DESCRIPTION OF HARD DRIVE COMPONENTS

The Platters: The platters are the actual disks inside the drive that store the magnetized data. Traditionally platters are made of a light aluminum alloy and coated with a magnetizable material such as a ferrite compound that is applied in liquid form and spun evenly across the platter or thin metal film plating that is applied to the platter through electroplating, the same way that chrome is produced. Newer technology uses glass and/or ceramic platters because they can be made thinner and also because they are more efficient at resisting heat. The magnetic layer on the platters has tiny domains of magnetization that are oriented to store information that is transferred through the read/write heads. Most drives have at least two platters and the larger the storage capacity of the drive, the more platters there are. Each platter is magnetized on each side, so a drive with 2 platters has 4 sides to store data.

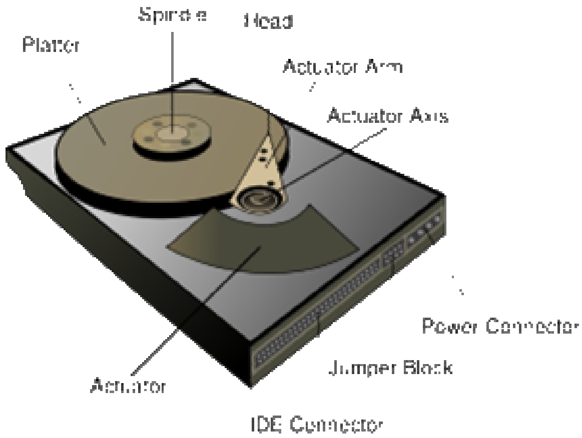

The Spindle and Spindle Motor The platters in a drive are separated by disk spacers and are clamped to a rotating spindle that turns all the platters in unison. The spindle motor is built right into the spindle or mounted directly below it and spins the platters at a constant set rate ranging from 3,600 to 7,200 RPM. The motor is attached to a feedback loop to ensure that it spins at precisely the speed it is supposed to.

The photo is the inside of the Hard drive showing the Heads in a park position or loading zone, the platters, the Spindle, and Spindle Motor assembly. The HDD's spindle system relies on air pressure inside the enclosure to support the heads at their proper flying height while the disk rotates. Hard disk drives require a certain range of air pressures in order to operate properly. The connection to the external environment and pressure occurs through a small hole in the enclosure (about 0.5 mm in diameter), usually with a carbon filter on the inside (the breather filter, see below). If the air pressure is too low, then there is not enough lift for the flying head, so the head gets too close to the disk, and there is a risk of head crashes and data loss.

<u>The Read/Write Heads:</u> The read/write heads read and write data to the platters. There is typically one head per platter side, and each head is attached to a single actuator shaft so that all the heads move in unison. When one head is over a track, all the other heads are at the same location over their respective surfaces. Typically, only one of the heads is active at a time, i.e., reading or writing data. When not in use, the heads rest on the stationary platters, but when in motion the spinning of the platters create air pressure that lifts the heads off the platters. The space between the platter and the head is so minute that even one dust particle or a fingerprint could disable the spin. This necessitates that hard drive assembly be done in a clean room. When the platters cease spinning the heads come to rest, or park, at a predetermined position on the heads, called the landing zone.

The read write heads HD heads are kept from contacting the platter surface by the air that is extremely close to the platter; that air moves at, or close to, the platter speed. The record and playback head are mounted on a block called a slider, and the surface next to the platter is shaped to keep it just barely out of contact. The write head magnetizes a region by generating a strong local magnetic field, the read and write elements are separate, but in close proximity, on the head portion of an actuator arm. The read element is typically magneto-resistive while the write element is typically thin-film inductive. Due to the extremely close spacing between the heads and the disk surface, any contamination of the read-write heads or platters can lead to a head crash — a failure of the disk in which the head scrapes across the platter surface, often grinding away the thin magnetic film and causing data loss. Head crashes can be caused by electronic failure, a sudden power failure, physical shock, wear and tear, corrosion, or poorly manufactured platters and heads.

Here is a close-up of the read/write heads on the disk platters. These ride on a thin layer of air generated by the rotating platters. . The platters are made from a non-magnetic material, usually aluminium alloy or glass, and are coated with a thin layer of magnetic material. The read-and-write head is used to detect and modify the magnetization of the material immediately under it. There is one head for each magnetic platter surface on the spindle, mounted on a common arm.

The Head Actuator: All the heads are attached to a single head actuator, or actuator arm, that moves the heads around the platters. Older hard drives used a stepper motor actuator, which moved the heads based on a motor reacting to stepper pulses. Each pulse moved the actuator over the platters in predefined steps. Stepper motor actuators are not used in modern drives because they are prone to alignment problems and are highly sensitive to heat. Modern hard drives use a voice coil actuator, which controls the movement of a coil toward or away from a permanent magnet based on the amount of current flowing through it. This guidance system is called a servo.
The platters, spindle, spindle motor, head actuator and the read/write heads are all contained in a chamber called the head disk assembly (HDA). Outside of the HDA is the logic board that controls the movements of the internal parts and controls the movement of data into and out of the drive.

The hard drive's electronics control the movement of the actuator and the rotation of the disk, and perform reads and writes on demand from the disk controller. Feedback of the drive electronics is accomplished by means of special segments of the disk dedicated to servo feedback. These are either complete concentric circles (in the case of dedicated servo technology), or segments interspersed with real data (in the case of embedded servo technology). The servo feedback optimizes the signal to noise ratio of the GMR sensors by adjusting the voice-coil of the actuated arm. The spinning of the disk also uses a servo motor. Modern disk firmware is capable of scheduling reads and writes efficiently on the platter surfaces and remapping sectors of the media which have failed. . An actuator arm (or access arm) moves the heads on an arc (roughly radically) across the platters as they spin, allowing each head to access almost the entire surface of the platter as it spins. The arm is moved using a voice coil actuator or (in older designs) a stepper motor. As data density increased, read heads using magneto resistance (MR) came into use; the electrical resistance of the head changed according to the strength of the magnetism from the platter.

The landing zone: Is an area of the platter usually near its inner diameter (ID), where no data is stored. This area is called the Contact Start/Stop (CSS) zone. Disks are designed such that either a spring or, more recently, rotational inertia in the platters is used to park the heads in the case of unexpected power loss. In this case, the spindle motor temporarily acts as a generator, providing power to the actuator.

Spring tension from the head mounting constantly pushes the heads towards the platter. While the disk is spinning, the heads are supported by an air bearing and experience no physical contact or wear. In CSS drives the sliders carrying the head sensors (often also just called heads) are designed to survive a number of landings and takeoffs from the media surface, though wear and tear on these microscopic components eventually takes its toll. Most manufacturers design the sliders to survive 50,000 contact cycles before the chance of damage on startup rises above 50%. However, the decay rate is not linear: when a disk is younger and has had fewer start-stop cycles, it has a better chance of surviving the next startup than an older, higher-mileage disk (as the head literally drags along the disk's surface until the air bearing is established). For example, the Seagate Barracuda 7200.10 series of desktop hard disks are rated to 50,000 start-stop cycles, in other words no failures attributed to the head-platter interface were seen before at least 50,000 start-stop cycles during testing.

Load/Unload technology relies on the heads being lifted off the platters into a safe location, thus eliminating the risks of wear and stiction altogether. The first HDD RAMAC and most early disk drives used complex mechanisms to load and unload the heads. Modern HDDs use ramp loading, first introduced by Memorex in 1967 [36], to load/unload onto plastic "ramps" near the outer disk edge.
All HDDs today still use one of these two technologies listed above. Each has a list of advantages and drawbacks in terms of loss of storage area on the disk, relative difficulty of mechanical tolerance control, non-operating shock robustness, cost of implementation, etc.
Addressing shock robustness, IBM also created a technology for their ThinkPad line of laptop computers called the Active Protection System. When a sudden, sharp movement is detected by the built-in accelerometer in the Thinkpad, internal hard disk heads automatically unload themselves to reduce the risk of any potential data loss or scratch defects. Apple later also utilized this technology in their PowerBook, iBook, MacBook Pro, and MacBook line, known as the Sudden Motion Sensor. Sony [37], HP with their HP 3D DriveGuard [38] and Toshiba [39] have released similar technology in their notebook computers.
This accelerometer based shock sensor has also been used for building cheap earthquake sensor networks.

Architecture: The motor has an external rotor; the stator windings are copper-colored. The spindle bearing is in the center. To the left of center is the actuator with a read-write head under the tip of its very end (near center); the orange stripe along the side of the arm, a thin printed-circuit cable, connects the read-write head to the hub of the actuator. The flexible, somewhat 'U'-shaped, ribbon cable barely visible below and to the left of the actuator arm is the flexible section, one end on the hub that continues the connection from the head to the controller board on the opposite side.
The head support arm is very light, but also rigid; in modern drives, acceleration at the head reaches 250 gs.

A hard disk drive with the platters and spindle motor hub removed showing the copper colored motor coils surrounding a bearing at the center of the spindle motor. The silver-colored structure at the upper left is the top plate of the permanent-magnet and moving coil "motor" that swings the heads to the desired position.

 Removing the platter motor is tricky - you have to give the spindle sticking out underneath the drive a tap using a hammer and a center punch. Beneath this plate is the moving coil, attached to the actuator hub, and beneath that is a thin neodymium-iron-boron (NIB) high-flux magnet. That magnet is mounted on the bottom plate of the "motor".

 Here is the platter motor removed from the drive casing.

 Here is a close-up of the platter motor coils. The coil, itself, is shaped rather like an arrowhead, and made of doubly-coated copper magnet wire. The inner layer is insulation, and the outer is thermoplastic, which bonds the coil together after it's wound on a form, making it self-supporting. Much of the coil, sides of the arrowhead, which points to the actuator bearing center, interacts with the magnetic field to develop a tangential force to rotate the actuator. Considering that current flows (at a given time) radially outward along one side of the arrowhead, and radially inward on the other, the surface of the magnet is half N pole, half S pole; the dividing line is midway, and radial.

 This is a particle catcher. It's purpose it to catch any small particles tossed off the disk as it rotates. On a good drive this will be clean but sometimes you find small metal particles on it when the drive has suffered a head crash.

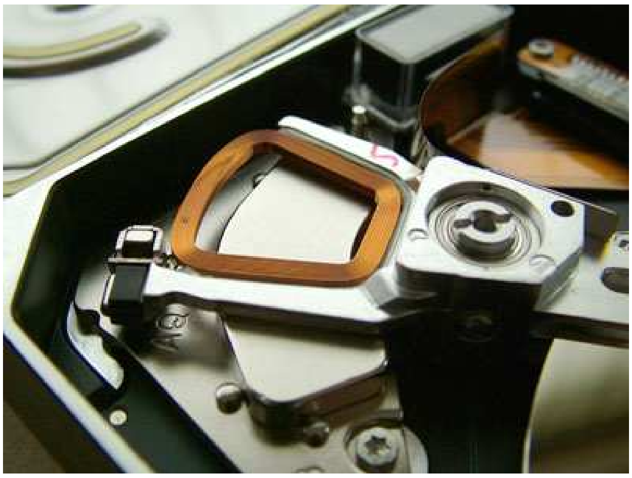 This is the actuator that controls the read/write heads and moves them across the drive.

 To unto the actuator you undo the screw holding it. However, you also need to undo the platters at the same time.

 Un-doing the platter screws. These are TORX No. 6 too.

 Platters removed along with the actuator and read/write heads.

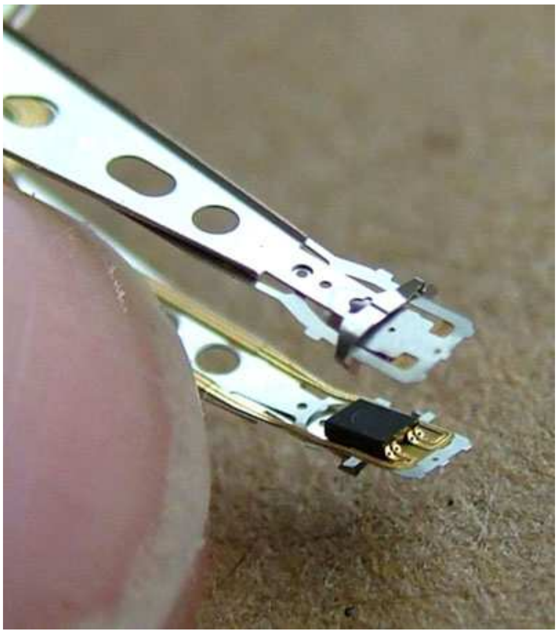 Here's a close-up of the bit that makes a hard drive work. This is the read/write head that writes data to the platters and reads it off - these are an amazing feat of engineering. These heads skim the surface of the platters incredibly close as it goes about reading and writing data - if they touch the platters when operating they will smash the delicate magnetic coating or even themselves, rendering the drive scrap.

 This is the ultra-strong magnet powering the motor for the actuator arm

 Here is a filter that cleans the air entering the drive. A drive isn't airtight, instead it allows the pressure to be equalized in the drive but the air coming into the drive has to be clean.

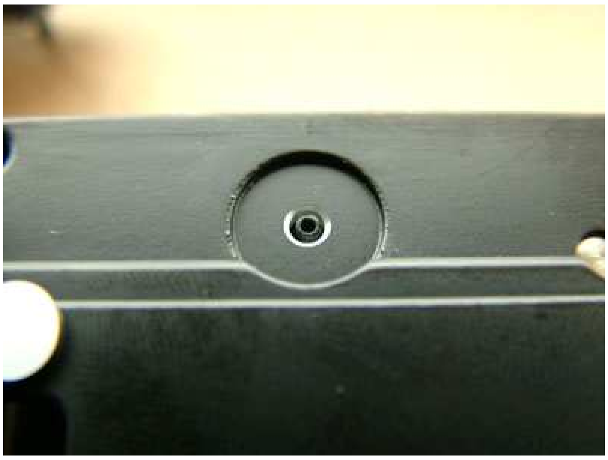 Here is the vent hole on the exterior of the drive.

 Finally, here's a look at all the parts that go into a drive.

<u>Determining the capacity or LBA of a Hard drive:</u> The capacity of an HDD can be calculated by multiplying the number of cylinders by the number of heads by the number of sectors by the number of bytes/sector (most commonly 512). Drives with the ATA interface and a capacity of eight gigabytes or more behave as if they were structured into 16383 cylinders, 16 heads, and 63 sectors, for compatibility with older operating systems. Unlike in the 1980s, the cylinder, head, sector (C/H/S) counts reported to the CPU by a modern ATA drive are no longer actual physical parameters since the reported numbers are constrained by historic operating-system interfaces and with zone bit recording the actual number of sectors varies by zone. Disks with SCSI interface address each sector with a unique integer number; the operating system remains ignorant of their head or cylinder count.

Hard disk drive manufacturers specify disk capacity using the SI prefixes mega-, giga- and tera-, and their abbreviations M, G and T. Byte are typically abbreviated B.
Most operating-system tools report capacity using the same abbreviations but actually use binary prefixes. For instance, the prefix mega-, which normally means 106 (1,000,000), in the context of data storage can mean 220 (1,048,576), which is nearly 5% more. Similar usage has been applied to prefixes of greater magnitude. This results in a discrepancy between the disk manufacturer's stated capacity and the apparent capacity of the drive when examined through most operating-system tools. The difference becomes even more noticeable for a gigabyte (7%), and again for a terabyte (9%). For a petabyte there is a 11% difference between the SI (10005) and binary (10245) definitions. For example, Microsoft Windows reports disk capacity both in decimal-based units to 12 or more significant digits and with binary-based units to three significant digits. Thus a disk specified by a disk manufacturer as a 30 GB disk might have its capacity reported by Windows 2000 both as "30,065,098,568 bytes" and "28.0 GB". The disk manufacturer used the SI definition of "giga", 109 to arrive at 30 GB; however, because Microsoft Windows, Mac OS and some Linux distributions use "gigabyte" for 1,073,741,824 bytes (230 bytes), the operating system reports capacity of the disk drive as (only) 28.0 GB.