

Hard facts

Text by PowerQuest, photos and drawings Vlado Damjanovski © 2004

What is a hard disk?

A hard disk or hard drive is the part of your computer responsible for long-term storage of information. Unlike volatile memory (often referred to as RAM) which loses its stored information once its power supply is shut off, a hard disk stores information permanently, allowing you to save programs, files, and other data. Hard disks also have much greater storage capacities than RAM; in fact, current hard disks may contain over 250 GB of storage space.

Basic components of a hard disk

A hard disk is comprised of four basic parts: platters, a spindle, read/write heads and integrated electronics.

- Platters are rigid disks made of metal or plastic. Both sides of each platter are covered with a thin layer of iron oxide or other magnetizable material.

- The platters are mounted on a central axle or spindle, which rotates all the platters at the same speed.

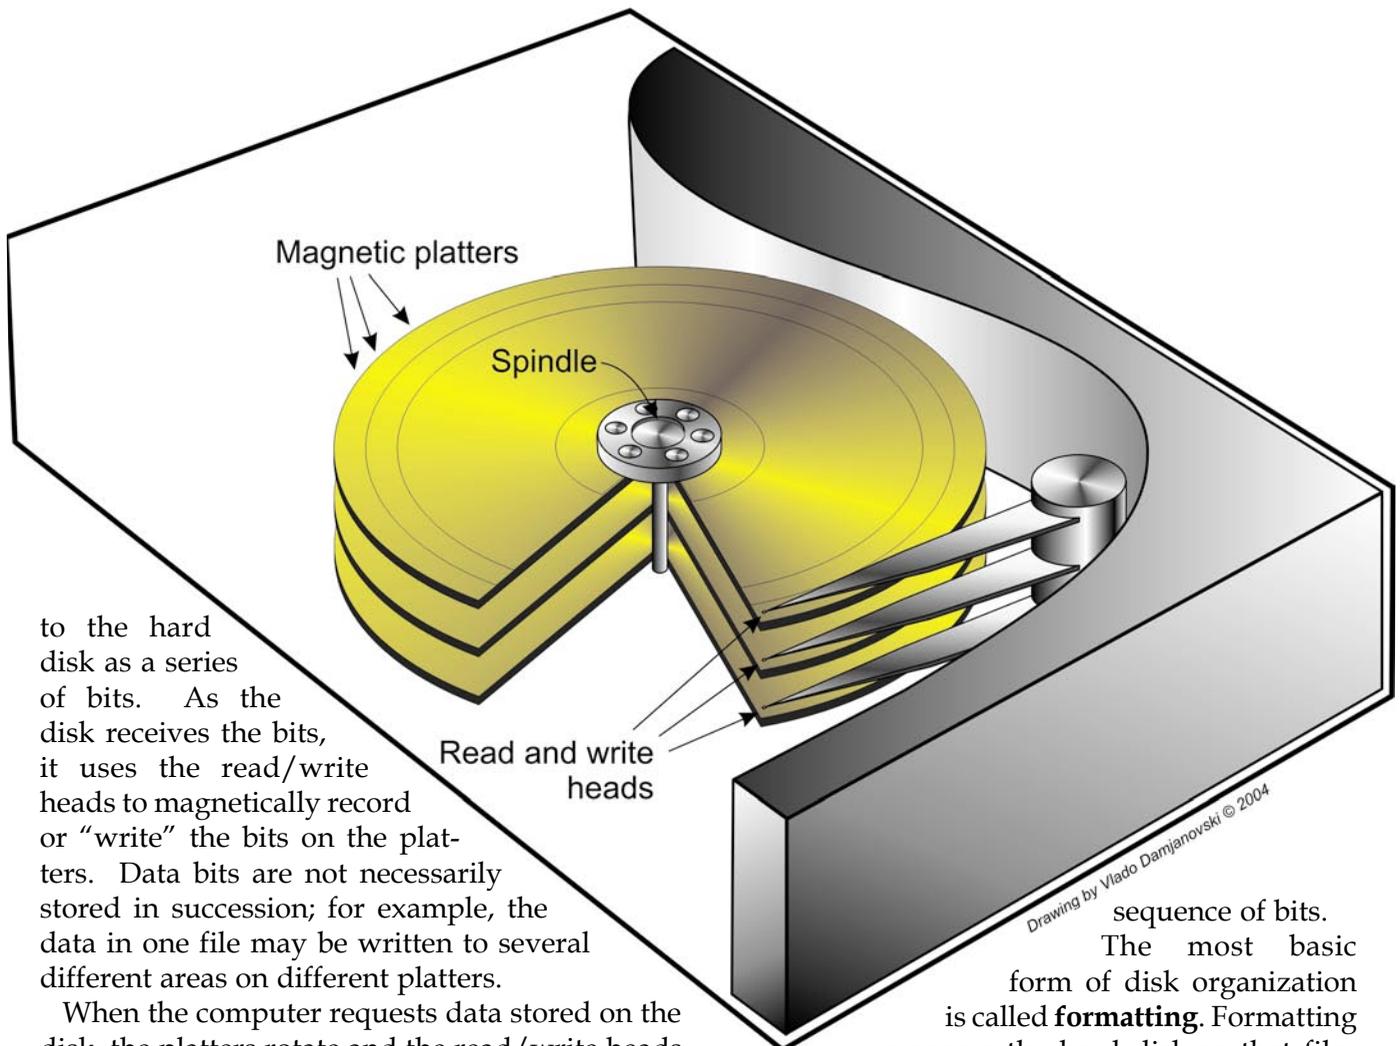
- Read/write heads are mounted on arms that extend over both top and bottom surfaces of each disk. There is at least one read/write head for each side of each platter. The arms jointly move back and forth between the platters' centers and outside edges; this movement, along with the platters' rotation, allows the read/write heads to access all areas of the platters.

- The integrated electronics translate commands from the computer and move the read/write heads to specific areas of the platters, thus reading and/or writing the needed data.

How is data stored and retrieved?

Computers record data on hard disks as a series of binary bits. Each bit is stored as a magnetic charge (positive or negative) on the oxide coating of a disk platter.

When a computer saves data, it sends the data



to the hard disk as a series of bits. As the disk receives the bits, it uses the read/write heads to magnetically record or "write" the bits on the platters. Data bits are not necessarily stored in succession; for example, the data in one file may be written to several different areas on different platters.

When the computer requests data stored on the disk, the platters rotate and the read/write heads move back and forth to the specified data area(s). The read/write heads read the data by determining the magnetic field of each bit, positive or negative, and then relay that information back to the computer.

The read/write heads can access any area of the platters at any time, allowing data to be accessed randomly (rather than sequentially, as with a magnetic tape). Because hard disks are capable of random access, they can typically access any data within a few millionths of a second.

What is disk formatting?

Computers must be able to access needed information on command; however, even the smallest hard disk can store millions and millions of bits.

How does the computer know where to look for the information it needs?

To solve this problem, hard disks are organized into discrete, identifiable divisions, thus allowing the computer to easily find any particular

sequence of bits. The most basic form of disk organization is called **formatting**. Formatting prepares the hard disk so that files can be written to the platters and then quickly retrieved when needed.

Hard disks must be formatted in two ways: physically and logically.

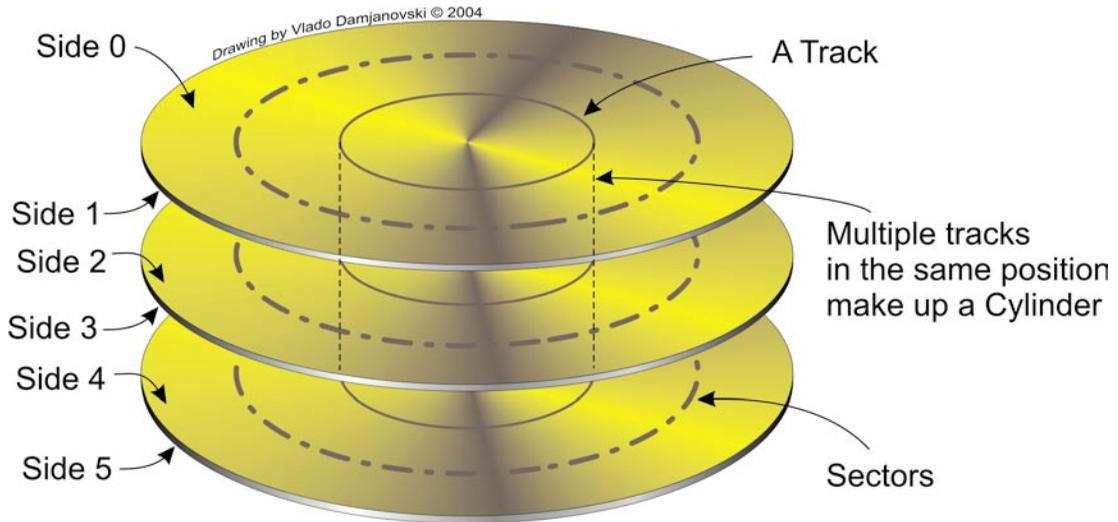
- *Physical Disk Formatting*
- *Logical Disk Formatting*

Physical disk formatting

A hard disk must be physically formatted before it can be logically formatted. A hard disk's physical formatting (also called low-level formatting) is usually performed by the manufacturer.

Physical formatting (pictured below) divides the hard disk's platters into their basic physical elements: tracks, sectors, and cylinders. These elements define the way in which data is physically recorded on and read from the disk.

· Tracks are concentric circular paths written on each side of a platter, like those on a record or compact disc. The tracks are identified by



number, starting with track zero at the outer edge.

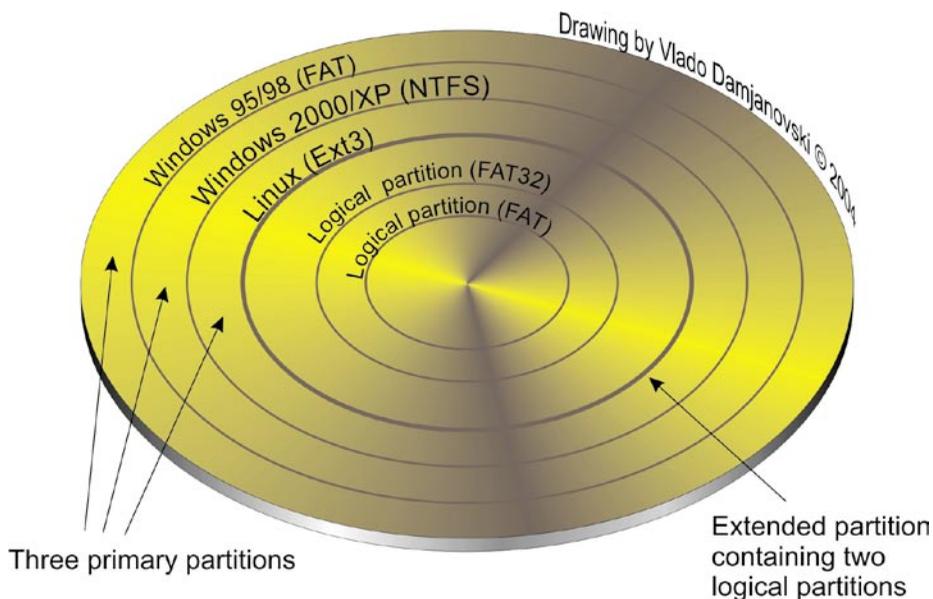
- Tracks are divided into smaller areas or sectors, which are used to store a fixed amount of data. Sectors are usually formatted to contain 512 bytes of data (there are 8 bits in a byte).

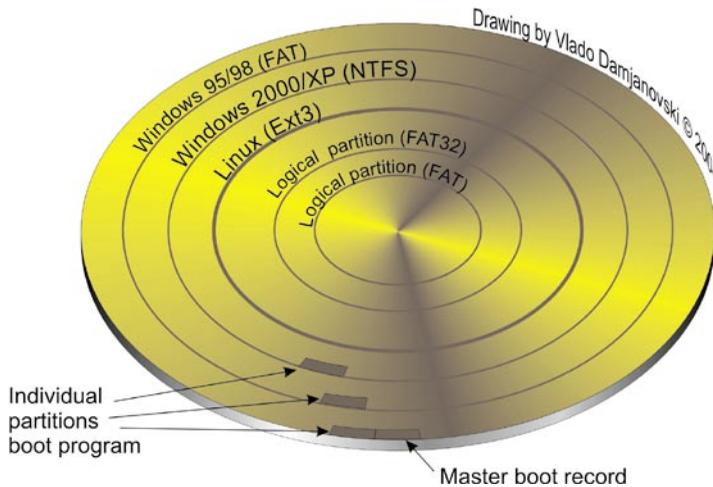
- A cylinder is comprised of a set of tracks that lie at the same distance from the spindle on all sides of all the platters. For example, track three on every side of every platter is located at the same distance from the spindle. If you imagine these tracks vertically connected, the set forms the shape of a cylinder.

Computer hardware and software frequently work using cylinders. When data is written to a disk in cylinders, it can be fully accessed without having to move the read/write heads.

Because head movement is slow compared to disk rotation and switching between heads, cylinders greatly reduce data access time.

After a hard disk is physically formatted, the magnetic properties of the platter coating may gradually deteriorate. Consequently, it becomes more and more difficult for the read/write heads to read data from or write data to the affected platter sectors. The sectors that can no longer be used to hold data are called bad sectors. Fortunately, the quality of modern disks is such that bad sectors are rare. Furthermore, most modern computers can determine when a sector is bad; if this happens, the computer simply marks the sector as bad (so it will never be used) and then uses an alternate sector.





Logical Disk Formatting

After a hard disk has been physically formatted, it must also be logically formatted. Logical formatting places a file system on the disk, allowing an operating system (such as DOS, Windows, or Linux) to use the available disk space to store and retrieve files. Different operating systems (OS) use different file systems, so the type of logical formatting you apply depends on the OS you plan to install.

Formatting your entire hard disk with one file system necessarily limits the number and types of operating systems you can install on the disk. Luckily, there is a solution to this problem. Before a disk is logically formatted, it can be divided into partitions. Each partition can then be formatted with a different file system, allowing you to install multiple operating systems. Dividing your hard disk into partitions also allows you to use disk space more efficiently.

Understanding File Systems

All file systems consist of structures necessary for storing and managing data. These structures typically include an operating system boot record, directories, and files.

A file system also performs three main functions:

- 1) tracking allocated and unused space,
- 2) maintaining directories and filenames, and

3) tracking where each file is physically stored on the disk.

Different file systems are used by different operating systems. Some operating systems (such as Windows) can recognize only one file system, while others (such as Linux) can recognize several. Some of the most common file systems are:

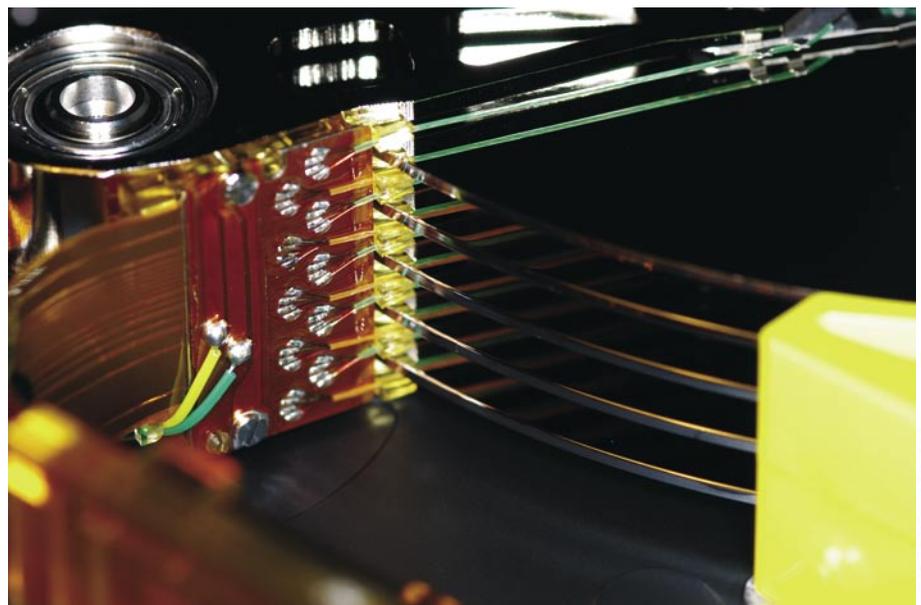
- FAT (File Allocation Table)
- FAT 32 (File Allocation Table 32)
- NTFS (New Technology File System)
- Linux Ext2, Ext3 and Linux Swap

FAT (File Allocation Table)

The FAT file system is used by DOS, Windows 3.x, and Windows 95 (in most installations). The FAT file system can also be used by Windows 98/Me/NT/2000/XP.

The FAT file system is characterized by the use of a file allocation table (FAT) and clusters. The FAT is the heart of the file system; for safety, the FAT is duplicated to protect its data from accidental deletion or corruption. Clusters are the FAT system's smallest unit of data storage; one cluster consists of a fixed number of disk sectors. The FAT records which clusters are used, which are unused, and where files are located within the clusters.

The FAT file system supports disk or partition sizes up to 2 GB, but only allows a maximum of 65,525 clusters. Therefore, whatever the size of the hard disk or partition, the number of sectors



in one cluster must be large enough so that all available space can be included within 65,525 clusters. The larger the available space, the larger the cluster size must be.

In general, large clusters tend to waste more space than small clusters. For more information on managing cluster size, click See Also on the Help toolbar.

The FAT file system also uses a root directory. This directory has a maximum allowable number of entries and must be located at a specific place on the disk or partition. Operating systems that use the FAT file system represent the root directory with the backward slash character (\), and initially display this directory at boot-up.

The root directory stores information about each sub-directory and file in the form of individual directory entries. For example, a file's directory entry holds information such as the filename, the size of the file, a date and time stamp that indicates when the file was last changed, the starting cluster number (which cluster holds the first portion of the file), and the file's attributes (for example hidden, system and so forth).

FAT 32 (File Allocation Table 32)

FAT32 is a file system that can be used by Windows 95 OEM Service Release 2 (version 4.00.950B), Windows 98/2000/XP. However, DOS, Windows 3.x, Windows NT 3.51/4.0, and earlier versions of Windows 95 cannot recognize FAT32, and are thus unable to boot from or use files on a FAT32 disk or partition.

FAT32 is an enhancement of the FAT file system and is based on 32-bit file allocation table entries, rather than the 16-bit entries used by the FAT system. As a result, FAT32 supports much larger disk or partition sizes (up to 2 terabytes).

The FAT32 file system uses smaller clusters than the FAT file system, has duplicate boot records, and features a root directory that can be any size and can be located anywhere on the disk or partition.

NTFS (New Technology File System)

The New Technology File System (NTFS) is accessible by Windows NT/2000/XP. NTFS is not recommended for use on disks less than 400 MB because it uses a great deal of space for system structures.

The central system structure of the NTFS file system is the master file table (MFT). NTFS keeps multiple copies of the critical portion of the MFT to protect against corruption and data loss.

Like FAT and FAT32, NTFS uses clusters to store data files; however, the size of the clusters is not dependent on the size of the disk or partition. A cluster size as small as 512 bytes can be specified, regardless of whether a partition is 500 MB or 5 GB. Using small clusters not only reduces the amount of wasted disk space, but also reduces file fragmentation, a condition where files are broken up over many noncontiguous clusters, resulting in slower file access. Because of its ability to use small clusters, NTFS provides good performance on large drives.

Finally, the NTFS file system supports hot fixing, a process through which bad sectors are automatically detected and marked so that they will not be used.

Linux Ext2, Ext3, and Linux Swap

The Linux Ext2 and Linux Swap file systems were developed for the Linux OS (a freeware version of UNIX). The Linux Ext2 file system supports a maximum disk or partition size of 4 terabytes. Linux Swap is used for the Linux swap file. The Ext3 file system is a journaling extension to the Ext2 file system. Journaling dramatically reduces time spent recovering a file system after a crash, thus allowing faster recovery times on single machines and recovery of a machine's file system on another machine when a cluster of nodes share a disk. [•]

