

# Mass Storage – Disks

We are concerned here with storing large amounts of data (and programs) **permanently**, hence another name for these devices is *long term storage devices*. In relation to RAM, or *primary memory*, they are also sometimes called *secondary memory*.

Many of the devices discussed here can also be attached externally, which is an important option for laptops.

## Magnetic Media

Magnetic media for computers are digital. They save bits of data by altering the polarity of a section on the surface of the medium. This change lasts indefinitely (like magnetizing a pin), so it is considered permanent. Magnetic media includes hard disk drives, tapes, and a number of legacy devices. They also work similarly to the magnetic strip on the back of a credit or ATM card.


## Hard Disk Drives

A hard disk drive (HDD) is a non-volatile, random access device for digital data. It features rotating rigid platters on a motor-driven spindle within a protective enclosure. Data is magnetically read from, and written to, the platter by read/write heads that float on a film of air above the platters.


Introduced by IBM in 1956, hard disk drives have fallen in cost and physical size over the years while dramatically increasing in capacity. Hard disk drives have been the dominant device for secondary storage of data in general purpose computers since the early 1960s. They have maintained this position because advances in their areal recording density have kept pace with the requirements for secondary storage. Today's HDDs operate on high-speed serial interfaces; i.e., serial ATA (SATA) or serial attached SCSI (SAS).

A hard-disk drive is normally located inside the computer case. Additional drives may be added either inside the case or attached externally via a port.

Hard disks come in various speeds. The key specifications to be considered are average *access time* and *data transfer rate*. The **access time** includes both the *seek time* (time to move the arm & head to the correct location) and *latency* (time for the desired location to rotate under the read/write head). For a particular size (0.85"-8") of disk platter, the faster the disk



Hard Disk Drive with labeled parts



Multi-platter Hard Disk Drive

spins (3,000-15,000 RPM), the lower the access (latency) time will be and the faster the data transfer rate.

As shown in the image at right:

- A) platters,
- B) actuator arm,
- C) read/write head,
- D) segment – an arc on the track,
- E) track,
- F) sector.


The data organization is essentially the same as the floppy disk shown below, except that all the tracks that line up one above the other are call a *cylinder*. A *cluster*, not shown here, is two or more contiguous segments that hold data for a single file. (A file might have multiple clusters.)

## Multiple Drives & RAID


Most PCs have a single hard disk installed. A system can be purchased with a second hard disk or the user may install the second as an upgrade. With two disks, one is designated the master and the other the slave, although neither is in control of the other. Both appear as separate hard disks to the system. (Note that this is different from disk partitions which appear as separate disks to the system, but are actually a single physical hard drive.) RAID disks combine disks differently. In this instance, multiple physical disks are present, but they appear to the system to be a single device.

RAID is the abbreviated term for a Redundant Array of Independent Drives. RAID is now used as an umbrella term for computer data storage schemes that can divide and replicate data among multiple physical disk drives. The physical disks are said to be in a RAID array, which is addressed by the operating system *as one single disk*. The different schemes or architectures are named by the word RAID followed by a number (e.g., RAID 0, RAID 1). Each scheme provides a different balance between two key goals: increase data reliability and increase input/output performance.


RAID 0 uses the concept of striping to evenly split data between 2 or more drives. This allows the computer to access



Screenshot showing a single physical disk with multiple partitions



Raid 0




RAID 1

multiple drives simultaneously to increase data transfer rate and response time. The disadvantages of this setup is reliability. If one of the drives fail, all your data is lost. This is because fragments of each file are stored on different disks. With two drives, you are "twice" as likely to experience a failure.

RAID 1, is the setup in which two (or more) disks are used to create a copy, assuring no data loss if one or the other drive fails. Each disk has a full copy. Performance wise, there is no gain.

RAID 2 uses striping and parity bits to correct errors or recover all data when a single disk fails. Successively higher levels use different schemes to increase performance and/or recover from more failures.

Wikimedia Commons  
and By Paul Mullins: constructed



## Tape Drives

A tape drive is a data storage device that reads and writes data on a magnetic tape. It is typically used for offline, archival data storage. Tape media generally has a favorable unit cost and long archival stability.

A tape drive provides sequential access storage, unlike a disk drive, which provides random access storage. A disk drive can move its read/write head(s) to any random part of the disk in a very short amount of time, but a tape drive must spend a considerable amount of time winding tape between reels to read any one particular piece of data. As a result, tape drives have very slow average seek times. Despite the slow seek time, tape drives can stream data to and from tape very quickly. For example, popular Linear Tape-Open drives can reach, as of 2010, continuous data transfer rates of up to 140 MB/s, which is comparable to hard disk drives.



DDS tape drive with 4 types of magnetic data storage media on top, from right hand side: 3.5" floppy disk (1.44 MB), 310m QIC DC tape (250 MB, ~6 mm wide, 1020 ft length), 112m Data8 tape (5 GB, 8 mm, 367 ft), 150m DDS-4 tape (20 GB, 4 mm, 492 ft)

## Optical Media

### Blu-ray

Blu-ray Disc (BD) is an optical disc storage medium designed to supersede the DVD format. The standard physical medium is a 12cm plastic optical disc, the same size as DVDs and CDs. Blu-ray Discs contain 25 GB per layer, with dual layer discs

(50 GB) the norm for feature-length video discs and additional layers possible in the future.

Optical media operates by creating a single spiral track that records data bits (zeroes & ones) as sections of different reflectivity. The sections (or the transition from one to another) represent bits permanently, although the media itself (the disc) is easily damaged. The single spiral is sequential in nature, but the device can easily move radially (unlike a tape), making it a random-access medium. The laser follows a track as the disc spins, sensing differences in reflectivity.

Discs are available in read only (BD), record once (BD-R) and rewritable (BD-RE) formats. Recording to a disc involves using the laser in high power mode to alter the reflectivity of a die on the disk. For rewritable disks, the die can be changed back to its initial state and reused. As indicated above, these dies may be layered with the laser focused on a particular layer to store more data.

Most BD devices are backward compatibles with DVDs and CDs, although you should check the specifications to be sure before purchase.

## DVD

The concepts for storing data on a DVD, or Digital Video Disc, are similar to those of BD. However, the laser used is not capable of focusing as finely, so the data area for each bit is larger (which means less capacity) and DVD players cannot read or write BDs. DVD devices are generally backward compatible with CDs.

The capacity for a DVD-R or DVD-RW is 4.5 GB per layer. They are available as double-sided (two layers) for 9 GB recordable.


DVD-RAM is a rewritable 4.7 GB random access disc used primarily for camcorders. It is not related to Random Access Memory (RAM).

## CD

The CD or Compact Disc, which might be considered a legacy device, is still widely used. It is capable of hold 700 MB (0.7 GB) of data. CD devices cannot read or write either DVD or BDs.



Underside of a DVD-R disc



Schematic representation of a cross-section through a 3D optical storage disc (yellow) along a data track (orange marks).



Inside a DVD player, tray opened—

## #15 How optical storage works (alternate)

**Solid State Drives** A solid-state drive (SSD) is a data storage device that uses solid-state memory to store persistent data with the intention of providing access in the same manner of a traditional hard disk drive. SSDs are distinguished from traditional hard disk drives (HDDs), which are electromechanical devices containing spinning disks and movable read/write heads. SSDs, in contrast, use microchips which retain data in non-volatile memory chips and contain no moving parts. Compared to electromechanical HDDs, SSDs are typically less susceptible to physical shock, are silent, and have lower access time and latency. SSDs use the same bus interface as hard disk drives, thus easily replacing them in most applications.

SSDs gain their performance edge over hard disks because they do not have the lengthy access time of the mechanical device (and no "spin up" time). Speeds for SSD are roughly equivalent to traditional RAM. The lack of moving parts also provides the edge in terms of physical damage and failure possibilities.

The primary disadvantages are that there are a limited number of writes to a flash device and they are not yet as big, in terms of data storage capacity, as hard disks. There is no limit (except for eventual device failure) to the number of writes on a hard disk.

The latest version (Fall 2010) MacBook Air has neither a hard disk drive nor an optical drive.



The disassembled components of a hard disk drive (left) and of the PCB and components of a solid-state drive (right)



Fall 2010 MacBook Air

---

## RAM Disks

RAM-disks are a kind of solid state device, based on volatile memory, such as DRAM. They are characterized by ultrafast data access and are used primarily to accelerate applications that would otherwise be held back by the latency of flash SSDs or traditional HDDs. RAM-disks (DRAM-based SSDs) usually incorporate either an internal battery or an external AC/DC adapter and backup storage systems to ensure data persistence while no power is being supplied to the drive from external sources. If power is lost, the battery provides power while all information is copied from the RAM-disk to backup storage (HDD or SSD). When the power is restored, the information is copied back to the RAM-disk from the back-up storage, and the RAM-disk resumes normal operation (similar to the hibernate function used in modern operating systems).

RAM-disks of this type are usually fitted with DRAM modules of the same type used in regular PCs and servers, which can be swapped out and replaced by larger modules. Like RAM (primary memory),

a RAM-disk does not suffer the limited number of writes problem associated with flash SSD devices.

## Legacy Devices 1

A legacy system is an old method, technology, computer system, or application program that continues to be used, typically because it still functions for the users' needs, even though newer technology is now available. Reasons for not changing to a newer system may include cost, planned retirement (that includes change), and resistance to change from users.

A legacy system may include procedures or terminology which are no longer relevant in the current context, and may hinder or confuse understanding of the methods or technologies used.

The term "legacy" may have little to do with the size or age of the system itself – current mainframes run 64-bit Linux and Java alongside 1960s vintage code.

## Floppy Disks

A "floppy disk" (or diskette) is a form of magnetic data storage; thin, flexible, soft, flat piece of mylar plastic, packaged in a 3.5 inch plastic case. The reason it was called a floppy is because obsolete (8 and 5.25 inch formats) types would "flop" as you waved them. Floppy disks were invented by IBM and were a popular form of data storage from the 1970's to the 1990's. The key to their widespread use was their inexpensive cost and ease of portability. Information could be transferred to a floppy disk, stored, disk removed, then inserted into another system to then be accessed.

Although there was a variety of different sizes of floppy disks produced by various manufacturers, the most widely used was the 3.5 inch. It is a form of secondary "permanent" storage and can hold approximately 1.44 MB. The disk is inserted into a floppy disk drive, a device that holds, spins, reads data from, and writes data to a floppy disk.

With the new systems now available, the demise of the floppy disk was inevitable. It has been replaced by CD, CD-ROM and (USB) flash drives.

The data organization shown to the right is the same as a single platter in a hard disk drive with: a *track* (A), a *sector* (B),



(Legacy) Commodore 64 Game System



3.5" Floppy Disk




5.25" Floppy Disk



8" Floppy Disk

a *segment* or sector of a track (C), and a *cluster* of segments (D).

The floppy disk was renowned, when in common use, as a primary means of spreading viruses from computer to computer. Even when networked computers became more common, many users shared data and programs on floppy disks. A virus attached to a specific program would only be activated when that program was run by the user, however, a boot sector virus required only that the PC be rebooted with an infected disk in place. At that time, it was common for the computer to attempt to boot first from the floppy disk drive (a recovery feature when the hard disk unbootable). Current viruses tend to spread through network enabled programs, such as email.



Disk structure  
*public domain image*

### Zip Disks

A zip drive is a secondary storage device (internal or external) that uses zip disks. Zip disks are cartridges similar to floppy disks but capable of storing 70-500 times more memory. Zip disks are disks with a special high quality magnetic coating that have a capacity of 100, 250, or 750 MB.

Zip drives are almost obsolete amongst today's students as CDs or flash memory sticks are much more convenient since they can be read by almost all computers (zip disks need special zip drives).





Zip Disk

5 1/4" hard drive (CAV) MiniDisc (CAV) 8" floppy disk (CAV)



lomega Zip (CAV)



audio Compact Disc (CLV\*)





Digital Versatile Disc (CLV\*)

Blu-ray Disc (CLV)



12" vinyl record (CAV)

LaserDisc (CLV/CAA/CAV)



Disk storage comparison

Source: [http://cs.sru.edu/~mullins/cpsc100book/module03\\_internalHardware/module03-04\\_internalHardware.html](http://cs.sru.edu/~mullins/cpsc100book/module03_internalHardware/module03-04_internalHardware.html)