

Types of Storage Devices

Physical components or materials on which data is stored are called storage media. Hardware components that read/write to storage media are called storage devices.

Two main categories of storage technology used today are magnetic storage and optical storage.

Primary magnetic storage

- o Diskettes
- o Hard disks (both fixed and removable)
- o High capacity floppy disks
- o Disk cartridges
- o Magnetic tape

Primary optical storage

- o Compact Disk Read Only Memory (CD ROM)
- o Digital Video Disk Read Only Memory (DVD ROM)
- o CD Recordable (CD R)
- o CD Rewritable (CD RW)
- o Photo CD

Magnetic Storage Devices

Purpose of storage devices → to hold data even when the computer is turned off so the data can be used whenever needed. Storage involves writing data to the medium and reading from the medium. Writing data → recording the data on the surface of the disk where it is stored for later use. Reading data → retrieving data from the surface and transferring it into the computers memory for use.

Diskette drives, hard drives and tape drive all use the same type of medium → use similar techniques for reading/writing data. Surfaces of diskettes and magnetic tape are all coated with a magnetically sensitive material such as iron oxide.

The principle use to store data is that of polarisation – all the ions in the magnetic material align themselves in one direction. Just as a transistor can represent binary “on” or “off”, the orientation of the magnetic field can be used to represent data. A magnet has one important advantage over a transistor → it can maintain it’s state without a continual supply of electricity.

Surfaces of disks are coated with millions of tiny iron particles so data can be stored on them. Each particle acts as a magnet, taking on a magnetic field when subjected to an electromagnet. The read/write heads of a disk drive contain electromagnets, which generate magnetic fields in the iron on the storage medium as the head passes over the disk. The read/write heads store a string of 1s and 0s by alternating the direction of the current in the electromagnet.

To read data the process is reversed. The read/write heads pass over the disk while no current is flowing through the electromagnet. The storage medium has a charge and the electromagnet does not → storage medium changes the magnet in the head which

causes a small current to flow through the head in one direction or the other depending on the polarity of the field. The disk drive senses direction of flow and data is sent from the head into memory.

Magnetic Disks

Diskette drives and hard disk drives are the most commonly used storage devices in PCs. Both fall into the magnetic storage category because they record data as magnetic fields.

Fundamental differences and similarities between the two:

- A diskette contains a single flat piece of plastic (the disk) coated with iron oxide enclosed in vinyl or plastic cover. A hard disk contains one or more rigid metal platters coated with iron oxide permanently enclosed in a hard disk drive.
- Diskettes are small and portable (they can be removed from diskette drives). Hard disks are usually built into the computer and are not portable (unless the computer is). Exceptions are removable hard disks and external hard drives which can be detached from the system.
- Floppy disks store only 1.44 MB although special floppy disks offer higher capacity. New hard disks can store several thousand times as much data as a diskette.
- Hard drives are much faster than diskettes, their disks spin faster and they locate data on the disks surface in much less time.

Almost all PCs sold today come with a hard disk and one disk drive. Some computers also feature a third built in magnetic device – a device that uses high capacity floppy disks.

How data is organised on a magnetic disk

Before the computer can use a diskette to store data, the disks surface must be magnetically mapped so that the computer can go directly to a specific point without searching through all the data. This process of mapping a disk is called **formatting** or **initialising**. When you purchase a new disk, they should be formatted for either PC or Mac.

It may be helpful to reformat disks from time to time as this deletes all the data on disk. During the formatting you process you can also determine whether the disk has any faulty spots and you can copy important system files onto the disk.

Hard disks must also be formatted so that the computer can locate data on them. When you buy a computer, the hard disk has already been formatted correctly and probably contains some programs and data. You can format your hard disk if necessary but the process is different to that for a diskette.

Modern diskettes store data on both sides of the disk (numbered side 0 and side 1) and each side has its own read/write head. When formatting a disk, the disk creates a set of magnetic concentric circles called **tracks**, on each side of the disk. The number of tracks required depends on the type of disk. Most high -density diskettes have 80

tracks on each side. A hard disk may have several hundred tracks on each side of each platter. Each track is a separate circle. These are numbered from the outermost circle to the innermost, starting with zero.

Each track on a disk is also split into smaller parts. Imagine slicing a disk as you would a pie. Each slice cuts across all the tracks resulting in short segments or sectors. A sector can contain up to 512 bytes. All the sectors are numbered in one long sequence so the computer can access each small area on the disk with a unique number. This scheme simplifies a 2 dimensional set of co-ordinates into a single numeric address.

If the diskette has 80 tracks on each side and each track contains 18 sectors → 1440 sectors per side → 2880 sectors per disk. This configuration is true regardless of the length of the track. The diskettes outermost track is longer than the innermost one but each track is still divided into the same number of sectors. Regardless of physical size all the sectors hold the same number of bytes i.e. the shortest, innermost sectors hold the same amount of data as the longest outermost ones.

A disks allocation of sectors per track is somewhat wasteful because the outer longer tracks could theoretically hold more data than the shorter inner ones. For this reason hard disks allocate more sectors to the longer tracks on the disks surface. As you move towards the disks centre, each subsequent track has fewer sectors. This arrangement takes advantage of the hard disks potential capacity and enables a typical hard disk to store data more efficiently than a floppy disk. Because many hard disks allocate sectors in this manner, their sectors per track specification are often given as an average. Such hard disks are described as having an *average of x sectors per track*.

Because files are not usually a size that is an even multiple of 512 bytes, some sectors contain unused space after the end of the file. In addition the Windows OS allocates a group of sectors called a **cluster** to each file stored on a disk. Cluster sizes vary on the size and type of disk, but can range from 4 → 64 sectors. A small file that contains only 50 bytes will only use a portion of the first sector of a cluster leaving the remainder of the sector and cluster allocated but unused.

The different capacities of disks are generally a function of the number of sides, tracks and sectors per track. Floppy disks have only a few standard types and formats but hard disks come in a variety of sizes so their capacities vary greatly.

A sector is the smallest unit with which any disk drive can work. Each bit and byte within a sector can have different values but the drive can read or write only whole sectors at a time. If the computer needs to change just 1 byte out of 512 it must rewrite the whole sector.

How the OS finds data on a disk

A computer's Operating system can locate data on a disk because each track and sector is labelled, and the location of all data is kept in a special log on the disk. The

labelling of tracks and sectors is called performing a **logical format** (also called **low level format** or **soft format**). A commonly used logical format by windows creates 4 disk areas

1. Master boot record (MBR)
2. File – allocation table (FAT)
3. Root folder or directory
4. Data Area

Master boot record is a program that runs when you first start the computer. It determines whether the disk has the basic components necessary to run the operating system correctly. If the program determines that the correct files are present and disk has a valid format, it transfers control to one of the OS programs that continues the system start up (known as **booting**). Boot record also describes characteristics such as number of bytes per sector and number of sectors per track.

FAT is a log that records the location of each file and the status of each sector. When you write to a disk, the OS checks the FAT for an open area, stores the data and then identifies the file and its location in the FAT. The FAT solves a common filing problem, what happens when you load a file; increase its size by adding to it and then save it again? The OS checks the FAT for free areas and places pointers that in it that links together the nonadjacent parts of the file. In other words it splits the file by allocating new space for the overflow, the pointers in FAT enable the OS to put the file back together again.

When a file is saved in this way it becomes **fragmented**. Fragmented files cause side effects such as taking longer to save and load.

Users do not see the information in FAT but often use the information. A **folder** (also called a **directory**) is a tool for organising files on a disk. Folders can contain files or other files so it is possible to set up hierarchical systems of folders on your computer. The top folder on any disk is called the **root folder** (or **root directory**). When you use the OS to view the contents of a folder, the OS lists specific information about each file in the folder such as name, size, date of creation or last modification and so on.

The part of the disk that remains free after the boot sector, FAT and root folder have been created is called the **data area**, because that is where the data is stored!

Diskettes (Floppy Disks)

The diskette drive includes a motor that rotates the disk on a spindle and the read/write heads that can move to any spot on the surface of the disk as it spins. This allows the heads to access data randomly rather than sequentially – the heads can skip from one spot to another without having to scan through all the data in between.

Diskettes spin at approx. 300 revolutions per minute. The longest it can take to position a point on the diskettes under the read/write heads is the amount of time for one revolution – 0.2 second. The farthest the heads have to move is from the centre of the diskette to the outside edge (or vice versa). The heads can do this in less time – about 0.17 seconds.

Because both operations occur at the same time (spinning the disk and moving the heads over the surface), the maximum time to position the heads over a given location on the diskette – known as maximum access time – remains the greater of the two times, 0.2 second.

The maximum access time for diskettes can be even longer, however, because they do not spin when they are not being used. It can take about 0.5 second to rotate the disk from a dead stop.

Most common uses for disks

- Moving files between computers not connected by a network or communication hardware – one of the easiest ways to move data from one computer to another is to copy it onto diskette and use the diskette in another computer.
- Loading new programs onto a system – Large programs are usually delivered by CD-ROM but many programs are still sold on diskette. You install it by copying the contents of the disks onto the hard drive and then running a small program, which installs the files automatically.
- Backing up data or programs – Primary copy of data or programs are stored on the hard drive, Backing up is the process of creating a duplicate copy for safekeeping. To protect against data loss it is wise to back up a hard disk. Due to the limited storage capacity diskettes are used to back up small data files rather than programs or entire hard disks.

Types of Diskettes

1980's → most PCs used 5.25-inch diskettes.

Now → almost completely replaced by the 3.5-inch.

The size refers to the diameter of the disk and is not an indication of capacity.

The 5.25 type is encased in a flexible vinyl envelope with an oval cutout that allows read/write heads access to the disk. The 3.5 type is encased in a hard plastic shell with a sliding shutter. On insertion into the drive the shutter slides back to expose the surface to the read/write heads.

Diskette capacities

Both types have evolved from lower to higher densities. The density is a measure of the capacity of the surface of the disk, the higher the density → the more closely the iron oxide particles are packed → the more data a disk can store.

Early versions of disks were double density (DD). As diskette media improved → storage capacity improved. DD diskettes have been replaced by high density (HD) ones providing significantly more storage.

Diameter (inches)	Type	Tracks	Sectors/ track	Sectors	Bytes/ sector	Total bytes
5.5	DD	40	9	720	512	360k
5.5	HD	80	15	2400	512	1200k
3.5	HD	80	9	1440	512	720k
3.5	HD	80	18	2880	512	1440k
3.5	HD	80	36	5760	512	2880k

Macs never used 5.25 diskettes.

DD Mac disk holds 800k not 720k.

The Extra density disk holds more data than a regular 1.44 MB disk but needs a special hard drive.

Hard Disks

Hard disk is still the most common storage device for all computers. Like diskettes hard drives store data in tracks divided into sectors. Physically however they look quite different to diskettes.

Includes one or more metal platters mounted on a central spindle, like a stack of rigid diskettes. Each platter is covered with a metal coating and the entire unit is contained in a sealed chamber. The hard disk and drive are a single unit which includes the hard disk, the motor that spins the platters and a set of read/write heads. Because you cannot remove the disk from the drive the terms hard disk and hard drive are used interchangeably.

Hard drives have become the primary storage devices for PCs because they are convenient and cost effective. They outperform diskettes in both speed and capacity. Hard disks offer capacities from several hundred MB and more. Most entry level PCs now come with hard disks of at least 6.8MB.

Two important physical differences between hard disks and diskettes account for the differences in performance.

1. Hard disks are sealed in a chamber
2. Hard disk consists of a rigid metal platter (usually aluminium) rather than flexible Mylar.

Hard disk spin between 3600 rpm and 7200 rpm compared to a diskette 300 rpm. The speed at which the disk spins is a major factor in overall performance. The rigidity of the disk and the high-speed rotation allows more data to be recorded on the surface. Waving a magnet past an electric coil causes a current to flow. The faster you wave the magnet and the closer the magnet is to the coil, the larger the current generated in the coil. The disk that spins faster can use smaller magnetic charge to make current flow to the read/write head. The drives heads can also use a lower density current to record data on the disk.

Not only do hard disks pack the data closer together they also hold more data because they usually include multiple platters stacked one on top of each other. This configuration means that the disk has more than 2 sides, in addition to side 0 and 1 there are sides 2, 3, 4, and so on. Larger capacity disks may use 12 platters but both side of every platter are not always used.

The number of sides that the disk uses is specified by the number of read/write heads. A particular disk may have 6 disks platters (12 sides) but only 11 heads → one side is not used to store data. Unused side is often the bottom one.

The term **cylinder** is often used to refer to the same track across all the disks. Track 0 (outermost track) on every disk is cylinder 0.

Hard disk generally store 512 bytes of data in a sector, but can have more sectors per track – 54, 63 or even more sectors per track are not uncommon.

Computation of a hard disks capacity is identical to that for other disks but the numbers are larger. Breakdown of capacity for a 50 GB hard disk, assuming 11 platters, 264,528 tracks, 369 sectors per track:

$$\begin{array}{rcll} 12024 \text{ cylinders} & \times & 22 \text{ heads (sides)} & \\ & & & = 264528 \text{ tracks} \\ 264528 \text{ tracks} & \times & 369 \text{ avg. sectors/track} & \\ & & & = 97,610,823 \text{ sectors} \\ 97610823 \text{ sectors} & \times & 512 \text{ bytes} & \\ & & & = 49,976,754,984 \text{ bytes} \\ & & & \text{(Approx)} \end{array}$$

If the actual numbers of sectors per track were taken into account instead of an average, the calculation would be more difficult and the resulting number of bytes would be higher.

In spite of capacity and speed advantages there is one major drawback → to achieve optimum performance the read/write heads must be extremely close to the surface of the disk without actually touching it. The read/write heads fly so close to the disk that a human hair, a dust particle or even a fingerprint would bridge the gap and cause the head to crash. A **head crash** – in which the head touches the disk – can destroy the data stored in the area of the crash. A severe head crash can destroy the read/write head.

Removable high capacity magnetic disks

- High capacity floppy disks
Disks with the same dimensions as 3.3-inch floppy diskette or slightly larger have a much greater capacity than a standard diskette. Popular high capacity floppy disks include
 1. A HiFD - (high capacity floppy drive) made by Sony Electronics Inc. Looks like a normal 3.5-inch disk but can store up to 20 MB. This drive can read and format double density and high-density floppy disks in addition to HiFD disks.

2. SuperDisk – made by Imation, same dimensions as standard 3.5-inch but formatted to hold 120 MB. Drives are compatible with high-density 3.5-inch diskettes.
 3. Zip Disks – made by Iomega Corp. Available in 100 MB and 250 MB capacities and require a special disk drive and utility software. Slightly larger than normal and are not compatible with ordinary floppy.
- Hot-Swappable Hard Disks
At the high end of price and performance. Sometimes used on high-end workstations or servers that require large amounts of storage. They allow the user to remove (swap out) a hard disk and insert (swap in) another while the computer is still on (hot). The removable disk contains the disk, drive and read/write heads in a sealed box.
 - Disk Cartridges
Most removable disks work like a diskette, a disk in a plastic case that is inserted into or removed from the drive. The disk and case are called disk cartridges. Do not offer the same storage as true hard disks, they hold more data than diskettes and are faster and just as portable.

Tape Drives

Read and write data to the surface of a tape the same way as an audiocassette – difference is that a computer tape drive writes digital data.

Best use of tape storage is for data that is not used often, such as backup copies of your hard disk. Because a tape is a long strip of magnetic material, the tape drive has to write the data to it serially – 1 byte after another. This is inherently slower than the direct access provided by media such as disks. When you want to access specific data on a tape the drive has to scan through all the data that you don't need to get to the data that you want. Result → Slow access time. The access time varies depending on the speed of the drive, length of tape and position on the tape to which the head wrote the data in the first place.

Despite the long access time tape drives are well suited for certain purposes – especially backing up a systems entire hard disk. Backing up using high capacity disks or disk cartridges is usually expensive. Many people now use recordable CDs to back up their system. This option is more expensive than tapes, and unless CD-RW are used the data cannot be overwritten. Capacities of tapes can be as high as 100 GB and more, tape offers an inexpensive way to store a lot of data on a single cassette.

Special type of tape drive uses **digital audiotape (DAT)** to achieve high storage capacities. DAT drives typically have multiple read and write heads built into a small wheel or cylinder that spins near the tape at high speed. The tape moves past the heads with a slower speed. The write heads on the spinning wheel each write the data with opposite magnetic polarities on overlapping areas of the tape. Each read head reads only one polarity or the other. The result is a high density per inch of tape. Although DAT tapes are relatively inexpensive the drives can be expensive when compared to standard tape drives.

Optical Storage Devices

The most popular alternative to magnetic storage systems are optical storage media. The most widely used type of optical storage medium is the compact disk (CD), which is used in CD-ROM, DVD-ROM, CDR, CDRW and PhotoCD systems.

Since the mid 1990s nearly all PCs have been sold with a built in CD-ROM drive. Consumers are now buying more systems with DVD-ROM drives rather than standard CD-ROM units. These devices fall into the optical storage category because they store data on a reflective surface so that it can be read by a beam of laser light. A laser uses concentrated, narrow beam of light, focused and directed with lenses, prisms and mirrors. The tight focus of the laser beam is possible because all the light is of the same wavelength.

CD-ROM

The familiar audio compact disk is a popular medium for storing music. In the computer world, the medium is called compact disk read only memory (CD-ROM). This uses the same technology used to produce music CDs.

The CD-ROM drive for music or data reads 0s and 1s from a spinning disk by focusing a laser on the disks surface. Some areas of the disk reflect the laser light into a sensor, other areas scatter the light. A spot that reflects the laser beam is interpreted as a 1 and the absence of a reflection is interpreted as a 0.

Data is laid out on a CD-ROM disk in a long, continuous spiral that starts at the outer edge and winds inwards towards the centre. Data is stored in the form of **lands**, which are flat areas on the metals surface, and **pits**, which are depressions or hollows. A land reflects the laser light into the sensor (a data bit of 1) and a pit scatters the light (a data bit of 0).

On a full CD-ROM the spiral of data stretches almost 3 miles long. A standard CD can store 650 MB of data or about 70 mins of audio.

CD-ROM Speeds

Compared to hard disk drives CD-ROM drives are quite slow, in part because the laser reads pits and lands one bit at a time. Another reason is the rotational speed of the disk. Like a track on a magnetic disk the track of an optical disk is split into sectors. However the sectors are laid out quite differently than they are on magnetic disks.

The sectors near the middle of the CD wrap farther around the disk than those near the edge. For the drive to read each sector in the same amount of time it must spin the disk faster when reading the sectors near the middle and slower when reading the sectors near the edge. Changing the speed of rotation takes time – enough to seriously impair the overall performance of the CD-ROM drive. The first CD-ROM drives read data at 150 KBps and were known as single speed drives. This rate is much slower than that of a typical hard drive which transfers data at rates of 5 – 15 MBps.

Presently CD-ROM drives read data 2 → 52 times faster than first models (300 → 7800 KBps).

Even with the changing speed of the disk reading data from an optical medium is a relatively simple undertaking. Writing data however is another matter. The medium is a foil disk that is physically pitted to reflect or scatter the laser beam. The disk is then covered in a plastic coating and it is difficult to alter the surface of the disk after it has been stamped.

CD-ROM Uses

The fact that you must use special technologies to write data to an optical disk does not mean that this storage medium is not useful. In fact many applications rely on huge volumes of data that rarely change. Because of the high precision and data density possible with CD-ROM, a single CD typically can hold about 650 MB of data.

Software companies can distribute their products on CD-ROM. Because of the high capacity and the fact that 1 CD much cheaper to produce than a set of diskettes, many software publishers regard CDs as the distribution medium of choice.

DVD-ROM

Digital video (or versatile) disk read only memory, is a high-density medium capable of storing a full-length movie on a single disk the size of a CD.

Achieves such high storage capacities by using both sides of the disk and special data compression technologies. The latest generation of DVD-ROM use layers of data tracks; the laser beam reads data from the first layer and then looks through it to read data from the second layer.

Each side of a standard DVD-ROM can hold up to 4.7 GB. Dual layer DVD-ROM can hold 17 GB of data.

CD-R, CD-RW, PhotoCD

CD-R allows you to create your own CD-ROM disks that can be read by any CD-ROM drive. After the information has been written to the CD it cannot be changed.

Using CD-RW drives the user can write and overwrite data onto CDs. With a CD-RW data can be revised in the same manner as a floppy disk.

One popular form of recordable CD is PhotoCD, a standard developed by Kodak for storing digitised photographic images on a CD. Many film-developing stores now have PhotoCD drives that can store your photos and put them onto a CD.

Measuring Drive Performance

Average Access Time

For a storage device the **average access time** (or **seek time**) is the amount of time the device takes to position its read/write heads over any spot in the medium. This measurement is an average because access times can vary greatly depending on the distance between the heads original location and their destination.

For storage devices access times are measured in **milliseconds (ms)**. For memory devices access times are measured in **nanoseconds (ns)**.

In a disk drive the access time depends on a combination of 2 factors.

1. Revolutions per minute (rpm).
2. Time taken to move the heads from one track to another.

<u>Device</u>	<u>Typical Access Time</u>
Static RAM (SRAM)	5 – 15 ns
Dynamic RAM (DRAM)	50 – 70 ns
ROM	55 – 250 ns
Hard disk drives	6 – 12 ms
CD-ROM drives	80 – 800 ms
Tape Drives	10 – 500 s

File Compression

One solution to the storage problem is to compress data. **File compression** or **data compression** is the technology for shrinking the size of a file, thereby freeing up space for more data and programs to reside on the disk.

Entire hard disks, floppy disks, or individual files can be compressed as much as a 3:1 ratio. File compression is performed by software that uses mathematical algorithms to squeeze the data into smaller chunks by removing information that is not vital to the file or data. When the file is returned to its original size this data is reinserted so that the original data is reproduced exactly as it was before compression.

A utility used to compress one or more files shrinks the selected files and stores them inside a new file with its own name. The resulting file is called an **archive file** because it stores the compressed data inside.

Data transfer rate

The speed at which a device can transfer data – how long it takes for one device to transfer data to another device. Speeds are expressed as a rate, or as some amount of data per unit time. When measuring a devices data transfer rate (**throughput**) the time is measured in seconds, but units of data may be measured in bytes, KB, MB or GB.

Speeds for hard disks are generally high, from around 15 MBps to 80 MBps and higher. CD-ROMs and diskettes are the slowest storage devices.

Drive Interface Standards

An important factor in determining how quickly a drive can read or write data is the type of controller that the drive uses. A **disk controller** connects the disk drive to the computers bus enabling the drive to exchange data with other devices.

EIDE

Enhanced integrated drive electronics (EIDE) is an enhanced version of IDE. It is a catch all term for drive interfaces based on this standard. Supports transfer rates of 33 MBps. The standard variants go by many different names including Fast IDE, ATA, Fast ATA and more.

SCSI

Small computer system interface (SCSI) was originally developed as a way to connect third party peripheral devices to mainframe computers. One way to think of SCSI is as an extension to the computers bus. As such all interface circuitry needed by the device has to be on the device itself. One benefit is to improve efficiency, allows higher transfer rates than EIDE.