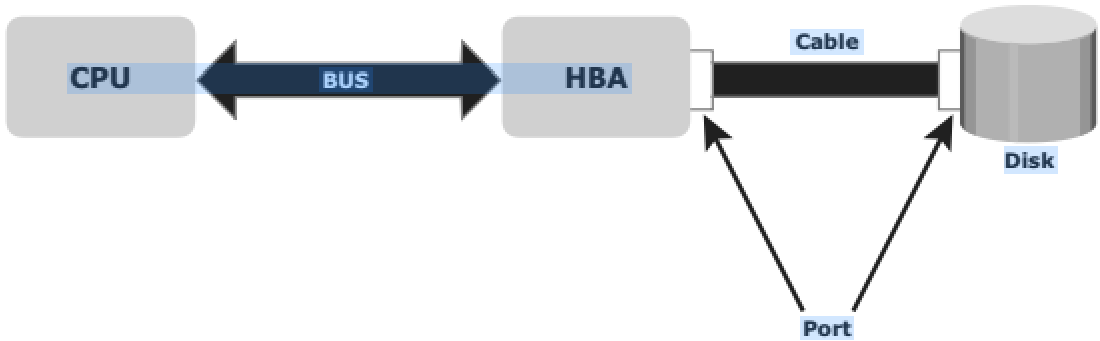

Figure 2-1: Physical components of connectivity

Logical Components of Connectivity

PCI

IDE/ATA

SCSI

Storage

Disk Drive Components

Key components of a disk drive are platter, spindle, read/write head, actuator arm assembly, and controller .

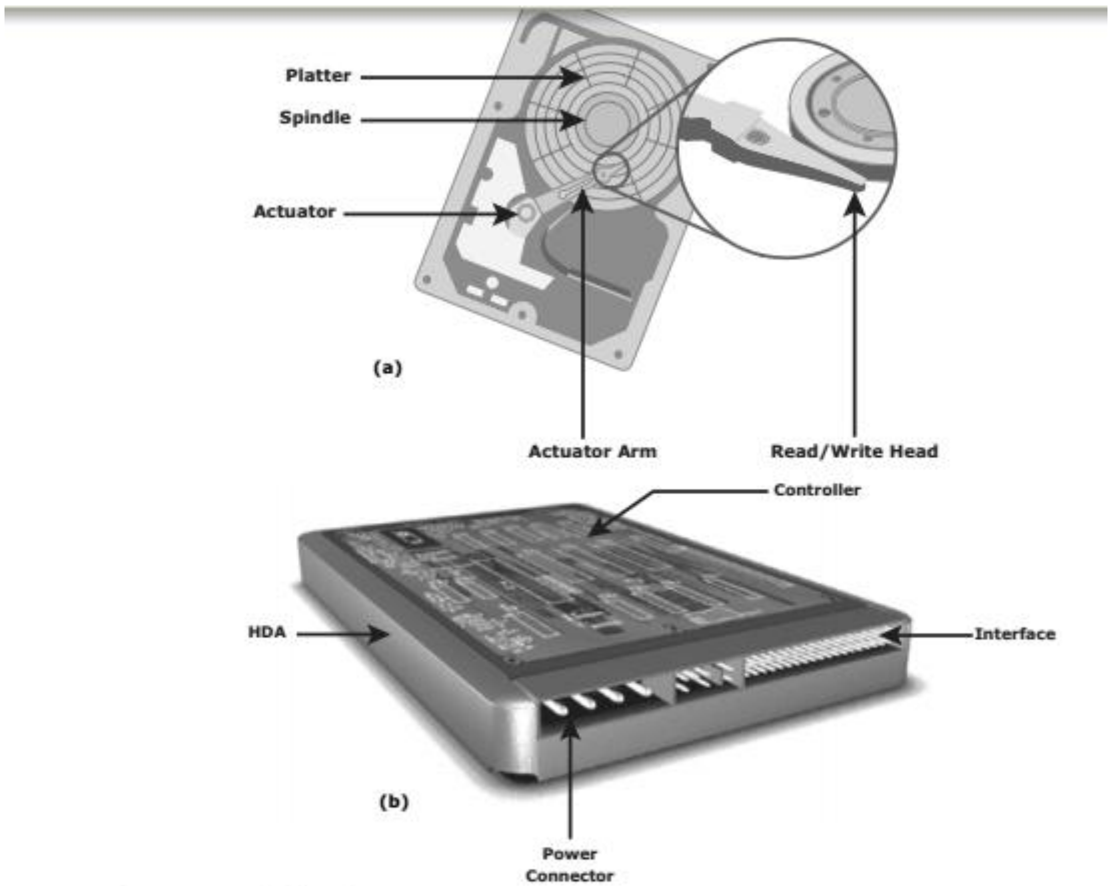

Figure 2-2: Disk Drive Components

Platter and Spindle

A typical HDD consists of one or more flat circular disks called platters.A spindle connects all the platters,

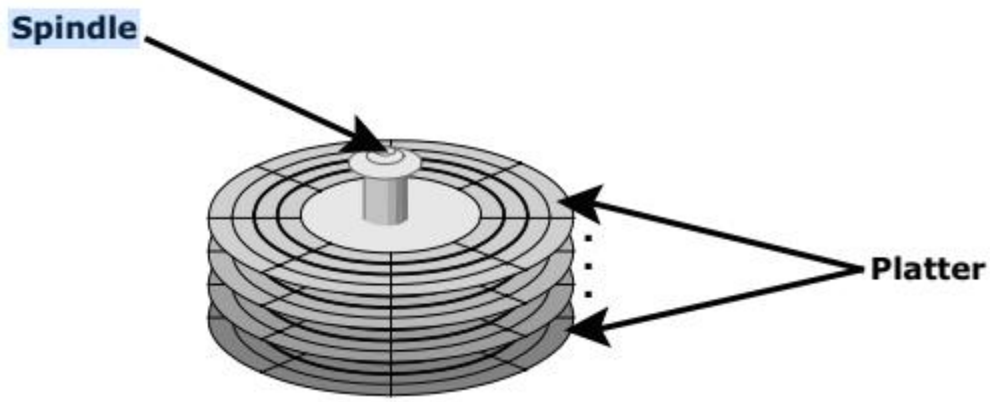

Figure 2-3: Spindle and platter

Read/Write Head

Read/Write (R/W) heads, shown in Figure 2-4, read and write data from or to a platter. Drives have two R/W heads per platter, one for each surface of the platter.

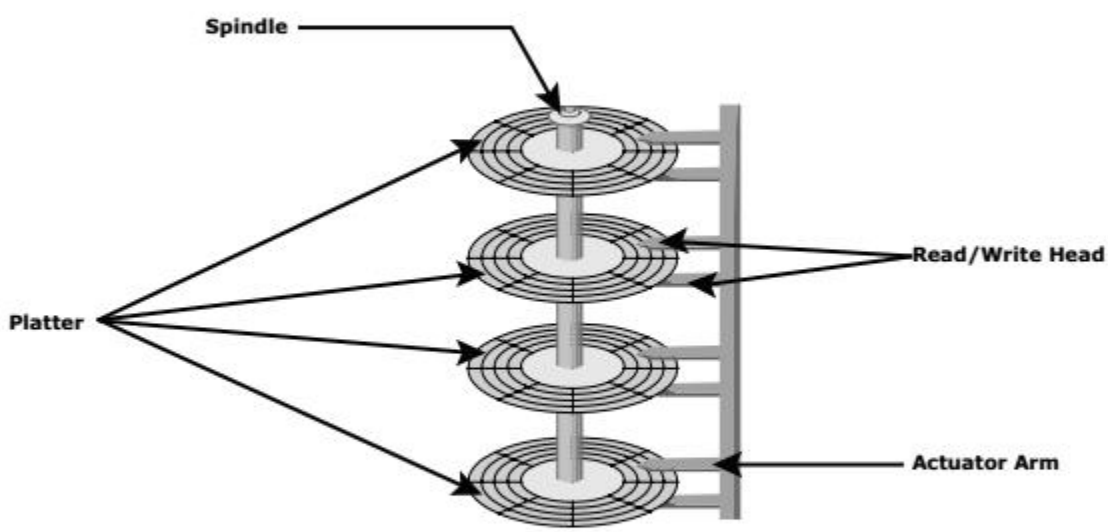

Figure 2-4: Actuator arm assembly

Physical Disk Structure

Data on the disk is recorded on tracks, which are concentric rings on the platter around the spindle, as shown in Figure 2-5. The tracks are numbered, starting from zero, from the outer edge of the platter. The number of tracks per inch (TPI)on the platter (or the track density) measures how tightly the tracks are packed on a platter.

Each track is divided into smaller units called sectors. A sector is the smallest, individually addressable unit of storage. The track and sector structure is written on the platter by the drive manufacturer using a formatting operation. The number of sectors per track varies according to the specific drive. The first personal computer disks had 17 sectors per track. Recent disks have a much larger number of sectors on a single track. There can be thousands of tracks on a platter, depending on the physical dimensions and recording density of the platter.

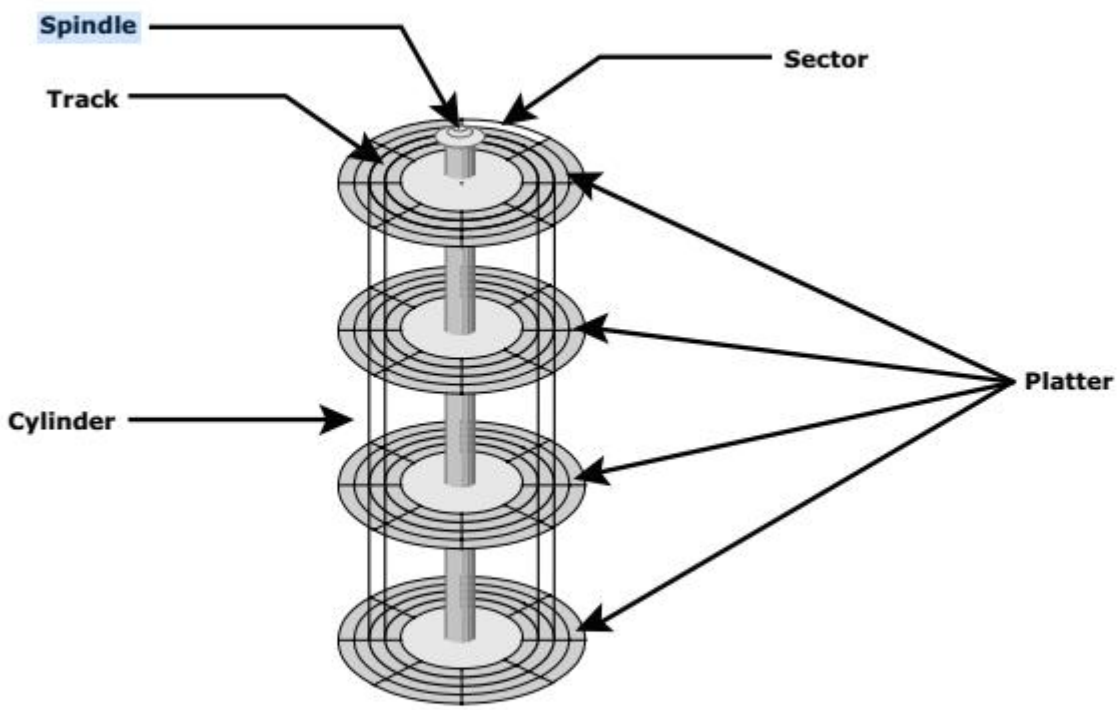

Figure 2-5: Disk structure: sectors, tracks, and cylinders

Zoned Bit Recording

Because the 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, as shown in Figure 2-6 (a). 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 inefficient use of available space.

Zone bit recording utilizes the disk efficiently. As shown in Figure 2-6 (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.

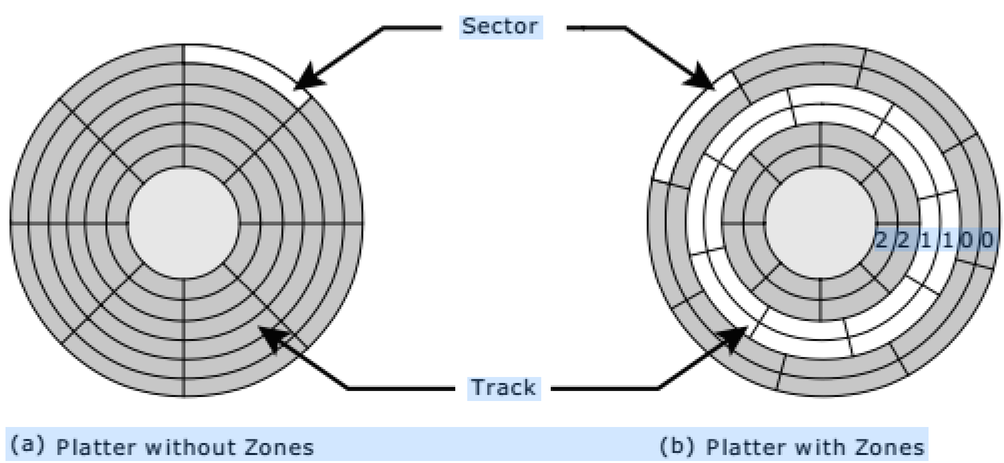

Figure 2-6: Zoned bit recording

Logical Block Addressing

Earlier drives used physical addresses consisting of the cylinder, head, and sector (CHS) number to refer to specific locations on the disk, as shown in Figure 2-7 (a), and the host operating system had to be aware of the geometry of each disk being used. Logical block addressing (LBA), shown in Figure 2-7 (b), simplifies addressing by using a linear address to access physical blocks of data. The disk controller translates LBA to a CHS address, and the host only needs to know 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.

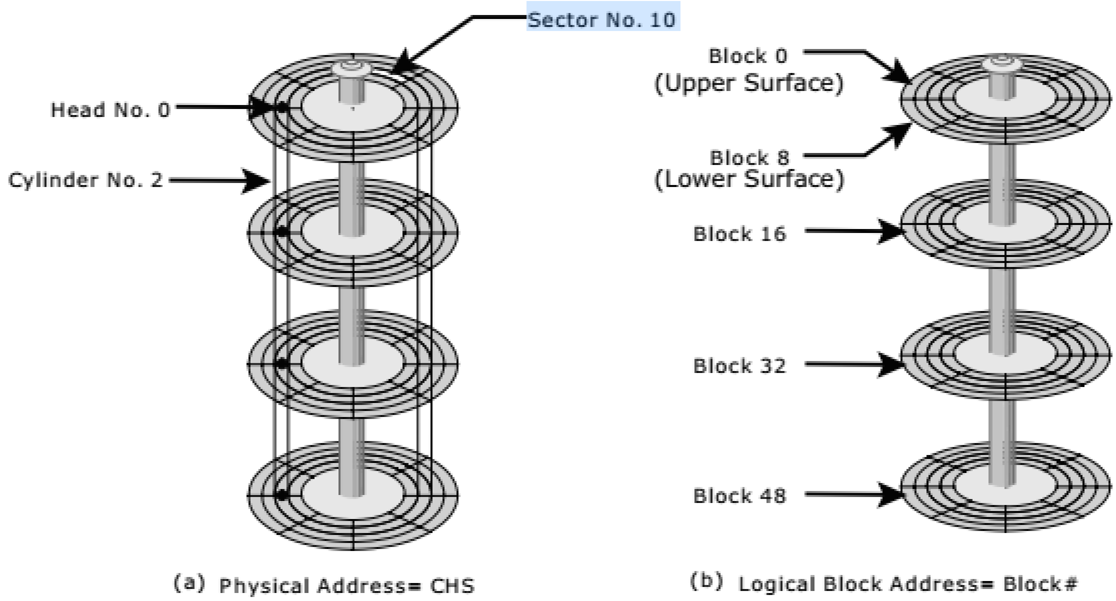

Figure 2-7: Physical address and logical block address

Disk Drive Performance

A disk drive is an electromechanical device that governs the overall perfor-mance 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. that contribute to service time on a disk drive are seek time, rotational latency, and data transfer rate.

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.

Each of these specifications is measured in milliseconds. The average seek time on a modern disk is typically in the range of 3 to 15 milliseconds. Seek time has more impact on the read operation of random tracks rather than adjacent tracks.

To minimize the seek time, data can be written to only a subset of the available cylinders. This results in lower usable capacity than the actual capacity of the drive. For example, a 500 GB disk drive is set up to use only the first 40 percent of the cylinders and is effectively treated as a 200 GB drive. This is known as short-stroking the drive.

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 34 Section I n Storage Systemthe R/W head is called rotational latency. This latency depends on the rotation speed of the spindle and is measured in milliseconds. The average rotational latency is one-half of the time taken for a full rotation. Similar to the seek time, rotational latency has more impact on the reading/writing of random sectors on the disk than on the same operations on adjacent sectors.Average rotational latency is around 5.5 ms for a 5,400-rpm drive, and around 2.0 ms for a 15,000-rpm drive.

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. It is important to first understand the process of read and write operations in order to calculate data transfer rates. 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.

The data transfer rates during the R/W operations are measured in terms of internal and external transfer rates,.

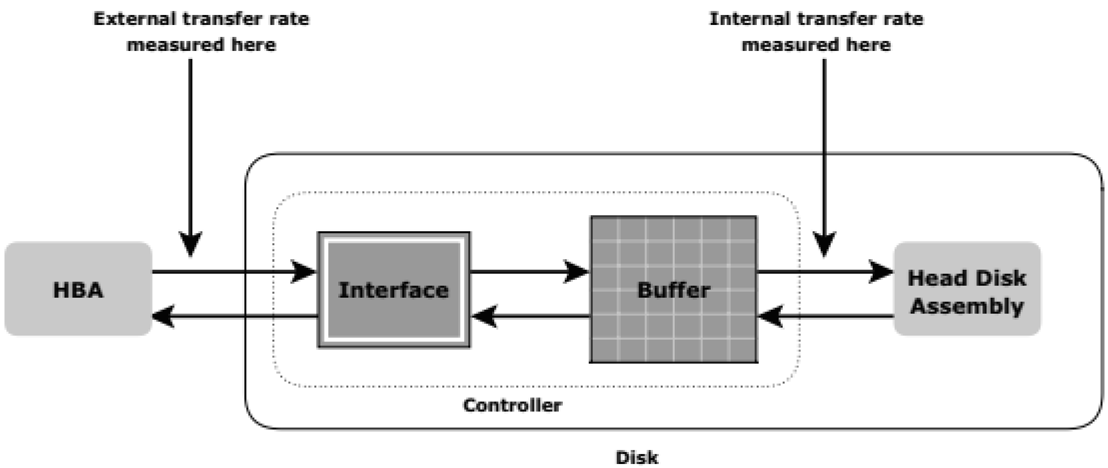

Figure 2-8: Data transfer rate

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. The sustained external transfer rate is lower than the interface speed.