HARD DRIVE Study Guide

DATA PROCESSING AND THE FIRST STORAGE DEVICE

Early computers did not incorporate HDD's for storage, because they were not invented yet. In fact, they did not even use floppy disks for storing data. During the 1950's - 1960's they used paper tape, and then paper cards. Each time you stored information in a different format, you had to manually program the computer to do what was needed. You did not have pre-loaded software programs to do the job for you --- not even on a floppy disk. Imagine having to pull out a panel and wire it each time you wanted to save information or data in a different format. As the computer operator typed in the data or information to be stored, the **IBM 024 80-Column Card Punch** computer machine would punch holes into cards that were fed from a stacked tray on the right. It worked similar to a typewriter, except instead of the machine typing letters as you hit the keys, it would punch holes. The cards were fed from right-to-left in front of your keyboard on the machine, so that you could actually see what was happening. When the card was filled, it would end up in another stacked tray on the left. A second computer operator would then gather the stack of cards and retype the same data or information on the exact same cards on another machine called the **IBM Card Verifier**, to make sure that there were no errors. If there were errors, the cards were sent back to the first computer operator and he or she had to redo the work. You obviously could not erase mistakes, and the high cost of operating the final mainframe computer did not allow for errors.

Keypunch Machine Computer

Keypunch Machine Card Verifier Machine

Card Sorter

From the COMPUTER MUSEUM OF AMERICA

| |

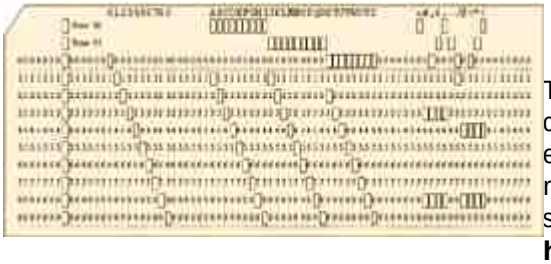

Keypunch Card

These cards would then be picked up to be sorted in another computer machine called the **IBM Card Sorter**, and then eventually read so that the information could be stored on reels of tape or printed. You could not just click a mouse to sort by date, name, size, etc. **To make it even harder, you had to manually wire special boards on these devices to tell it what to store and how to categorize the data.** The cards were then stored in boxes on shelves until they were needed again. Considering that there were only 80 columns and 12 rows in the cards, you can imagine the cumbersome task it presented in storing information for later use or retrieval, and the amount of cards that it required.

Interestingly, the first computer monitors were also designed to be 80 columns across, where only 80 letters would appear on the computer screen before it carried over to the next line. This coincided with the same amount of columns as the keypunch cards. **Currently, the default DOS prompt screen still only uses 80 letters across.** It has been quite a few years since I operated those machines, but they were computers, and it was more efficient than a typewriter and paper for storing, computing and

presenting information. Compared to today, it was more like assembly line computing. Instead of all the components and devices sitting inside of a desktop or tower, they were all at separate stations.

Actually, the first card punch computing system and machine was invented by **Herman Hollerith** around 1880. **This computing system was used to gather, compute, and tabulate data for the American Census in 1890.** The name "**Herman Hollerith**" probably does not ring a bell for you, but he is the main person responsible for the computers and hard drives that you use today. He eventually renamed his automatic tabulating/computing machine company in 1924 to **IBM** (*International Business Machines*). IBM THROUGH THE YEARS

Since paper storage has its limits and deteriorates from age, moisture, handling, etc., a new way to store information became necessary. After all, when those cards became old they would have to be remade all over again. Additionally, the users, and this new system of storing data with the help of computers grew, and so did the amount of information being stored. The use of keypunch computing eventually phased out, and the computers were then attached directly to the data entry stations.

| |

IBM experimented with many variations of what eventually became the first HDD. The very first one, even though it used magnetic storage similar to what we use today, did not use a flat disk. It used a metal drum shaped like a can. However, it was extremely large and hard to work with.

Eventually, IBM invented the first true disk storage drive, which contained 50 flat disks. Each disk was 24 inches in diameter. That means its form factor was 24 inches. Put two 12 inch rulers side by side and you will get an idea of how wide the disks were. Pile 50 of those disks on top of each other and you can begin to appreciate the enormous size of this HDD compared to those used today. Now consider the fact that this huge monster hard drive only could store a total of 5 megabytes of information or data. The name of this HDD was the **IBM 305 RAMAC** (*Random Access Method of Accounting and Control*), and was unveiled in 1956. IBM would not sell the RAMAC hard disk drive. **However, you could lease one of colossal giants for $35,000 a year. To appreciate how much the 5 meg hard drive cost to lease in 1956, understand that a new Cadillac could be bought for $4000 and a new Rolls Royce for $10,000.** (*The RAMAC hard drive is the large box that looks like a cabinet on the right side of the picture.)*

We have come a long way since those days, and worked down through the form factors of 5.25", 3.5", 2.5", 1.8", 1.3", and it continues to grow smaller. Additionally, the amount of storage space has increased from a mere 5 megabytes to the incredible capacity now measured in terabytes. **As the old saying goes, "We've come a long way baby", and we owe it all to Herman Hollerith of IBM.**

| |

THE MODERN DAY HARD DISK DRIVE

The HDD or hard disk drive is a hardware device used as the storage medium for the computer. It is used to store data, information, programs, images, and various other files permanently. This means that it stays on the disk after the computer is shut off, and can be accessed over and over, until or unless you decide to delete the files on the HDD. The files are written and stored on individual platters that reside

inside of the HDD case. These platters are covered with a special magnetic material on both sides. The platters of the HDD spin at very high speeds. These speeds are measured in **RPM** *(revolutions per minute - how many times an object spins in a complete revolution in one minute)*. HDD's spin at thousands of revolutions per minute. To store files on the HDD, the read/write heads, which are located on the actuator arm, send small magnetic charges onto the surface of the platter. Look at the picture and diagram below from PC Guide to see a close up of the parts inside a hard drive.

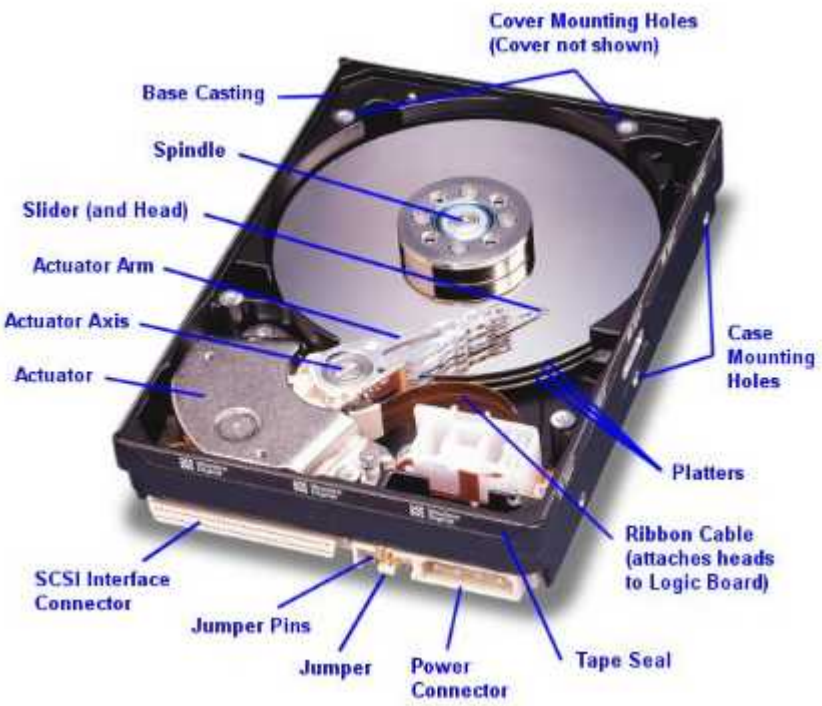

For a larger view CLICK on the HDD image above.

| |

Locate this list in the image above:

- **CASE & MOUNTING HOLES** (where the cover is attached to the case and also where the HDD is securely attached to the HDD bay in the computer case.)
- **POWER CONNECTOR** (where the power connector and cable attach to the power supply for powering the HDD.)
- **JUMPER PINS & JUMPER** (for setting the **MASTER**, **SLAVE**, **CABLE SELECT**, **SINGLE** , etc. settings for the hard drive.)
- **INTERFACE CONNECTOR** . (where the HDD connects to the motherboard or card controller with a flat ribbon cable.) In the picture above, an SCSI connector is shown. **In an IDE/EIDE HDD, there would be a connector similar, but with a different amount of connector pins and settings.**
- **PLATTERS** (both sides of each platter are used for storage of data.)
- **SPINDLE** (located in the center of all the platters, and is used to spin the platters.)
- **ACTUATOR** (a mechanical device for moving or controlling the actuator arm.)
- **ACTUATOR AXIS** (sits in the center of the actuator enabling it to move smoothly.)
- **ACTUATOR ARM** (a mechanical arm enabling the HDD read/write heads to extend over & in-between the platters for reading and writing to the HDD.)
- **SLIDER & HEAD** (located at the end of the actuator arm and is the mechanism used to read and write to the platters. A spindle motor spins the spindle - not shown above.)
- **INTERNAL RIBBON CABLE** (attaches heads to the logic board of the HDD.)

The **HDD cover** is not shown in the picture above. However, it fits securely over the case, and is secured by screws. It is sealed to prevent contaminants from entering inside the HDD, and a special seal (usually tape) is secured to where the case and cover join. In a sense, this seal acts like a gasket. This special seal is in place so that the company can tell if the HDD has been tampered with in case the HDD is still under warranty and needs to be replaced.

| |

CONTROLLER

The **CONTROLLER** is located on the motherboard, or sometimes (especially in very old HDD's) on a card that fits in an available adapter slot. Newer hard drives have a controller built into them, and are referred to as an onboard controller. The HDD is connected to the motherboard with a flat ribbon interface cable. As in most computer parts names, the definition is in its name. Therefore, the controller is the in-between part that controls the information, etc. going in and out of the HDD and the rest of the computer.

| |

BIOS

The **BIOS** of the system is the lowest-level interface between the HDD and the complete system . The control of access to hard disks is determined by several important factors. These are:

BIOS Interrupt Routines - provides a way for the software and HDD to communicate. Newer OS's bypass the BIOS and control the transfer of data and instructions within the operating system. However, the BIOS is still sometimes used for compatibility issues.

Hard Disk Detection and Configuration - The BIOS is used to configure *(automatically or manually)* the HDD, so that the rest of the computer knows how to work with the HDD. If the BIOS is too old to configure the HDD *(due to limitations in BIOS recognizing the size and special abilities of the HDD)* then the HDD manufacturers provide special software, such as Western Digital's Data Lifeguard Tools , to overcome these barriers.

Hard Disk Interface Mode Support - The BIOS works with the system chipset and I/O bus to control the different interface modes that the HDD can use. These modes are:

- **PIO MODE** - (Programmed in/out mode) The transfer of data is performed by the CPU (central processing unit). The CPU is handles the instructions that are executed for transferring the data to and from the hard drive. There are 5 mode channels available in PIO MODE (0, 1,2, 3 & 4). (see: http://www.pcguide.com/ref/hdd/if/ide/modes_PIO.htm)
- **DMA MODE -** (Direct Memory Access Mode) A hard drive that is capable of operating in DMA MODE transfers instructions and data without using the processor. This mode enables the hard disk to transfer data directly to memory. (see: http://www.pcguide.com/ref/hdd/if/ide/modes_DMA.htm)
- **BLOCK MODE** - Block mode is a performance enhancement that allows the grouping of multiple read or write commands over the IDE/ATA interface so that they can be handled on a single interrupt. (see: http://www.pcguide.com/ref/hdd/if/ide/modes_Block.htm)

| |

Int13h Interface - the interface or software interrupt that supports most commands given to the HDD. It had limitations and was replaced by the Int13h Extensions. However, it is still sometimes used by older systems and for compatibility issues. **Supported HDD's with a maximum size of 8.46 GB**

and a 24-bit wide path/interface.

(see: http://www.pcguide.com/ref/hdd/bios/biosInt13h-c.html)

Int13h Extensions - developed to overcome the traffic congestion of data/instructions trying to pass through the older 24-bit wide interface, and to allow for larger HDD's. **It allowed for larger HDD's of up to 9.4 TRILLION GIGABYTES to operate and incorporated a 64-bit wide path/interface.**

(see: http://www.pcguide.com/ref/hdd/bios/biosExtensions-c.html)

Direct Disk Access (Bypassing the BIOS) - Since the BIOS sits in the middle of the HDD and the rest of the computer, it has a tendency to slow down access rates and times, and therefore affects the peak performance of the HDD. To overcome this problem, Direct Disk Access was employed. Modern OS's (operating systems) now bypass the BIOS and provide 32-bit protected mode data access to and from the HDD. Also read Designing for 64-bit Windows at Microsoft.com http://www.microsoft.com/hwdev/64bitwindows/64bitsystems.htm

(see: http://www.pcguide.com/ref/hdd/bios/biosDirect-c.html)

| |

PHYSICAL BREAKDOWN OF HDD PLATTERS

TRACKS / SECTORS / CLUSTERS / CYLINDERS

To understand how a hard disk drive platter is divided up, it is easier to first compare it to a library filling up an entire wall. Picture each of the shelves containing rows and rows of books. Now picture those books being sets of encyclopedias, handyman books, dictionaries, etc. The very bottom shelves are left empty on purpose, so that you have a work area to place the books when you read them. (Keep in mind that when you first set up a hard drive, all the files are in order, just like the sets of books are on the shelves.)

- **Encyclopedias are lined up and grouped in order by letters A-Z.**
- **Handyman books are grouped in order by numbers 1-20.**
- **Dictionaries are grouped alphabetically in order by author.**

When you go to the library shelves to get an encyclopedia book on BEARS, it would be easy to find. All you would have to do is locate the encyclopedia set, and you would pull down the second book with the letter **"B"**, since **BEARS** fall under **"B " words**. When you are finished with the book, you would put it back where you found it. Right? Otherwise, it would be harder to find the next time you needed that particular book.

Imagine that you were not careful and just put the books you found in any open area on the shelves that the book would fit into *(but not where you found it)*. Then imagine doing this every time you pulled a new book out of the library. Now lets say that you continued this untidy pattern over a period of time until you have pulled down every book in the library. Well, all I can say is that I sure wouldn't want to come over to your house to do a homework assignment. I would spend more time looking for the books than I would working on the homework.

Hard disk drives are accessed very similar to this. However, it is not a very good librarian. When files stored on the hard drive are used, they are pulled out of their storage place and brought down to the empty storage spaces to work with. To operate a game, you might need a long list of files just to start the game. Each time you did something new in the game you would need more files, and each time

you saved your game you would be adding even more files to the hard drive. When you are finished playing the game, the files are kind of thrown back into the hard drive in any open space it finds, not always where they came from. After a while it becomes harder and harder for the hard drive to gather all those files together when you want to play the game again, because it cannot just go to one area of the hard drive. It must seek and search the entire set of PLATTERS, TRACKS, SECTOR AND CLUSTERS to get to them. Even though a RECORD is kept of where their location is, physically it takes longer and longer to locate them, because the hard drive platters have become very **DEFRAGMENTED**. Additionally, they may have become **CORRUPT** and will show up as errors, making them unreadable and unusable. All the extra work that the hard drive has to do can help it grow old and useless a lot sooner than it should. To check for errors you need to regularly use the **SCANDISK** tool, and to defragment hard drives you need to use the **DISK DEFRAGMENTER** tool. These can both be found in your operating system, along with instructions on how to use them. **Now that you understand this overly simple explanation of the storage of files on a hard drive, lets go on to identifying how it is really broken down into individual areas. Examine the diagram below.**

| |

PLATTERS

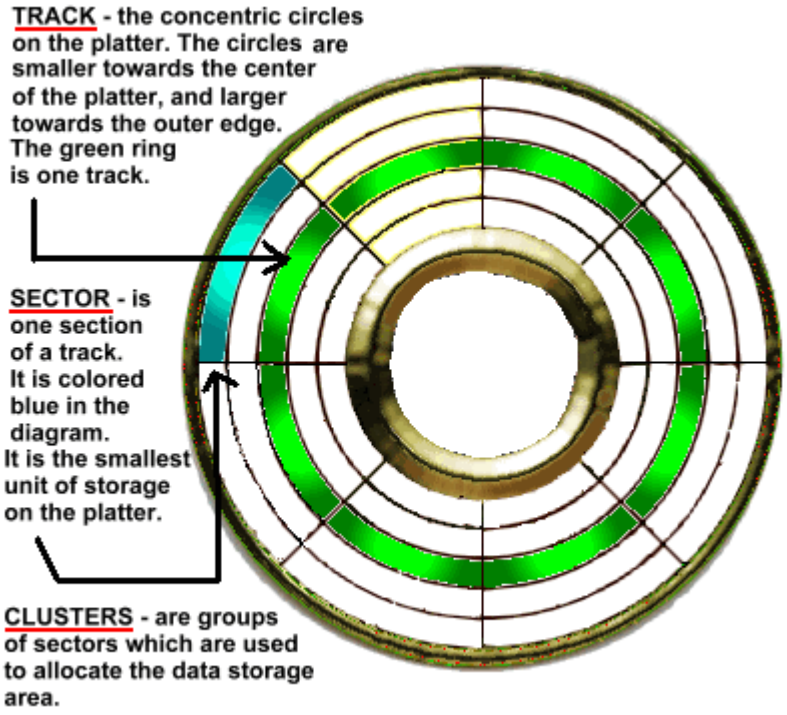

| |

TRACKS & SECTORS - The diagram above is basic so that you can understand the divisions easier. Hard disk drive platters are divided up into individual areas consisting of **TRACKS** & **SECTORS**. Look at the diagram above to understand the difference between a track and sector. You see that a **TRACK** is one complete ring around the platter, and that a **SECTOR** is just one part of a track. You can also see that there are many sectors inside of one track. Each track in the diagram above has 16 sectors. **(In actuality, there are thousands of tracks on each platter, and they are nowhere near as wide as in the diagram.)**

CLUSTERS - Additionally, a **CLUSTER** is a group of sectors. Files are stored in clusters. This means

that a single file can be stored in many sectors, because the file might be too large to fit in just one sector. When an HDD becomes fragmented, a single file can become scattered all over the platters. When the request for the file is sent, the HDD will start thrashing all over the platters trying to get all the parts of the file if the HDD is badly defragmented.

CYLINDERS - Understanding **CYLINDERS** is a bit harder. Imagine that there are 6 platters just like the one above in the diagram. Each platter is slid onto the SPINDLE so that they are all evenly stacked on top of each other with some space in-between each platter. Now imagine that each platter is identical and has one green ring on it (a track) just like in the diagram above. When the platters are stacked, this means that each green ring on each platter is located directly under or above each other. Imagine drawing an imaginary line from each section, around each track (the inner and outer circle of the track) on one platter, and connecting it to the same green track on the next platter underneath it. Then do the same to each of the other platters until all of the green tracks are connected together. If it were possible to take a can with the lid and bottom cut off, and somehow slip it through the platters to fit only in the space where the green tracks are, you would be inside of the **CYLINDER**. It is called a cylinder because of the imaginary cylindrical shape that it forms.

HEADS - Every hard drive has **2 read/write heads for each platter in the drive** . There is one head for each side of each platter. The head is attached to an **actuator arm**. It works similar to the same way that the arm on a phonograph record player acts, except on a record player the phonograph needle touches the phonograph records to read it. The heads **NEVER** touch the platters. If they did, the heads and the data on the platters would be destroyed and cause what is called a **"head crash"**, rendering the hard drive completely useless. The heads are used to read & write data to and from the hard disk platters using a magnetic process similar to the process of playing and recording music tapes.

| |

SPINDLE & MOTOR - A hard drive's motor, which rotates the platter (or platters) counter-clockwise at speeds measured in revolutions per minute (**RPM**). The speed varies depending on the individual hard drive and manufacturer. Spindles spin continuously unless the hard drive is set up to go into a sleep mode when it is inactive or idle (not being used).

ELECTRONIC BOARD - The hard drive is further controlled by electronics that are all soldered or otherwise connected on a small pc board about the size of an adapter card. It fits on top of the case, directly above the mechanical hard drive components. The connectors and jumpers on the back of the hard drive are soldered onto this board.

DATA INTERFACE CONNECTOR - This is the connector on the back of the HDD where the flat ribbon interface cable connects the HDD to the motherboard or the controller card. Remember from a previous lesson on CABLES & CONNECTORS , that the side with **the red stripe on the cable indicates that it needs to be connected to the lowest number pin on the board connector** (1 or 2 will usually be seen on the board for identification purposes). *(Older HDD's which used 20-pin and 34-pin connectors will not be discussed here, because they are no longer being used.)*

- **IDE hard drives** = 40-pin connector
- **Ultra ATA/33** - 40-pin connector
- **Ultra ATA/66** - 40-pin, 80-conductor IDE cable **(REQUIRES an Ultra ATA/66 compatible motherboard or controller)** *(Ultra ATA/66 hard drives are 100 percent backwards compatible with Ultra ATA & DMA.)*
- **Ultra ATA 100 Interface** - 40-pin, 80-conductor IDE cable **(REQUIRES an Ultra ATA/66 compatible motherboard or controller)** *(Ultra ATA/100 hard drives are 100 percent backwards compatible with Ultra ATA/66, Ultra ATA/33 and DMA, and with existing EIDE/IDE hard drives, CD-ROM*

drives and host systems.)

- **SCSI hard drives** = 50-pin, 68-pin or 80-pin connectors

LED CONNECTOR - (light emitting diode) Older hard drives had a face plate with a light connected to the front of it. When the computer was on you could see the light blink when the HDD was in use. The light was connected to the board a via 2-wire cable. Modern computers still use the LED, but they are now mounted on the computer case. The LED is connected directly to the motherboard on a special connector which still uses 2 wires. The motherboard manual will show where these wires are supposed to be connected.

| |

MORE ON HARD DISK DRIVES

FAT - FILE ALLOCATION TABLE - The **FAT** (hidden on a hard drive) is similar to a **TABLE OF CONTENTS** in a book. It is located on a special section of the hard drive. The OS (operating system) uses this table to locate files on a disk. Remember how I explained that files are stored in clusters of sectors, and that one file could be divided up into many sectors? When the hard drive becomes very fragmented, it becomes harder for the hard drive to work at locating the pieces of these files that are scattered all over the platters, but FAT still keeps an accurate record of where each piece is located. If the FAT becomes corrupt, the files could not be located. The FAT has changed with operating systems and so has its name.

- **DOS = FAT**
- **WINDOWS 95 = VFAT**
- **WINDOWS 98 & ABOVE = FAT32**

SMART DRIVES - Modern hard drives incorporate a system called a **"Self-Monitoring, Analysis, and Reporting Technology"**. This system was developed and incorporated into hard drives to monitor and test the **HDD** for signs that might lead to its failure. If the **SMART** system finds anything wrong, it notifies you with a warning when you boot up your computer. However, in order for this system to work you need to make sure that it is activated in your CMOS settings. Speaking of **CMOS** , the computer needs to know how many tracks, cylinders, heads and sectors there are, and where each of these are located. This information is also stored in **CMOS**, and needs to be accurate for the hard drive to work. The settings can usually be found on the cover plate of the hard disk drive. If it isn't located there, look in the manual that came with the hard drive. Otherwise, you will have to find the information at the manufacturers web site.

MTBF - **Mean Time Between Failure** - Hard drive manufacturers test the hard drives they make for reliability and the predicted life of a hard drive If the **MTBF** was 100,000 hours this means that the hard drive has been tested and found to run at least 100,000 hours before it failed or encountered disk errors.

SEEK TIME & FULL STROKE SEEK - Seek time is the is the amount of time required for the hard drives read/write heads to move between tracks over the surfaces of the platters. It is further broken up into 3 different measurements. A fast seek time is a desired and deciding factor when deciding on which hard drive to purchase for your computer. **However, you must take into considerations that a nearly empty hard disk drive will run much faster than one which is nearly full.**

- **Track-to-Track Seek Time** - The amount of time it takes the read/write heads to move from track to track.
- **Average Seek Time** - Since tracks vary in size depending where on the platter they are located

(towards the outer edge or close to the inner edge - by the spindle) they seek time between tracks is different. Therefore measurements are made randomly of all the tracks to find the Average Seek Time.

- **Full Stroke Seek Time** - This is the time it takes for the read/write heads to seek and read the entire disk platter, starting from the outer track to the innermost track, and all the tracks in-between. **If you open an old hard drive** *(one you no longer are able to use)* **and power it up, you can watch the actuator arm and heads complete a Full Stroke Seek when the hard drive is first turned on.**

DATA TRANSFER RATE - The data transfer rate is measured in **MBps** (megabytes per second), and is the measurement of data transfer between the HDD and the CPU. A faster data transfer means the user has to wait less time for a software program to load and operate.

CACHE - the cache, which is also called the **BUFFER**, is a temporary storage area for all newly created data. The data is held in the cache until the hard drive is able to access, rewrite and save data to the hard disk.

| |

HARD DRIVE TYPES

IDE - **Integrated Drive Electronics** - put most of the electronics on the drive itself, therefore it controls itself.

EIDE - **Enhanced Integrated Drive Electronics** - improved on the limitations of the IDE drive and allows up to 4 drives on one controller, including non hard disk drives such as CD-ROM drives.

SCSI DRIVES - (**Small Computer System Interface drives**) Also pronounced as scuzzy drives, they allow the addition of up to 7 devices, scanners or printers in one line, as long as they use the SCSI interface. They are much faster than **IDE/ATA** drives, but are also more expensive. These drives are connected in a row to each other in a daisy-chain fashion. Each SCSI device on the bus requires a unique SCSI number from 1-7, with the computer's SCSI number being 0. Included in SCSI Drives are the following types (also see DATA INTERFACE CONNECTOR above for information on connector types):

- **SCSI 1**
- **SCSI 2**
- **SCSI 3**

EXTERNAL HARD DRIVES - External hard drives are connected to the computer outside of the case via a connector interface on the back of the computer. An example is the QUICK DRIVE which connects to the parallel port connector, and is a way to add extra hard drives for storage, but does not run as fast as an internal hard drive. Other QUICK DRIVES connect via the USB or PC CARD such as those used on laptops. A newer type of external hard drive by Western Digital are called Fire Wire External Hard Drives . The **IEEE FIREWIRE** peripheral interface (also known as 1394 and iLINK) has 30 times more bandwidth than USB, making it the perfect technology for high-speed devices like the latest hard drives, Digital Video (DV) camcorders, CD-RW drives, printers, and scanners.

| |

MEASUREMENT OF HARD DRIVE CAPACITIES

Modern Hard Drives capacities are measured in megabytes and gigabytes. Although computer technology has climbed to capacities as large as **terabytes**, it is currently not available in the market for the everyday user.

One important note to keep in mind is that most hard drive manufacturers list hard drives of

gigabyte and megabyte capacities by using the 1000 multiplier rather than the true multiplier of 1024 in deciding the capacity of the HDD. For instance, let's say you go to the store and buy a hard drive that is listed as 8 GB (gigabyte). After you install it and access the property sheet for your computer and hard drive in Windows, you notice that you are missing some storage space on the hard drive. Instead of it showing the HDD as having 8 GB of space, it shows 7. something gigabytes. In reality , you are not missing any space on the HDD, and there is nothing wrong with it. The problem, as previously stated, is the incorrect way that manufacturers compute/list the capacity.

A quick refresher on bits and bytes:

- 8 bits = 1 byte
- 1024 bytes = 1 kilobyte (KB)
- 1024 kilobytes = 1 megabyte (MB)
- 1024 megabytes = 1 gigabyte (GB)
- 1024 gigabytes = 1 terabyte (TB)

Using the chart above, change the multiplier of 1024 to 1000, and you will understand how the HDD manufacturers are incorrectly stating the size of the HDD.

| |

LINKS TO EXPLORE & LEARN MORE

- Hard Disk Drive- Webopedia.com
- Hard Drive Related Terms and Tricks
- PC Guide - Reference Guide - Hard Disk Drives
- TheRef (tm) Drive and Controller Guide
- Howstuffworks How Hard Disks Work
- PC Tech Guide - Hard Disk Glossary of Terms - HARD DISKS
- IBM Archives 1956
- PC Mechanic - Hard Drives / Hard Drive, Heal Thyself / Common Hard Drive Problems

| |