


HOW HARD DISK DRIVE WORKS

4/11/2012

Khoder Ahmed Hachem

ASK -PC Academy



A- Summary:

A **hard disk drive (HDD)**; also **hard drive**, **hard disk**, or **disk drive**) is a device for storing and retrieving digital information, primarily computer data. It consists of one or more rigid (hence "hard") rapidly rotating discs (often referred to as platters), coated with magnetic material and with magnetic heads arranged to write data to the surfaces and read it from them.

B- Acknowledgement:

Thanks and Giving

I would like to thank God for blessing , and that helped me to be succeed in the preparation of this issue and modest effort; which I hoped to give benefits from it for all people.

In addition, I send my special thanks to the academic administration and all teachers for help and assistance; which they have given to me all necessary sites for this topic.

C- Introduction:



Hard disks store digital information in a relatively permanent form. Figure (C.1).

Nearly every desktop computer and server in use today contains one or more hard-disk drives. Every mainframe and supercomputer is normally connected to hundreds of them. These billions of hard disks do one thing well -- they store changing digital information in a relatively permanent form. They give computers the ability to remember things when the power goes out. (Figure C.1)

Before the hard drive, magnetic tape on reels spun frantically on mainframe computers. With magnetic disk recording, you can move the read/write head more-or-less directly to where the data is--allowing you random access and a much quicker process than waiting for a thousand feet of tape to spin under the read/write head.

D- Index:

Contents	Pages
Acknowledgement	2
Introduction	2
Analysis of practical and theoretical:	
Definition	4
Magnetic Recording	4-5
Differences between Cassettes	
Tapes (Magnetic Recording) and Hard disks	5
Capacity and Performance	6-7
Inside The Board	7-8
Inside: Platters and Heads	8-9-10
Access and Interfaces	11-12-13
External Removable Drives	13-14
Market Segments	14
Manufacturers and Sales	14-15
The Future of Hard Drives	15
Conclusion	15
References and Biographies	16

E- Analysis of Practical and Theoretical:

1-Definition:

Hard disks were invented in the 1950s. They started as large disks up to 20 inches in diameter holding just a few megabytes. They were originally called "fixed disks" or "Winchesters" (a code name used for a popular IBM product). They later became known as "hard disks" to distinguish them from "floppy disks."

A hard drive is a storage device that rapidly records and reads data represented by a collection of magnetized particles on spinning platters. Its long-term memory-preserving data programs and your operating system even while the machine is asleep or off.

They are classified as non-volatile, random access, digital, magnetic, data storage devices. They have been the dominant device for secondary storage of data in general purpose computers since the early 1960s. Most hard disk drives in the early 1980s were sold to PC end users as an add on subsystem. As the 1980s began, hard disk drives were a rare and very expensive additional feature on personal computers (PCs); however by the late '80s, their cost had been reduced to the point where they were standard on all but the cheapest PC.

2-Magnetic Recording:

Data is written and read as a series of bits, the smallest unit of digital data. Bits are either a 0 or a 1, or on/off state. These bits are represented on a platter's surface by the longitudinal orientation of particles in the magnetically sensitive coating that are changed (written) or recognized (read) by the magnetic field of the read/write head. The drive's firmware adds extra bits to the data that allow the drive to detect and correct random errors.

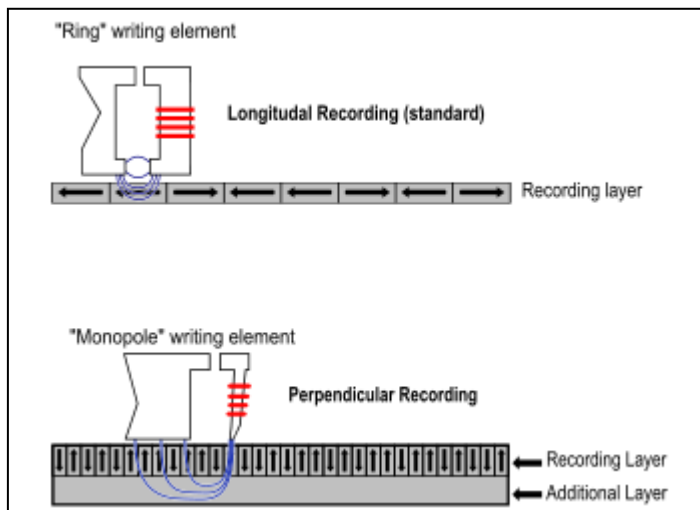


Figure (2.1)

In this type of recording, the particles are arranged perpendicular to the platter's surface. Figure (2.1) in this orientation they can be packed closer together for greater density, with more data per square inch. More bits per inch also mean more data flowing under the read/write head for faster throughput.

Information is written to and read from both sides of the platters using mechanisms mounted on arms that are

moved mechanically back and forth between the center of the platter and its outer rim. This movement is called seeking, and the speed at which it's performed is the seek time. Tracks are divided up into logical

units called sectors. Each sector has its own address (track number plus sector number), which is used to organize and locate data. In the event a drive's read/write head doesn't arrive at the track it's seeking, you may experience what's called latency or rotational delay. This delay occurs before a sector spins underneath the read/write head, and after it reaches the proper track.

In the low-level formatting a drive establishes the tracks and sectors on the platter. The starting and ending points of each sector are written onto the platter. This process prepares the drive to hold blocks of bytes. High-level formatting then writes the file-storage structures, like the file-allocation table, into the sectors. This process prepares the drive to hold files. Modern drives make extensive use of Error Correcting Codes (ECCs), particularly Reed–Solomon error correction. These techniques store extra bits, determined by mathematical formulas, for each block of data; the extra bits allow many errors to be corrected invisibly. But allow higher recording densities to be employed without causing uncorrectable errors, resulting in much larger storage capacity. The perpendicular recording densities may be supplanted by other magnetic recording technologies. Toshiba is promoting bit-patterned recording (BPR), while Xyratex are developing heat-assisted magnetic recording (HAMR).

3- Differences between Cassette tapes (Magnetic Recording) and Hard disks:

At the simplest level, a hard disk is not that different from a cassette tape. Both hard disks and cassette tapes use the same magnetic recording techniques described in How Tape Recorders Work. Hard disks and cassette tapes also share the major benefits of magnetic storage -- the magnetic medium can be easily erased and rewritten, and it will "remember" the magnetic flux patterns stored onto the medium for many years.

The big differences between cassette tapes and hard disks:

- The magnetic recording material on a cassette tape is coated onto a thin plastic strip. In a hard disk, the magnetic recording material is layered onto a high-precision aluminum or glass disk. The hard-disk platter is then polished to mirror-type smoothness.
- With a tape, you have to fast-forward or reverse to get to any particular point on the tape. On a hard disk, you can move to any point on the surface of the disk almost instantly.
- In a cassette-tape, the read/write head touches the tape directly. In a hard disk, the read/write head "flies" over the disk, never actually touching it.
- The tape in a cassette-tape moves over the head at about 2 inches (about 5.08 cm) per second. A hard-disk platter can spin underneath its head at speeds up to 3,000 inches per second (about 170 mph or 272 kph)!
- The information on a hard disk is stored in extremely small magnetic domains compared to a cassette tapes. The size of these domains is made possible by the precision of the platter and the speed of the medium.

Because of these differences, a modern hard disk is able to store an amazing amount of information in a small space. A hard disk can also access any of its information in a fraction of a second.

4- Capacity and Performance:

The capacity of an HDD may appear to the end user to be a different amount than the amount stated by a drive or system manufacturer due to amongst other things, different units of measuring capacity, capacity consumed in formatting the drive for use by an operating system and/or redundancy.

A typical desktop machine will have a hard disk with a capacity of between 10 and 40 [gigabytes](#). Data is stored onto the disk in the form of **files**. A file is simply a named collection of [bytes](#). The bytes might be the [ASCII codes](#) for the characters of a text file, or they could be the instructions of a software application for the computer to execute, or they could be the records of a data base, or they could be the pixel colors for a GIF image. , however, a file is simply a string of bytes. When a program running on the computer requests a file, the hard disk retrieves its bytes and sends them to the [CPU](#) one at a time. There are two ways to measure the performance of a hard disk:

- **Data rate** - The data rate is the number of bytes per second that the drive can deliver to the CPU. Rates between 5 and 40 megabytes per second are common.
- **Seek time** - The seek time is the amount of time between when the CPU requests a file and when the first byte of the file is sent to the CPU. Times between 10 and 20 milliseconds are common.

The other important parameter is the **capacity** of the drive, which is the number of bytes it can hold.

a. Units of storage capacity:

The capacity of hard disk drives is given by manufacturers in [megabytes](#) (1 MB = 1,000,000 bytes), [gigabytes](#) (1 GB = 1,000,000,000 bytes) or [terabytes](#) (1 TB = 1,000,000,000,000 bytes). *Mega-* and *giga-* denote powers of 1000, where prefixes like *kilo-* and *mega-* mean powers of 1024.

- $kilo = 2^{10} = 1024^1 = 1024,$
- $mega = 2^{20} = 1024^2 = 1,048,576,$
- $giga = 2^{30} = 1024^3 = 1,073,741,824,$
- $terra = 2^{40} = 1024^4 = 1,099,511,627,776,$

b. HDD formatting:

Modern HDDs, such as SASⁱ and SATAⁱ drives, appear at their interfaces as a contiguous set of logical blocks; typically 512 bytes long but the industry is in the process of changing to 4,096 byte logical blocks. The process of initializing these logical blocks on the physical disk platters is called *low level formatting* which is usually performed at the factory and is not normally changed in the field. *High level formatting* then writes the [file system](#) structures into selected logical blocks to make the remaining logical blocks available to the host [OS](#) and its applications.^[50] The operating system [file system](#) uses some of the disk space to organize files on the disk, recording their file names and the sequence of disk areas that represent the file. Examples of data structures stored on disk to retrieve files include the [MS DOS file allocation table](#) (FAT) and [UNIX inodes](#), as well as other operating system data structures.

c. Redundancy:

In modern HDDs spare capacity for defect management is not included in the published capacity; however in many early HDDs a certain number of sectors were reserved for spares, thereby reducing capacity available to end users. In some systems, there may be hidden partitions used for system recovery that reduce the capacity available to the end user. For RAID drives, data integrity and fault-tolerance requirements also reduce the realized capacity. For example, a RAID1 drive will be about half the total capacity as a result of data mirroring. For RAID5 drives with x drives you would lose 1/x of your space to parity. Most RAID vendors use some form of [checksums](#) to improve data integrity at the block level. For many vendors, this involves using HDDs with sectors of 520 bytes per sector to contain 512 bytes of user data and 8 checksum bytes or using separate 512 byte sectors for the checksum data.

d. HDD parameters to calculate capacity:

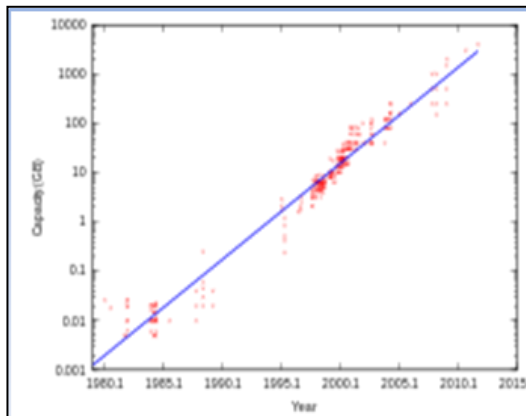


Figure (4.d.1)

PC hard disk drive capacity in [GB](#) over time. The vertical axis is [logarithmic](#), so the fit line corresponds to [exponential growth](#). The gross capacity of older HDDs can be calculated by multiplying for each [zone](#) of the drive the number of [cylinders](#) by the number of heads by the number of [sectors/zone](#) by the number of bytes/sector (most commonly 512) and then summing the totals for all zones. Some modern [ATA](#) drives will, Figure (4.d.1);

Also report [cylinder, head, sector](#) (C/H/S) values to the CPU but they are no longer actual physical parameters since the reported numbers are constrained by historic operating-system interfaces.

The old C/H/S scheme has been replaced by [logical block addressing](#). In some cases, to try to "force-fit" the C/H/S scheme to large-capacity drives, the number of heads was given as 64, although no modern drive has anywhere near 32 platters

5- Inside the Board:



Underneath the board are the connections for the motor that spins the platters, as well as a highly-filtered vent hole that lets internal and external air pressures equalize .Figure (5.1).

Removing the cover from the drive reveals an extremely simple but very precise interior:

Figure (5.1)



Figure (5.2)

In this picture figure (5.2) you can see:

- The **platters** - These typically spin at 3,600 or 7,200 rpm when the drive is operating. These platters are manufactured to amazing tolerances and are mirror-smooth.
- The **arm** - This holds the read/write heads and is controlled by the mechanism in the

upper-left corner. The arm is able to move the heads from the hub to the edge of the drive. The arm and its movement mechanism are extremely light and fast. The arm on a typical hard-disk drive can move from hub to edge and back up to 50 times per second -- it is an amazing thing to watch!

6- Inside the Platters and Heads:

A hard disk uses rigid rotating **platters**. Each platter has a planar magnetic surface on which digital data may be stored. Information is written to the disk by transmitting an electromagnetic flux through an antenna or **read-write head** that is very close to a magnetic material, which in turn changes its polarization due to the flux. The information can be read by a read-write head which senses electrical change as the magnetic fields pass by in close proximity as the platter rotates.

In order to increase the amount of information the drive can store, most hard disks have multiple platters. This drive has three platters and six read/write heads:



The mechanism that moves the arms on a hard disk has to be incredibly fast and precise. It can be constructed using a high-speed linear motor.

Figure (6.1)

Figure (6.1)



Figure (6.2)

Many drives use a "voice coil" approach -- the same technique used to move the cone of a [speaker](#) on your stereo is used to move the arm.

Figure (6.2)

The platters of a hard drive actually hold your data. The image on the left illustrates what a platter looks like after it has been scored by a [head crash](#).

The magnetic surface of each platter is divided into magnetic regions less than a micrometer in size. Each of these regions makes up a single

binary unit of information. Platter surfaces have a mirror finish. The surface of the platter is especially susceptible to finger prints with even the slightest touch. That is why careful handling of a hard drive is critical during the recovery process, and all recovery work must be completed in a clean room environment.

When a hard drive suffers a [physical failure](#), such as a [head crash](#),

When a hard drive suffers a [head crash](#), electronics damage, or spindle motor problem, these are typically classified as physical failures. The most common symptom of a physical failure with the [read-write heads](#) is the dreaded "[Click of Death](#)". Physical failures are some of the most complex problems associated with hard drives, and will nearly double the cost of the recovery when compared to a [logical failure](#). , this could mean we need to swap out the [head](#) assembly.

In the event a spindle motor has failed, we will need to move the [platters](#) from your drive over to the case of another drive, we then perform a [logical recovery](#) from the cloned copy we created. Figure (6.3)



Scored platter due to a [hard drive crash](#). Notice the debris radiating around the cut in the platter. This debris often causes even more damage to the platter. Figure (6.3).

a. Disk Read write Heads :

Disk read/write heads are mechanisms that read data from or write data to disk drives. The heads have gone through a number of changes over the years. In a hard drive, the heads 'fly' above the disk surface with clearance of as little as 3 nanometers. The "flying height" is constantly decreasing to enable higher areal density. The flying height of the head is controlled by the design of an air-bearing etched onto the disk-facing surface of the slider. The role of



Figure (6.a.1)

the airbearing is to maintain the flying height constant as the head moves over the surface of the disk. If the head hits the disk's surface, a catastrophic [head crash](#) can result. Figure (6.a.1).

b. Head Crash :

A head crash is a specific type of [physical hard disk failure](#), and occurs when the read-write head of a hard disk drive touches its rotating platter resulting in catastrophic damage to the magnetic media on the platter surface. Following a head crash, particles of material scraped free of the drive surface greatly increase the chances of further head crashes or damage to the platters.

This damage to the platters is commonly known as "scoring" and data in this area of the platter is completely unrecoverable. In fact, a scored platter is the leading cause of unrecoverable data. There are cases when [platters](#) with very light scoring can be recovered from, but only in instances where the scoring has not affected the system area of the drive, which is vital to the basic functionality of the hard drive.

c. Click of Death :

Click of death is a term that became common in the late 1990s referring to the clicking sound in disk storage systems that signals the device has failed, often catastrophically. The phrase "Click of death" was coined by Mac journalist Tim Robertson, and later publicized by David Pogue. Originally the term was used to describe failing Zip drives, but [hard disk drives](#) also became prone to the symptom. Clicking hard drives are commonly the result of a [mechanical failure](#). However, there are times when this symptom arises due to issues within the firmware, or system area, which prevent the drive from calibrating properly. In [physical failures](#), the most common problem involves the [read-write heads](#).

d. Logical Data Recovery :

Generally speaking when a hard drive is received by a data recovery firm, the first thing they do is evaluate it, and determine what recovery solution will be necessary. If the drive failure is a logical issue then the process starts by creating either an image file of your drive or cloning your drive sector-by-sector another drive. All recovery work is then performed on the image, or cloned copy. After the imaging is complete, a scan will be performed to try and repair the file system corruption. Logical failures usually result in major corruption of the partition and file system. They may or may not be caused by an underlying mechanical issue with the drive. Some logical failures are so bad, that no computer system will even boot with the drive attached. In most cases a logical failure is not the result of any physical problems with the drive. However, there are some instances where file structures become corrupt, or sectors become unrecoverable either due to the [read - write heads](#) being dirty, or [platter](#) coating degradation. Only a data recovery professional will be able to help you determine what exactly has occurred. Even with the drive operating in a degraded state, there are many options we have available that, in most cases, will allow us to fully recover the data without having to take the drive into our clean room for a [physical recovery](#).

7- Access and Interfaces:

Hard disk drives are accessed over one of a number of bus types, including as of 2011 parallel [ATA](#) (PATA, also called IDE or [EIDE](#); described before the introduction of SATA as ATA), [Serial ATA](#) (SATA), [SCSI](#), [Serial Attached SCSI](#) (SAS), and [Fiber Channel](#). Bridge circuitry is sometimes used to connect hard disk drives to buses with which they cannot communicate natively, such as [IEEE 1394](#), [USB](#) and [SCSI](#). Typically, PCs rely on either a PATA (Parallel Advanced Technology Attachment) or SATA (Serial ATA) connection to a hard drive. You might even have both: Most modern motherboards offer both interfaces during the current period of transition from PATA to SATA; this arrangement is likely to continue for some time, as the PATA interface will remain necessary for connecting internal optical drives to the PC. The *parallel* in PATA means that data is sent in parallel down multiple data lines. SATA sends data serially up and down a single twisted pair. For the now obsolete [ST-506](#) interface, the first ST-506 disks used [Modified Frequency Modulation](#) (MFM) encoding, and transferred data at a rate of 5 [megabits](#) per second. Later controllers using 2,7 [RLL](#) (or just "RLL") encoding caused 50% more data to appear under the heads compared to one rotation of an MFM drive, increasing data storage and data transfer rate by 50%, to 7.5 megabits per second.

[Enhanced Small Disk Interface](#) (ESDI) also supported multiple data rates (ESDI disks always used 2, 7 RLL, but at 10, 15 or 20 megabits per second), but this was usually negotiated automatically by the disk drive and controller; most of the time, however, 15 or 20 megabit ESDI disk drives were not downward compatible (i.e. a 15 or 20 megabit disk drive would not run on a 10 megabit controller). PATA drives (also commonly called *IDE* drives). The original ATA interface of the 1980s supported a maximum transfer rate of 8.3MB per second--which was very fast for its time. ATA-2 boosted the maximum throughput to 16.6MBps. Subsequently, Ultra ATA arrived in 33MBps, 66MBps, 100MBps, and 133MBps flavors referred to as Ultra DMA-33 (Direct Memory Access) through Ultra DMA-133 or Ultra ATA-33 through Ultra ATA-133. The odds are overwhelming that you have Ultra ATA-66 or better unless your PC is more than seven years old. (Parallel) ATA interfaces were designed to support two drives on each channel, connected as master and slave on a single cable. When ATA's specification was standardized and the details began to be cleaned up, but still causes problems occasionally, especially with CD-ROM and DVD-ROM disks, and when mixing [Ultra DMA](#) and non-UDMA devices. Recognize by its 2-inch-wide 40-wire or 80-wire cables, though some 40-pin cables are round. Desktop drives typically use a 40-pin connector; the extra wires on 80-wire cables are to physically separate the data wires to prevent crosstalk at ATA-100 and ATA-133 speeds. Notebooks with 2.5-inch drives use a 44-pin connector, and 1.8-inch drives use a 50-pin connector.

- **SATA:**

SATA drives have a much thinner cable and smaller connectors than ATA drives, which allows for more connectors on motherboards and better airflow inside cases. And SATA simplifies setup by using a point-to-point topology, allowing one connection per port and cable. So gone are the jumpers and master/slave connections of PATA drives, where one cable would be used to connect two drives. And unlike PATA, SATA is also suitable for direct-attached external drives, allowing up to 2-meter-long cables on an interface (referred to as external SATA, or eSATA) that's significantly faster than USB 2.0 or FireWire. Serial ATA supports one drive per channel and per cable, with its own set of I/O ports, avoiding master/slave problems.

- **FireWire/IEEE 1394:**

FireWire/IEEE 1394 and USB (1.0/2.0/3.0) hard drives consist of enclosures containing generally ATA or Serial ATA disks with built-in adapters to these external buses.



Figure (7.1)

Several Parallel ATA hard disk drives. Each control cable could serve two or more drives, while a dedicated (and smaller) data cable served each drive:

- ST506 used [MFM](#) (Modified Frequency Modulation) for the data encoding method.
- ST412 was available in either MFM or [RLL](#) (Run Length Limited) encoding variants.
- [Enhanced Small Disk Interface](#) (ESDI) was an industry standard interface similar to ST412 supporting higher data rates between the processor and the disk drive. Figure (7.1).

Modern **bit serial interfaces** connect a hard disk drive to a host bus interface adapter (today typically integrated into the "[south bridge](#)") with one data/control cable.

- [Fiber Channel](#) (FC) is a successor to parallel SCSI interface on enterprise market. It is a serial protocol. In disk drives usually the [Fiber Channel Arbitrated Loop](#) (FC-AL) connection topology is used. FC has much broader usage than mere disk interfaces, and it is the cornerstone of [storage area networks](#) (SANs). Recently other protocols for this field, like [iSCSI](#) and [ATA over Ethernet](#) have been developed as well.
- [Serial Attached SCSI](#) (SAS). The SAS is a new generation serial communication protocol for devices designed to allow for much higher speed data transfers and is compatible with SATA.
- [Integrated Drive Electronics](#) (IDE), later renamed to ATA, with the alias P-ATA or PATA ("parallel ATA") retroactively added upon introduction of the new variant [Serial ATA](#).
- EIDE was an unofficial update (by Western Digital) to the original IDE standard, with the key improvement being the use of [direct memory access](#) (DMA) to transfer data between the disk and the computer without the involvement of the [CPU](#), an improvement later adopted by the official ATA standards.
- [Small Computer System Interface](#) (SCSI), originally named SASI for Shugart Associates System Interface, was an early competitor of ESDI. SCSI disks were standard on servers, workstations, [Commodore Amiga](#), and [Apple Macintosh](#) computers through the mid-1990s, by which time most models had been transitioned to IDE (and later, SATA) family disks.

ATA drives can either use their native CHS parameters (only on very early drives; hard drives made since the early 1990s use zone bit recording, and thus don't have a set number of sectors per track), use a "translated" CHS profile (similar to what SCSI host adapters provide), or run in ATA LBA mode, as specified by ATA-2. To maintain some degree of compatibility with older computers, LBA mode generally has to be requested explicitly by the host computer. ATA drives larger than 8 GiB are always accessed by LBA, due to the 8 GiB limit described above.

Acronym or abbreviation	Meaning	Description
SASI	Shugart Associates System Interface	Historical predecessor to SCSI.
SCSI	Small Computer System Interface	Bus oriented that handles concurrent operations.
SAS	Serial Attached SCSI	Improvement of SCSI uses serial communication instead of parallel.
ST-506	Seagate Technology	Historical Seagate interface.
ST-412	Seagate Technology	Historical Seagate interface (minor improvement over ST-506).
ESDI	Enhanced Small Disk Interface	Historical; backwards compatible with ST-412/506, but faster and more integrated.
ATA (PATA)	Advanced Technology Attachment	Successor to ST-412/506/ESDI by integrating the disk controller completely onto the device. Incapable of concurrent operations.
SATA	Serial ATA	Modification of ATA uses serial communication instead of parallel.

8- External Removable Drives:

[Toshiba](#) 1 [TB](#) 2.5" external [USB 2.0](#) hard disk drive

3.0 [TB](#) 3.5" [Seagate Free Agent Go Flex plug and play](#) external [USB 3.0-compatible](#) drive (left), 750 [GB](#) 3.5" [Seagate Technology push-button](#) external [USB 2.0](#) drive (right), and a 500 [GB](#) 2.5" [generic brand plug and play](#) external [USB 2.0](#) drive (front).

External removable hard disk drives offer independence from [system integration](#), establishing communication via connectivity options, such as [USB](#).

[Plug and play](#) drive functionality offers system compatibility, and features large volume data storage options, but maintains a portable design. These include:

- [Disk cloning](#)
- [Data storage](#)
- [Data recovery](#)

- [Backup](#) of files and information
- Storing and running [virtual machines](#)
- [Scratch disk](#) for [video editing](#) applications and [video recording](#)
- Booting [operating systems](#) (e.g. [Linux](#), [Windows](#), [Windows To Go](#) – a.k.a. [Live USB](#))

External hard disk drives are available in two main sizes (physical size), 2.5" and 3.5".

Features such as biometric security or multiple interfaces are available at a higher cost.

9- Market Segments:

- **Desktop** HDDs typically store between 60 [GB](#) and 4 TB and rotate at 5400 to 10,000 [rpm](#), and have a media transfer rate of 0.5 Gbit/s or higher (1 GB = 10⁹ bytes; 1 Gbit/s = 10⁹ bit/s). As of September 2011, the highest capacity consumer HDDs is 4 [TB](#).
- **Mobile** HDDs or laptop HDDs, smaller than their desktop and enterprise counterparts, tend to be slower and have lower capacity. Mobile HDDs spin at 4200 rpm, 5200 rpm, 5400 rpm, or 7200 rpm, with 5400 rpm being typical.
- **Enterprise** HDDs are typically used with multiple-user computers running [enterprise software](#). Examples are: transaction processing databases, internet infrastructure (email, webserver, e-commerce), scientific computing software, and near line storage management software.

The fastest enterprise HDDs spin at 10,000 or 15,000 rpm, and can achieve sequential media transfer speeds above 1.6 Gbit/s.

- **Consumer electronics** HDDs include drives embedded into [digital video recorders](#) and [automotive vehicles](#).

10- Manufacturers and Sales:

More than 200 companies have manufactured hard disk drives over time.

Worldwide revenue from shipments of HDDs is expected to reach \$27.7 billion in 2010, up 18.4% from \$23.4 billion in 2009 corresponding to a 2010 unit shipment forecast of 674.6 million compared to 549.5 million units in 2009. As of 2011, most hard drives are made by:

Manufacturer \ Market segment	Mobile	Desktop	Enterprise	Consumer electronics	Total market share in units shipped
Western Digital	30%				31.2%
Seagate	20%				29.2%
Hitachi GST ^[note 1]	19%				18.1%
Toshiba	17%				10.8%
Samsung ^[note 2]	14%				10.7%
Units shipped 4Q10 (millions)	72.70	64.06	8.16	22.39	167.31

1. ^ In the process of being purchased by Western Digital.
2. ^ Since been purchased by Seagate.

11- The Future of Hard Drives:

Eventually, all desktop and mobile hard drives will use the SATA interface and perpendicular magnetic recording. Any new PC you look for should have a SATA interface at least; you can upgrade to a perpendicular drive later when prices fall. Expect capacities to continue to grow exponentially, and for performance to grow moderately. Read "[The Hard Drive Turns 50](#)" for a look at where hard drives have been, and where they're going. *Longtime PC World contributor Jon L. Jacobi is a freelance writer based in San Francisco.*

F- Conclusion:

The hard disks drives are one from the principal features based in the computer.it consist for many types as: PATA P(ATA), SATA, eSATA, EIDE, IDE and/or Fiber channel. In addition: inside the hard drive there are three important functions: Platters, arms, and heads. Finally; the hard disk drive is dominate all information (data, information ...) stored in it as read/write from and to it as digital information. It's one from the requirements or /and basic devices in the computer or in the PC.

G- References & Biographies:

1. <http://computer.howstuffworks.com/hard-disk.htm> , by Brain, "**How Hard Disks Work**".
 2. http://www.pcworld.com/article/18693/how_it_works_hard_drives.html, By Jon L. Jacobi, PC World Sep 13, 2006 10:00 am, "**How It Works: Hard Drives**". How A Hard Drive Works
 3. <http://www.acsdata.com/how-a-hard-drive-works.htm>, "**How a Hard Drive Works**".
 4. http://en.wikipedia.org/wiki/Hard_disk_drive, From Wikipedia, the free encyclopedia: "**Hard disk drive**".
-
-