Hard Disk Quality and Reliability

- A hard disk's performance is its most important characteristic--**right up until the point where it stops working**. Then, suddenly, you don't care how fast it is--or rather, *was*. You just want it to start working again.

- Many people take their hard disk drives for granted, and don't think about their reliability much (other than worrying about their disk crashing some day). While the technology that hard disks use is very advanced, and reliability today is much better than it has ever been before, the **nature of hard drives is that every one will, some day, fail**. It is important to understand how drives fail and why, and how to interpret what manufacturers claims about reliability really mean.

- **This section takes** a comprehensive look at **hard disk quality and reliability**. First, I explain in some detail the different hard disk quality specifications that you are likely to encounter when shopping for a hard disk. Next, I discuss issues relevant to the matter of quality and reliability. I then discuss some of the features that hard disk manufacturers are including in their drives to improve quality and reliability. Finally, I discuss warranty issues, and what to do should you have a hard drive failure.

Hard Disk Quality and Reliability Specifications

- There are a number of different specifications used by hard disk drive manufacturers to indicate the quality and reliability of their products.
- Some of these, such as **MTBF**, are frequently discussed (but not always all that well understood).
- Others are obscure and typically of interest only to hard drive aficionados.
- All are important to those who care about hard disk quality--which should be anyone who stores data on a hard disk. :^) In this section I discuss the most important of these specifications, what they mean, and perhaps most importantly, what they don't mean! You'll also find some discussion of specifications in the section on quality and reliability issues, particularly temperature specifications and noise specifications.

Mean Time Between Failures (MTBF)

- The most common specification related to drive reliability is *mean time between failures* or ***MTBF***. This value, usually measured in hours, is meant to **represent the average amount of time that will pass between random failures on a drive** of a given type. It is usually in the range of 300,000 to 1,200,000 hours for modern drives today (with the range increasing every few years) and is specified for almost every drive.

- This number is very often misinterpreted and misused. Usually, the "analysis" goes like this: "Gee, a year contains 8,766 hours. That means my 500,000 MTBF drive should last 57 years." (I have even seen this *on the web site of a major hard disk manufacturer* that shall remain nameless to spare them the embarrassment!) After concluding that the MTBF means the drive will last for decades, amusingly, one of **two *opposite* things** usually happens: either the person actually thinks the drive will last half a century or longer, or the opposite: they realize this is crazy and so they write off the entire MTBF figure as "obvious exaggeration and therefore useless". The real answer of course is neither. (It is obviously impossible for any individual hard disk to be tested to anywhere near the amount of time required to provide a MTBF factor near even 100,000, never mind 500,000.)

Mean Time Between Failures (MTBF)

- To be interpreted properly, the MTBF figure is intended to be used **in conjunction** with the useful **service life** of the drive, the typical amount of time **before** the **drive enters the period** where failures due to component wear-out increase.

- MTBF only applies to the ***aggregate analysis*** *of large numbers of drives*; it says nothing about a particular unit. If the MTBF of a model is 500,000 hours and the service life is 5 years, this means that a drive of that type is supposed to last for five years, and that of a large group of drives operating within this timeframe, on average they will accumulate 500,000 of total run time (amongst all the drives) before the first failure of any drive.
- Or, you can think of it this way: if you used one of these drives and replaced it every five years with another identical one, in theory it should last 57 years before failing, on average

Mean Time Between Failures (MTBF)

- There are in fact 2 different types of MTBF figures. When a manufacturer is introducing a new drive to the market, it obviously has not been in use in the real world, so they have no data on how the drive will perform. Still, they can't just shrug and say "who knows?", because many customers want to know what the reliability of the drive is likely to be.

- To this end, the companies calculate what is called a ***theoretical MTBF figure***. This number is based primarily **upon the analysis of historical data**; for example: the historical failure rate of other drives similar to the one being placed on the market, and the failure rate of the components used in the new model. It's important to realize that these MTBF figures are estimates based on a theoretical *model* of reality, and thus are limited by the constraints of that model. There are typically as MTBF figure to be valid: the drive must be properly operating within allowable environmental limits, anc MTBF figures also cannot typically account for "ran conditions such as a temporary quality problem dur particular lot of a specific type of drive.

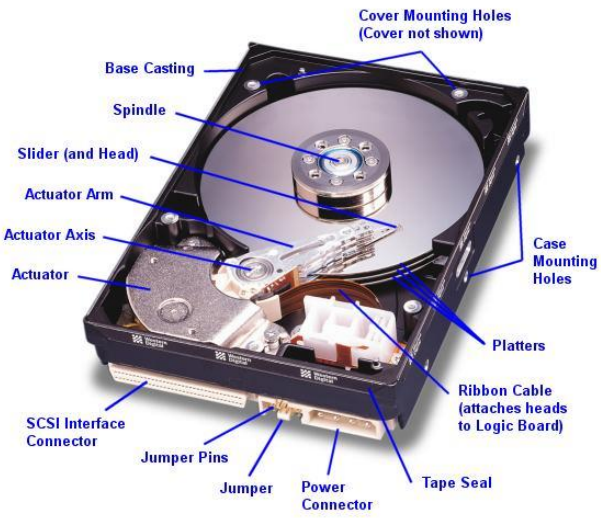

Mean Time Between Failures (MTBF)

- After a particular model of drive has been in the market for a while, **say a year**, the **actual failures** of the **drive can be analyzed** and a calculation made to determine the drive's ***operational MTBF***. This figure is derived by analyzing field returns for a drive model and comparing them to the installed base for the model and how long the average drive in the field has been running. Operational MTBFs are typically lower than theoretical MTBFs because they include some "human element" and "unforeseeable" problems not accounted for in theoretical MTBF.
- Despite being arguably more accurate, operational MTBF is rarely discussed as a reliability specification because most manufacturers don't provide it as a specification, and because most people only look at the MTBFs of new drives--for which operational figures are not yet available.

Mean Time Between Failures (MTBF)

- The key point to remember when looking at any **MTBF** figure is that it is **meant** to be an **average**, based on testing done on many hard disks **over** a **smaller period of time**. Despite the theoretical numbers sometimes seeming artificially high, they do have value when put in proper perspective; a drive with a much higher MTBF figure is probably going to be more reliable than one with a much lower figure. As with most specifications, small differences don't account for much; given that these are theoretical numbers anyway, **350,000** is not much different than 300,000.
- Overall, MTBF is what I consider a "reasonably interesting" reliability statistic--**not something totally useless**, but definitely something to be taken with a grain of salt.
- I personally view the drive's **warranty length** and **stated service life to be more indicative** of what the manufacturer really thinks of the drive. I personally would rather buy a hard disk with a stated service life of five years and a warranty of three years, than one with a service life of three years and warranty of two years, even if the former has an MTBF of 300,000 hours and the latter one of 500,000 hours.
- In the real world, the actual amount of time between failures will depend on many factors, including the operating conditions of the drive and how it is used; this section discusses component life. Ultimately, however, luck is also a factor, so keep those backups current.

Service Life

- Due to the **incredibly precise nature** of the technology used in hard disks (we're talking about millionths of an inch in many regards) and the fact that mechanical components are used, **it is impossible to guarantee the reliability of even the highest-quality disks for more than a few years**.

- Hard disks : they have a relatively high rate of "infant mortality" (or early failure), a period of very low failures for several years, and then a **gradual reduction in reliability** as they reach the end of their useful life. To ensure that those who rely on hard drives do not continue to use them as they enter the period of increased failure that results from component wear-out, every hard disk has defined as one of its specifications the manufacturer's intended *service life* for the product.

- The **service life** of a modern hard disk is usually **about three to five years**. In my opinion, what the manufacturer is telling you with this number is this: "If your drive gives you **good service** for this length of time without any trouble, you've gotten good value from it and you should consider getting a new drive if reliability is important to you; the drive isn't going to fail the instant it hits this age, but the chances of it failing will increase significantly the longer you use it past this point". This number is in many ways more important than the **"impressive"** MTBF numbers that are so eagerly thrown around, and is in fact intended to be used in conjunction with MTBF to present a realistic picture of drive life and reliability.

Service Life

- Interestingly, the claimed **service life** is **often longer than the warranty period** for the drive. For example, the service life might be five years but the warranty period only three years. Think about what this means. ;^) Basically, it says that the manufacturer thinks the drive should last five years, but they aren't going to bet on it lasting more than three! I personally think the warranty period is a better indication of a drive's true useful life--to a point--because the warranty period is where the manufacturer "puts their money where their mouth is". Sometimes warranty periods are made longer for marketing reasons, but no manufacturer will warrant a product for three years if they expect to see significant problems in that time, because they will lose their shirts.

- If reliability is important, you may want to make sure you get a new hard disk before your **warranty period is up**, or at the very least, before the drive reaches the end of its service life. Of course, at the rate technology changes, in three years you will probably want a new disk **for performance reasons as well**. Oh, and it goes without saying that anyone who cares about hard disk reliability should be performing regular backups.

- The real world life of a hard disk is almost always going to be significantly higher than the stated service life, which tends to be very **conservative**. One reasons why is that even if the warranty doesn't cover the entire stated service life, most companies don't want to see their drives failing within the period of time that they say they won't--it looks bad. Another reason is because the **service life only represents when the odds** of **failure increase**; there is still a statistical distribution of actual failures. I have used many hard disks that have been in operation for 10 or even 12 years, and these are older technology drives that never had anywhere near the stated reliability levels of today's disks. So there certainly is no reason to believe that the day after your drive's third anniversary, you need to yank it and replace it with a new one. But just remember that the longer you go past the drive's service life, the more the chances of a failure increase. The environmental conditions, and how the drive is used, will also have an important impact on its overall life.

Start/Stop Cycles

- The read/write heads of the hard disk float on a cushion of air over the surface of the platters. Each time the power is applied to the drive, the platters spin up to speed, and the buildup of air under the heads causes them to lift off the data surfaces. When the power is cut, the platters spin down, the cushion of air dissipates, and the heads drop back to the surface of the drive.

- **Each time the drive starts and stops** a small amount of **wear occurs** to the heads and also to other components such as the **spindle motor**. For this reason, hard drives are given a specification for the minimum number of start/stop cycles they are designed to handle during their service life. The value for a desktop drive is typically between 30,000 and 50,000 cycles (and remember that this is not an average, but a minimum). Notebook drives, which are more commonly spun up and down a great deal to save battery power, usually have even higher numbers.

- IBM drives that use its **head load/unload technology** are often given a specification for *minimum load/unload cycles* instead of start/stop cycles. This is basically the same concept, except that the numbers are typically much higher. Some IBM notebook drives are spec'ed for 300,000 load/unload cycles! Even if you started and stopped the drive 100 times a day it would take you over eight years to get to that number.

Error Rates

- Fortunately, true **errors are rarely encountered** when using hard disks. To help users understand better the rate at which errors will occur with a hard disk, manufacturers provide anywhere from one to **several error rate specifications**.

- The **most common** error rate spec is the drive's **unrecoverable error rate**, which is usually specified as "<**1 in 10^N bits**", where "N" is usually between 12 and 15. **"Unrecoverable"** means that the drive is **unable to use its error-correcting code**, retries or other techniques to recover from the error in reading the disk and thus properly recreate the data. If "N" is 14, then that means this will occur every 100 trillion (100,000,000,000,000) bits read from the disk. Not too shabby. :^)

- In fact, drives usually have several different error rate specifications; they are just usually not put into the short data sheets commonly distributed by hard disk manufacturers. In fairness, unrecoverable errors are the most important ones, but there are also **specifications for recoverable errors**, **errors recovered after multiple reads and so on**. To find these, you generally need to download the product manual for the drive or contact the manufacturer's technical support department. For a full explanation of these various errors and what they mean, see this full discussion of errors and error recovery.
- Error rate specifications are typically used to compare drives. Within the same general class of drive there are usually relatively few differences between manufacturers. The biggest difference in error rate can be seen by comparing newer drives to older ones--newer drives are usually significantly better despite pushing the design envelope with higher speed and much greater capacity.

Warranty Length

- Retail warranties for hard disks are usually either 3 years or 5 years in length. 3-year warranties are typical for consumer-grade (usually IDE/ATA), and 5-year warranties are the norm for "enterprise" class (usually SCSI)

- **Does** the fact that a drive has a **five-year warranty mean that it is higher in quality than one with a three-year warranty**? The answer is "yes and no". To some extent these more expensive drives are manufactured to more exacting standards because they are sold for much higher prices to clients who plan to use them in often-critical business applications. However, at least *some* of the motivation behind longer warranties is to sell the drives as being high quality and/or to match the warranty periods of other companies. It's a bit of a "chicken-and-egg" situation. And since every drive you look at it in a given class is likely to have the same length of warranty as any other, the length of the warranty doesn't do much to help you differentiate between drives. (Although it does establish a standard for a given class of drive and give you reason to be suspicious of any drive with a warranty that is much below that standard.)

- Here's something else that most people don't really think about--especially those who recommend high-end drives with five-year warranties over "cheaper" drives with "only" three-year warranties: are you really still going to be *using* that new drive in four years? Considering that most people who **buy high-end drives need high performance, and considering that the performance of today's drives will be a joke compared to that of drives made in four years, the answer for many people is "no".**

Hard Disk Quality and Reliability Issues

- **Quality and reliability specifications** are theoretical numbers that reflect the drive's quality characteristics under average--or in some cases "**ideal"--conditions**.
- In actual operation, however, the **reliability of a hard disk depends** as much on how the storage subsystem is implemented as it does on the characteristics of the drive itself. No implementation factors can make the drive ***more* reliable** than it is specified to be, but mistakes in the way the drive is used can sometimes make it *less* reliable. There are several different issues that impact upon this "real world" reliability.
- In this section I will take a look at topics related to reliability such as **cooling, noise and vibration, and power management**. I will also discuss determining when a drive is failing, and also look at some software causes for reliability problems (which are often mistakenly blamed on hardware!)

Hard Disk Quality: Luck Of The Draw?

- The simple fact of the matter is that **most major hard drive manufacturers make very high quality products**, and most hard disks provide their owners with years of very reliable service. However, all manufacturers make the occasional bad drive, and **sometimes, manufacturers will have a problem with a particular product**. If you happen to buy one of these, you will experience a failure, and in all likelihood you will hate that company and avoid their products from then on, perhaps with good reason. The problem is that many people will generalize this very small sample size into "Brand X sucks", when this very well may not be the case. They just may have been unlucky with what might in actuality be one of the best drives on the market.

- Meanwhile, they will post a frustrated-sounding message telling everyone **"DON'T BUY X!"** Human nature is such that there could be thousands of people using that particular model with no problems, but almost none will bother to post saying "I bought a new PC with an X hard drive in it and every day I use it, it doesn't fail!" :^)

- **There are many ways that a hard disk can fail**. The one that usually comes to mind first, the infamous "head crash" is not the only way or even necessarily the most common any more. There can also be problems with bad sectors showing up on the drive, for example. Many people don't think of this, but the integrated controller can also sometimes be the culprit for a bad drive. See this discussion of failure modes, and this section in the Troubleshooter for specifics on drive failures and possible solutions.

System and Drive Cooling

- In order to ensure long life and proper reliable operation, **all hard disks** are designed to function **only in specific temperature ranges**. The user of the hard disk must keep the drive within the specifications to be sure that the drive will continue to work well, and to maintain the drive's warranty status. You can find out the temperature allowances for a particular drive model by consulting the drive's product manual or data sheet, normally available for free download at the drive manufacturer's web site.
- There are in fact several different temperature limits that are specified for hard disks:
- **Non-Operating Temperature Range (okolina):** This is the range of acceptable temperatures for the drive when it is either in storage, or in a PC that is in transit or otherwise not operating. This range is normally very wide, much wider than the operating temperature limits, since the unit is much less sensitive to extremes of temperature when it is not functioning. A typical range would be -40°C (-40°F) to 70°C (158°F). Clearly few users will have a problem with these numbers.
- **Warning:** If a drive is allowed to go below freezing, or if it is quickly exposed to a large temperature change, it must be *acclimated* before use.

System and Drive Cooling

- **Minimum Operating Temperature:** The lowest acceptable temperature for the drive when in operation. A single number is normally provided for this value, with **5°C** (41°F) being typical. Again, due to the heat generated by the hard disk and the fact that almost all PCs are used indoors, this is rarely much of a concern.
- **Maximum Operating Temperatures:** The highest temperature allowed for the drive when in operation. Since the mechanical and electrical components within the hard disk--especially the spindle motor--produce heat, the biggest problem with keeping drives within operating parameters is not exceeding maximum allowable temperatures.

- A number of **different temperature values are usually provided**, depending on the drive manufacturer and model:
  - **Case Temperature:** The highest acceptable temperature allowed, measured at a specific point on the metal of the case.
  - **Component Temperatures:** Some manufacturers provide a number of different maximum temperature values, measured on the surface of various components on the drive (especially, the **hard disk's logic board**). This method is more precise and takes into account the fact that some components run hotter than others. In practice, due to the amount of work involved, these different temperature measurements are rarely used except in special applications.
  - **"Absolute" and "Reliability" Temperatures:** Some companies provide two sets of case and component temperature limits. The first set is the absolute maximum temperature(s) allowed for the drive. The second set is the maximum allowed for the drive in order to meet its reliability and MTBF specifications. Of course, the reliability temperatures are lower than the maximum temperature. Since reliability is so important for hard disks, the "reliability" temperatures are the ones to aim for.

- The temperature at which the drive operates is dependent on the temperature of the system--you can have a drive run too hot in a system that is otherwise quite cool, but if the rest of the system case is hot, the drive doesn't have a chance.

Noise and Vibration

- Due to their mechanical nature, there are several characteristics of hard disks that aren't relevant to most other PC components. Among these are the ***noise* and *vibration* created by the disk drive**. (Noise and vibration are of course two parts of the same phenomenon.) Since most devices inside the PC have no moving parts, they don't make noise. Hard disks on the other hand have very high-speed motors and actuators, both of which can make considerable noise and generate vibrations.
- Hard drive noise specifications are normally provided as two numbers. The values correspond directly to the two different mechanical devices that make noise in the drive:
- **Idle Noise:** This is the noise level the drive makes whenever the drive is in operation (and not in sleep or standby mode), regardless of what it is doing; it is caused by the spindle motor and the rotation of the platters.
- **Seek Noise:** This is the noise level when the drive is performing positioning tasks, particularly random seeks. It is caused by the movement of the actuator assembly (as well as the spindle motor and platters of course, since they continue to spin during seeks!)
- Hard drive noise is specified by the manufacturer in *bels* (or sometimes *decibels*; they are the same unit except that one bel equals ten decibels.) A "bel" is a logarithm of a ratio between two sound values. It is defined as follows:
- **Noise level (bels) = $\log_{10}$ (N1 / N0)**
- Where "log10" means a base-10 logarithm, "N1" is the noise level of the device being measured, and "N0" is a *reference* noise level. (Usually, N0 is taken as the threshold of human hearing, the quietest sound audible to the average person. The actual value of N0 is not important as long as all the noise figures are calculated using the same reference point.) The reason for the logarithm is notational convenience; the human ear can detect sound intensities in a *tremendous* range: up to a factor of 1,000,000,000,000 from the quietest to the loudest. Logarithmic notation expresses these sounds in a range of 0 to 12 bels, saving us all those zeroes. :^)

Noise and Vibration

- Companies always provide noise specifications, and some hard disk reviewers have started measuring noise levels with noise meters. Frankly, I don't put a lot of stock in these measurements and figures. I would even go so far as to say that they have *almost no bearing on real world noise perception* and by themselves are not useful for deciding if a drive model is going to give you a noise problem. Here are my reasons:
- **Noise is a very *subjective* and personal matter**. It's very possible to put two people in a room with the same system and have person A be "driven crazy" by the noise and person B not even realize the PC is turned on.
- Some types of noise bother certain people more than others, because humans have different sensitivity to sounds of various frequencies. For example, some people find certain drives have a high-pitched "whine" due to their spindle motors that others can't even detect; some find the spindle noise acceptable but the seek noise irritating, and some are the opposite.

- Hard disk noise levels depend a *great* deal on the case in which they are installed, and the location of the case relative to where the user is sitting. A small desktop case right in front of the user will make noise seem much more significant than a large tower case with the hard disk mounted in the rear, sitting on the floor next to a desk! And if a high-speed drive is installed in a server in a locked, ventilated, insulated computer room, does its noise level matter much at all?
- Absolute noise figures don't mean anything unless you take them in the context of the surroundings. If a drive has a high noise rating then it might be considered "noisy", but if you are using it in a crowded office with a lot going on, you're probably not going to hear it.
- Noise levels can sometimes vary between specific units even of the same model of drive.

Noise and Vibration

- While I trust manufacturers to make accurate sound readings, some of the individual hardware reviewers using sound meters haven't been properly trained in their use. Sound measurement is not a simple thing; you must compensate for the noise of other components and other effects to ensure the readings are accurate.
- So fine, look at the specifications for the drive. Clearly, if there is a big difference in the noise values, you are *more likely* to have a noise issue with the drive that has the higher specification. But in reality, it all depends on your personality, your surroundings, and your ears. If you aren't the type to care about the noise level of your PC, just don't worry about the whole issue. If you are, then you should if at all possible find a PC that has the drive you are considering installed into it, and listen to it yourself. That's the ideal trial step, but not always possible. A second choice is to look on the Internet for *a significant number of subjective comments*. If dozens of people are complaining about the noise level of a drive, then there is a good chance the drive may bother you; look for one that is commonly used and has many comments saying the drive is quiet.
- Another key issue with hard disk noise is that while seek noise *specifications* just consist of the actuator noise "added" to the spindle noise, they are very different kinds of noise! If you are sitting in front of a PC for hours, your brain will tend to "tune out" the background noise of the spindle. The noise of the actuator is intermittent however, and much louder. It is this noise that often causes the most objections, and again, it's a very personal thing.
- To some extent, you must realize that noise and vibration are a price you pay for increased performance. Idle noise is directly correlated to spindle speed, and seek noise is correlated to faster actuators (and hence reduced seek times). It is also true however that manufacturers are constantly improving their products, which means that the *generation* of a given technology is important. For example, the first 7200 RPM drives were rather noisy and hot; now 7200 RPM is a mature technology and the drives of this speed run *much* quieter and cooler than their ancestors.

- An important final note: to minimize vibration and noise in a PC, ensure that the drive has been *properly mounted*. If the drive is loose then its inherent vibrations can become amplified by the metal that touches it, making the entire PC seem to resonate in some cases!

Overclocking

- A popular activity undertaken by many PC hobbyists is the process of *overclocking*. This refers to running hardware above its standard and rated speed specifications in order to increase performance. I don't think overclocking makes sense for most people; I provide some background on overclocking, and my opinions on it, in this section.
- Overclocking a PC inherently reduces the reliability of virtually every part of the system. Since a PC is interconnected, and overclocking usually affects key components such as the system processor and the motherboard, in some ways everything in the system is involved. When overclocking involves running just the CPU at a higher speed, it has little *direct* reliability effect on the storage subsystem. **However, when overclocking the system bus, reliability concerns with hard drives often come into play.** Overclocking the PCI bus upon which the IDE/ATA interface runs can cause IDE hard disks to behave spuriously. Since most SCSI host adapters also run on the PCI bus, they can also behave strangely if the PCI bus is run faster than its normal speed.
- How a particular hard disk will react to overclocking is not easy to determine, and like most overclocking activities, requires trial and error. It also depends on how far you are pushing things, and especially, on the type of motherboard you are using and its integrated hard disk controller. However while it is true that "your mileage may vary", as the saying goes, it is also the case that some hard disk brands and models take to overclocking more than others. Research on any of the hundreds of overclocking sites, newsgroups and other resources can help you decide on a suitable drive if overclocking is one of your priorities.

Diagnostic Software

- **All hard disk manufacturers** write and maintain **special *diagnostic software* for their drives**. The purpose of diagnostic software is simple: **to test drives and diagnose potential problems with them**. These programs usually come with retail drives on a **floppy disk**, or can be **downloaded from the manufacturer's web site (at no charge).** In some cases a single program will function for any of the manufacturer's drives, but in other cases a different program is needed depending on the drive model; check the manufacturer's instructions.

- If you suspect a problem with your hard disk, you should always run a diagnostic program to check the status of the drive. For one thing, the technical support department of most manufacturers will require the information provided by the diagnostic tool in order to help you with any problems you experience with the drive. Some will require an error code or other status information before agreeing to declare a drive defective and issue an RMA for it.

- Often, the **diagnostic functions of these utilities are combined with other features for erasing the disk and are sometimes (incorrectly) called "low-level format utilities".**

Failure Modes and Failure Determination

- There are a large number of different ways that a hard disk can fail;
- The most "famous" way that a hard disk can go is via the infamous **head crash**. In reality, head crashes are the "airline disasters" of the hard disk world: they are dramatic, well-known and feared, but actually responsible for a small percentage of total drive problems.
- This short list will give you a flavor for the types of failures that occur:
- **Mechanical Failures:** These include component failures related to the mechanics of the disk. Problems of this type usually relate to the **spindle motor or bearings**, such as motor burnout, "stuck" bearings, excessive heat, or excessive noise and vibration. Actuator problems would also fit into this general category. Unsurprisingly, mechanical failures of hard disks comprise a large percentage of total problems.
- **Head and Head Assembly Failures:** The infamous head crash fits in here, as do other problems related to the heads: improper flying height, head contamination, defects in head manufacture, excessive errors on reads or writes, bad wiring between the heads and the logic board. These too comprise a large percentage of total failures.
- **Media Failure:** This class of problems relates to trouble with the platters and the magnetic media, formatting, servo operation and the like. This would include drives that fail due to read or write errors, poor handling, scratches on the media surface, errors in low-level formatting, etc. They are relatively uncommon.
- **Logic Board or Firmware Failures:** These are problems related to the drive's integrated logic board, its chips and other components, and the software routines (firmware) that runs it. Again, problems in this area of the disk are fairly uncommon compared to the other categories.

Failure Modes and Failure Determination

- Determining that your drive has a failure isn't always the simplest thing to do. Some failures manifest themselves rather clearly and obviously; others are quite a bit more subtle.
- Here are some of the most common signs that your hard disk *may* be failing--not definitely, but possibly.
- **Errors:** If the drive is returning read errors or "**sector not found**" errors, that is an obvious sign of a problem. Sometimes errors appear at a "higher level"; for example, if you experience repeated corruption of the Windows registry,
- **Changed Behavior:** If the drive has suddenly changed its behavior in a dramatic way due to nothing you (believe) you have done, this may be a bearer of bad tidings. This is especially true of noise: if the drive always made a particular noise then that noise is probably normal for the drive type But if the drive starts making *new* noises then you should be concerned. If the drive starts vibrating more than it used to, or spins up and down a lot more than before, these are also potential warning signs.
- **Scary Noises:** The above said, there are certain noises a drive should *never* make, even when new. If you hear something scraping as the drive turns on, or the drive makes repeated clunking sounds during operation, or emits a very high-pitched whine, then there may well be a problem.
- **Drive Not Recognized:** If the drive is not recognized at all by the BIOS of a motherboard to which it is connected *where it was recognized by that BIOS in the past*, this is a sign of a failed drive. (Non-recognition of a new drive can be due to BIOS compatibility issues.)
- **SMART Alarm:** If your hard disk and system support the SMART feature and you receive a SMART alert, consider this fair warning that the drive may have a problem.
- **Warning:** If you suspect that your hard disk is failing, *take action*. Back up your data immediately, and contact the drive manufacturer for instructions.

Hard Disk Quality and Reliability Features

- Hard **disk manufacturers** recognize the great importance of their customers' data, and therefore, the importance of the quality and reliability of their hard disk drives.

- In addition to generally improving the quality of their processes and their products, most manufacturers have added special features to their hard disks to help enhance their reliability.
- Most of these efforts are predicated on the principle that "an ounce of prevention equals a pound of cure". They are intended to either reduce the chances of a problem occurring, or at least let the user know if a problem is likely to occur, before data loss happens.

Hard Disk Quality and Reliability Features

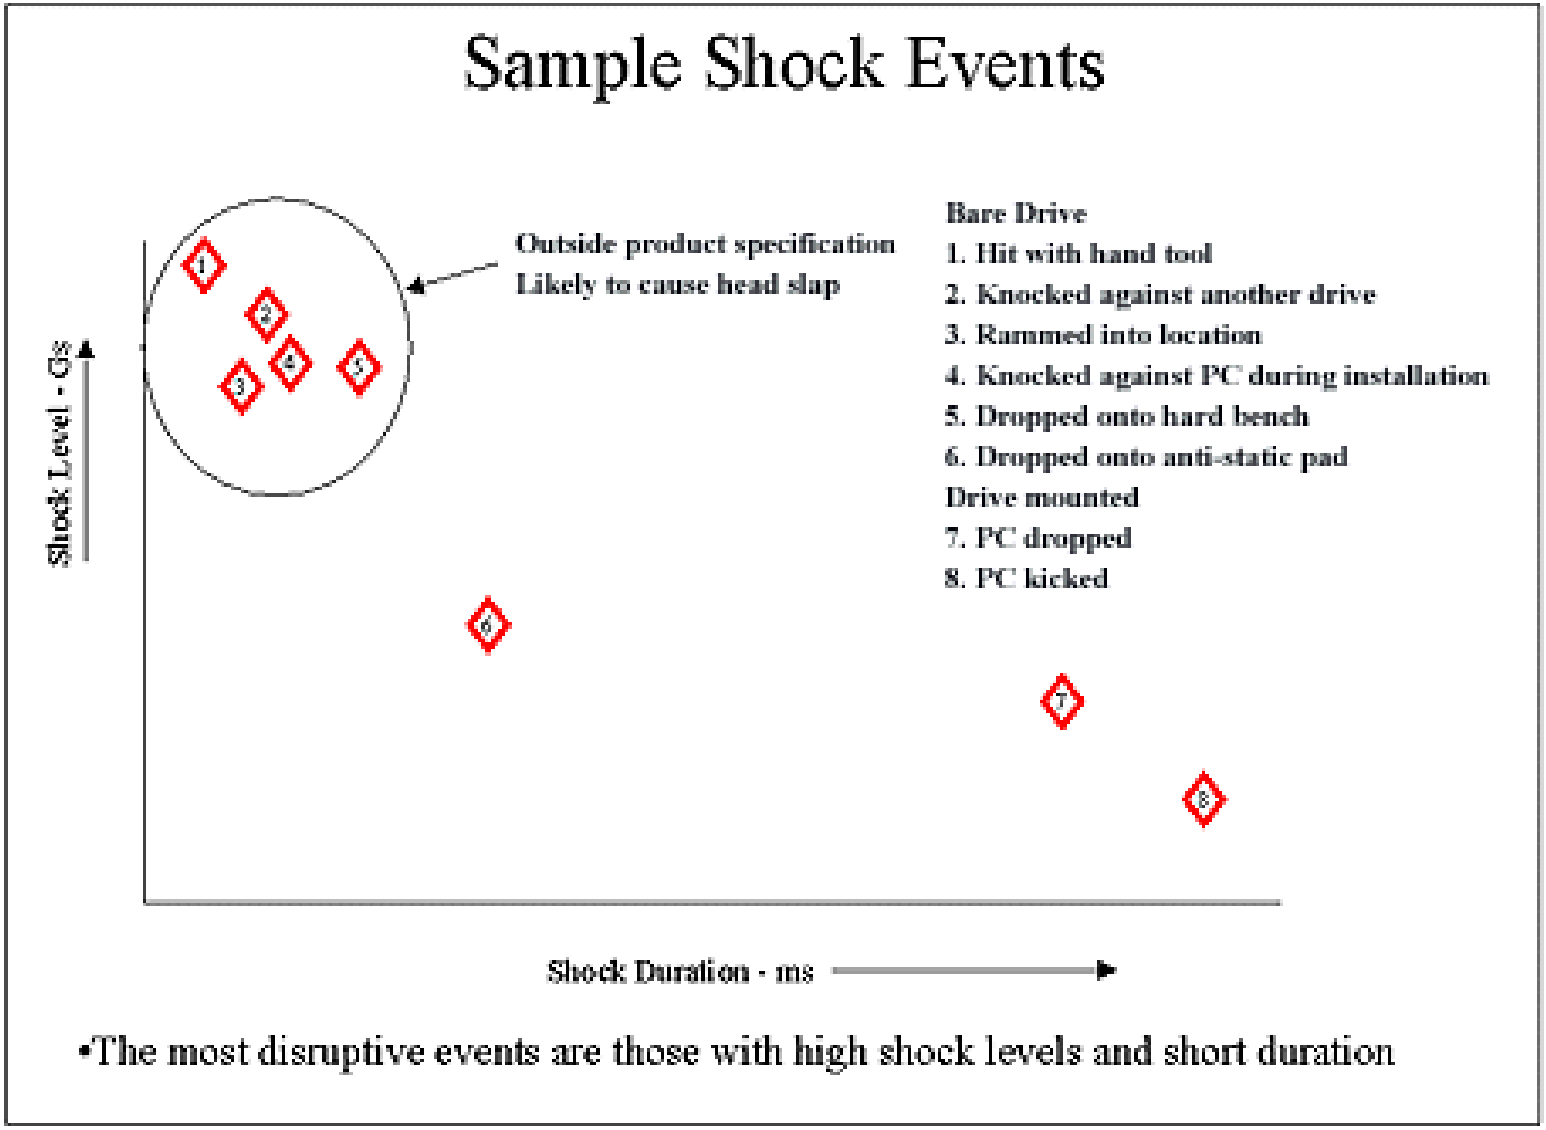

Hard Disk Quality and Reliability Features

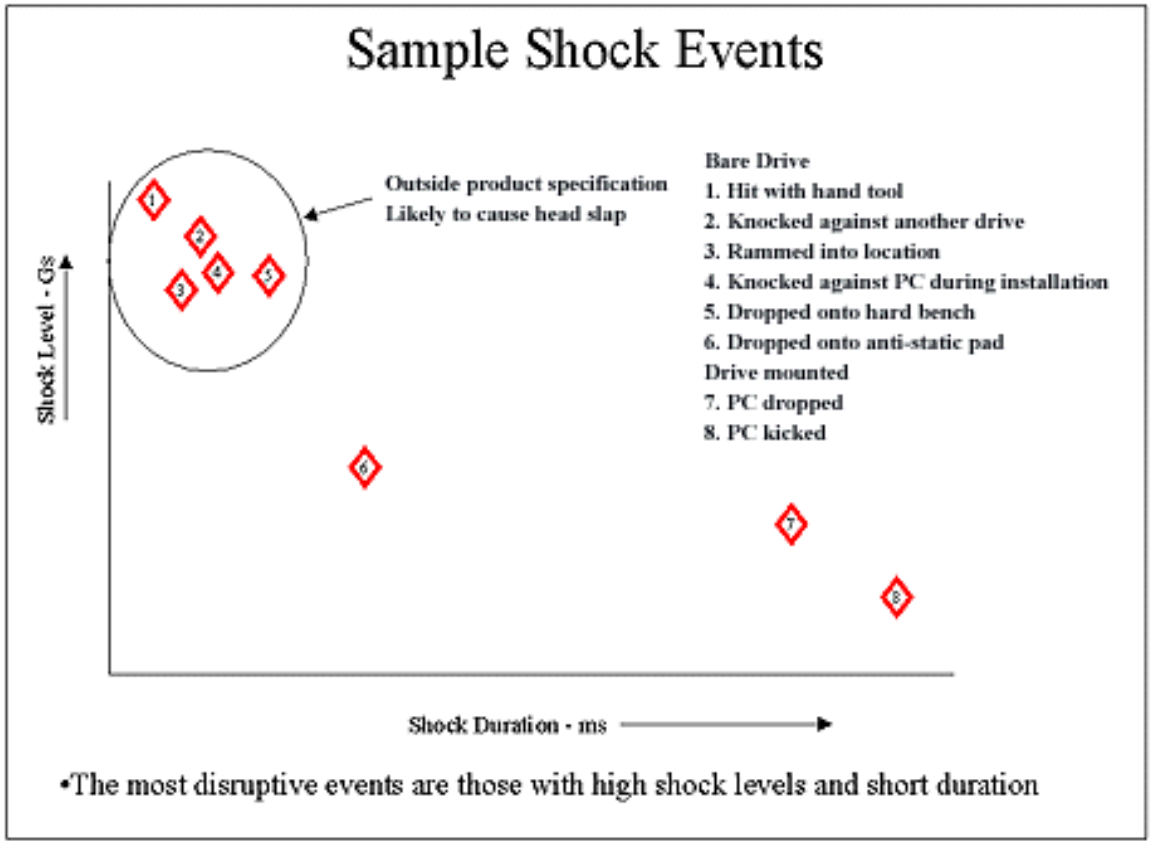

- This chart shows several sample shock events that can affect hard disks, mapped based on their duration and intensity. Notice that the more serious shocks are those delivered to the drive by itself, while those caused to the PC as a whole are less critical to the disk. Modern drives are equipped with many reliability features that can reduce the likelihood of damage in the event of common shock occurrences, but no advances will protect adequately against outright abuse of such delicate hardware.

Self-Monitoring Analysis and Reporting Technology (SMART)

- **In an effort to help users avoid data loss**, drive manufacturers are now incorporating logic into their drives that acts as an "**early warning system**" for pending drive problems. This system is called ***Self-Monitoring Analysis and Reporting Technology* or *SMART***. The hard disk's integrated controller **works** with various sensors to monitor various aspects of the drive's performance, determines from this information if the drive is behaving normally or not, and makes available status information to software that probes the drive and look at it.

- The fundamental principle behind SMART is that many **problems with hard disks don't occur suddenly**. They result from a **slow degradation of various mechanical or electronic components**. SMART evolved from a technology developed by IBM called *Predictive Failure Analysis* or *PFA*. PFA **divides failures** into **2 categories**: those that **can be predicted** and those **that cannot**. Predictable failures occur slowly over time, and often provide *clues* to their gradual failing that can be detected. An example of such a predictable failure is spindle motor bearing burnout: this will often occur over a long time, and can be detected by paying attention to how long the drive takes to spin up or down, by monitoring the temperature of the bearings, or by keeping track of how much current the spindle motor uses. An example of an unpredictable failure would be the burnout of a chip on the hard disk's logic board: often, this will "just happen" one day.

Self-Monitoring Analysis and Reporting Technology (SMART)

- The main principle behind failure prediction is that some failures cause gradual changes in various indicators that can be tracked to detect trends that may indicate overall drive failure.

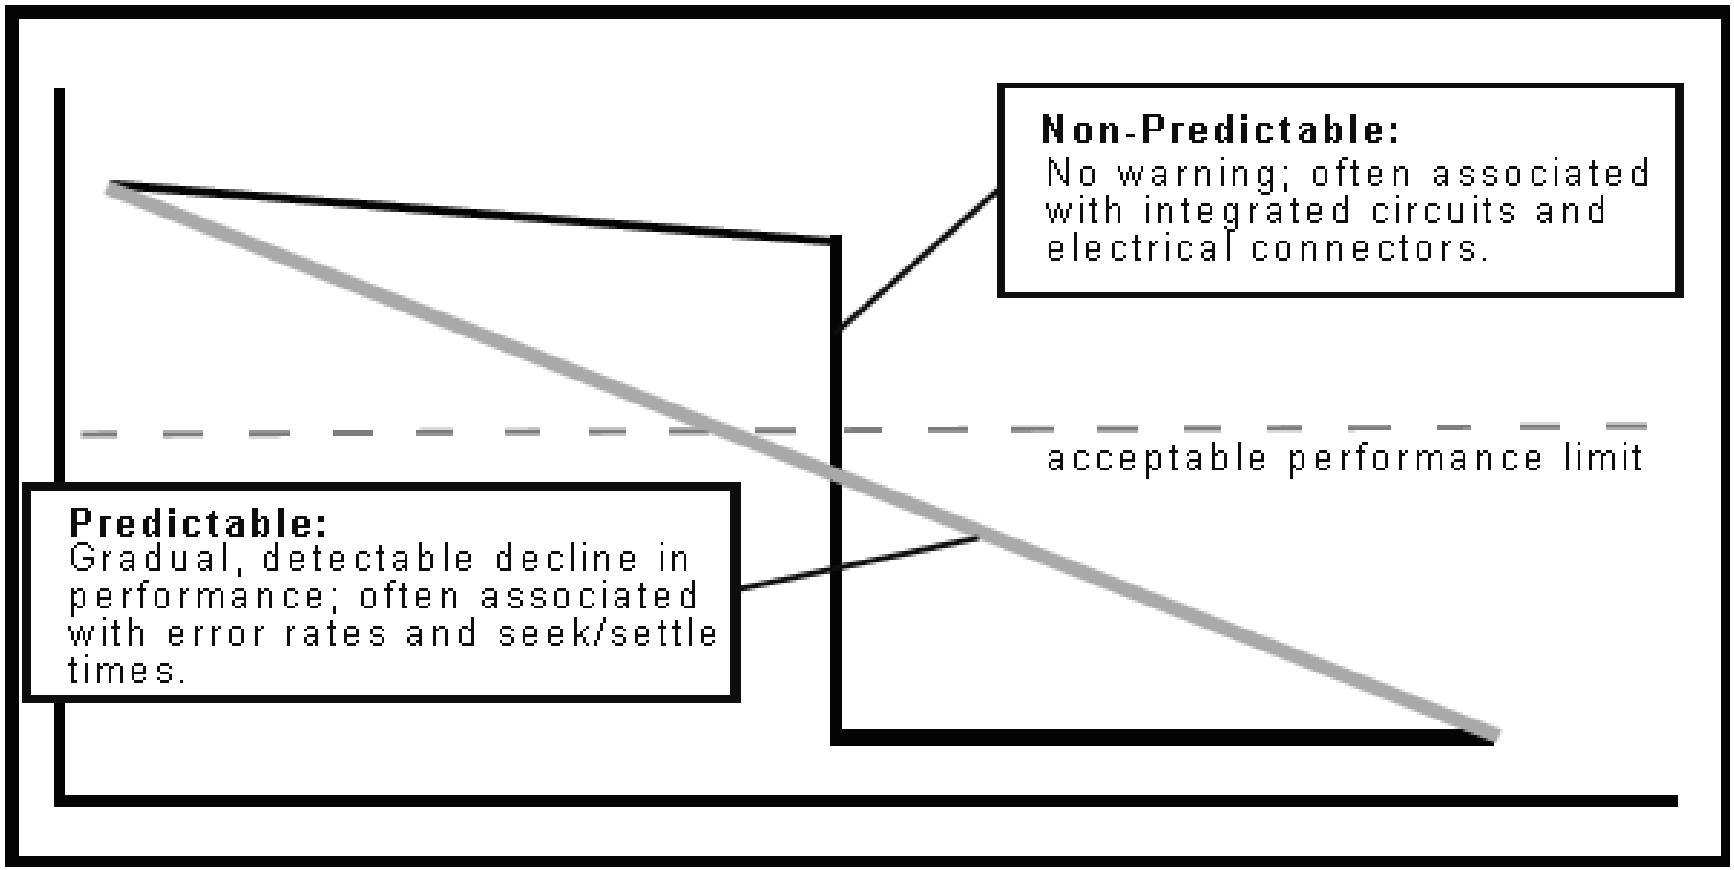

Self-Monitoring Analysis and Reporting Technology (SMART)

- The **drive manufacturer's reliability engineers analyze failed drives** and various mechanical and electronic characteristics of the drive to determine various correlations: relationships between predictable failures, and values and trends in various characteristics of the drive that suggest the possibility of slow degradation of the drive.
- The exact characteristics monitored depend on the particular manufacturer and model. Here are some that are commonly used:
- **Head Flying Height:** A downward trend in flying height will often presage a head crash.
- **Number of Remapped Sectors:** If the drive is remapping many sectors due to internally-detected errors, this can mean the drive is starting to go.
- **ECC Use and Error Counts:** The number of errors encountered by the drive, even if corrected internally, often signal problems developing with the drive. The trend is in some cases more important than the actual count.
- **Spin-Up Time:** Changes in spin-up time can reflect problems with the spindle motor.
- **Temperature:** Increases in drive temperature often signal spindle motor problems.
- **Data Throughput:** Reduction in the transfer rate of the drive can signal various internal problems.
- Using statistical analysis, the "acceptable" values of the various characteristics are programmed into the drive. If the measurements for the various attributes being monitored fall out of the acceptable range, or if the *trend* in a characteristic is showing an unacceptable decline, an alert condition is written into the drive's SMART status register to warn that a problem with the drive may be occurring.

Self-Monitoring Analysis and Reporting Technology (SMART)

- SMART requires a hard disk that supports the feature and some sort of software to check the status of the drive. All major drive manufacturers now incorporate the SMART feature into their drives, and most newer PC systems and motherboards have BIOS routines that will check the SMART status of the drive. So do operating systems such as Windows 98. If your PC doesn't have built-in SMART support, some utility software (like Norton Utilities and similar packages) can be set up to check the SMART status of drives. This is an important point to remember: the hard disk doesn't generate SMART alerts, it just makes available status information. That status data must be checked regularly for this feature to be of any value.
- Clearly, SMART is a useful tool but not one that is foolproof: it can detect some sorts of problems, but others it has no clue about. A good analogy for this feature would be to consider it like the warning lights on the dashboard of your car: something to pay attention to, but not to rely upon. You should not assume that because SMART generated an alert, there is definitely a drive problem, or conversely, that the lack of an alarm means the drive cannot possibly be having a problem. It certainly is no replacement for proper hard disk care and maintenance, or routine and current backups.
- If you experience a SMART alert using your drive, you should immediately stop using it and contact your drive manufacturer's technical support department for instructions. Some companies consider a SMART alert sufficient evidence that the drive is bad, and will immediately issue an RMA for its replacement; others require other steps to be performed, such as running diagnostic software on the drive. *In no event should you ignore the alert*. Sometimes I see people asking others "how they can turn off those annoying SMART messages" on their PCs. Doing that is, well, like putting electrical tape over your car's oil pressure light so it won't bother you while you're driving! :^)

Idle Time Error Checking

- To exploit this principle, some manufacturers are now incorporating routines to do ***proactive* error checking of the drive**. **During idle periods**, the drive's logic board performs reads over the surface of the disk to verify that the data can be read without errors. Any errors encountered are typically treated the same way an error would be addressed during a user read: the data may be **rewritten to "refresh" it,** or the sector where the error is relocated may be **remapped**. Drives that incorporate this feature usually integrate it with SMART; if an extensive number of errors occurs during such idle checks then this can signal problems with the drive.

- This is a useful reliability feature, but one that can sometimes cause confusion on the part of hard disk users. The reason is that this scanning activity can cause the drive to make "active" sounds even when the PC not doing anything involving the hard disk; the hard disk activity LED will also not be lit. (Since this scanning is purely an internal activity, it will not cause the hard disk controller or motherboard to light the LED.) If your drive is one that incorporates this feature and you see this "ghost" activity, there's nothing to be concerned about. There also should be no impact on the performance of the drive, since the scanning occurs only when the drive is idle.

Enhanced Shock Protection

- Several major types of drive failures can be caused by excessive shock applied to the drive, especially when it is in operation.
- Despite the fragility of hard disk technology, we do expect to be able to treat them, well, if not roughly, then at least without kid gloves. We want to put hard disks into laptop computers and carry them around, or use removable drive trays, or incorporate tiny hard disks into devices ranging from handheld PCs to digital cameras. And even with desktop machines, we certainly don't expect the drive to fail if we bump into the case while the PC is on, or similar minor mishaps occur.

- To address the increased needs for portability, and to improve reliability and reduce sensitivity to shocks in a general way, hard disk manufacturers have been hard at **work improving the shock resistance of drives**. One great way to do this would be to basically pack cushioning material around the drive; unfortunately, this would create other problems, ranging from needing to enlarge the drive's form factor, to making cooling more difficult. Therefore, shock protection is primarily done by increasing the rigidity of the most important components in the drive, such as the actuator assembly and spindle motor. To reduce head slap in particular, the head arms are made stronger. Structural support and some cushioning for the disk-head assembly is also provided to help reduce the transmission of shock energy from the outside of the drive to the platters and heads, where damage could occur.

Enhanced Shock Protection

- Anatomy of a head crash. Shock applied to the drive causes the heads to lift off the surface of the platters, then snap back causing damage. Newer drives using **shock protection** (which Quantum calls *SPS*) stiffen the head arms to greatly reduce the chances of head slap, and reduce the damage caused if it occurs.
  (Note that the **red** rectangles are actually the *sliders*, not the heads...)

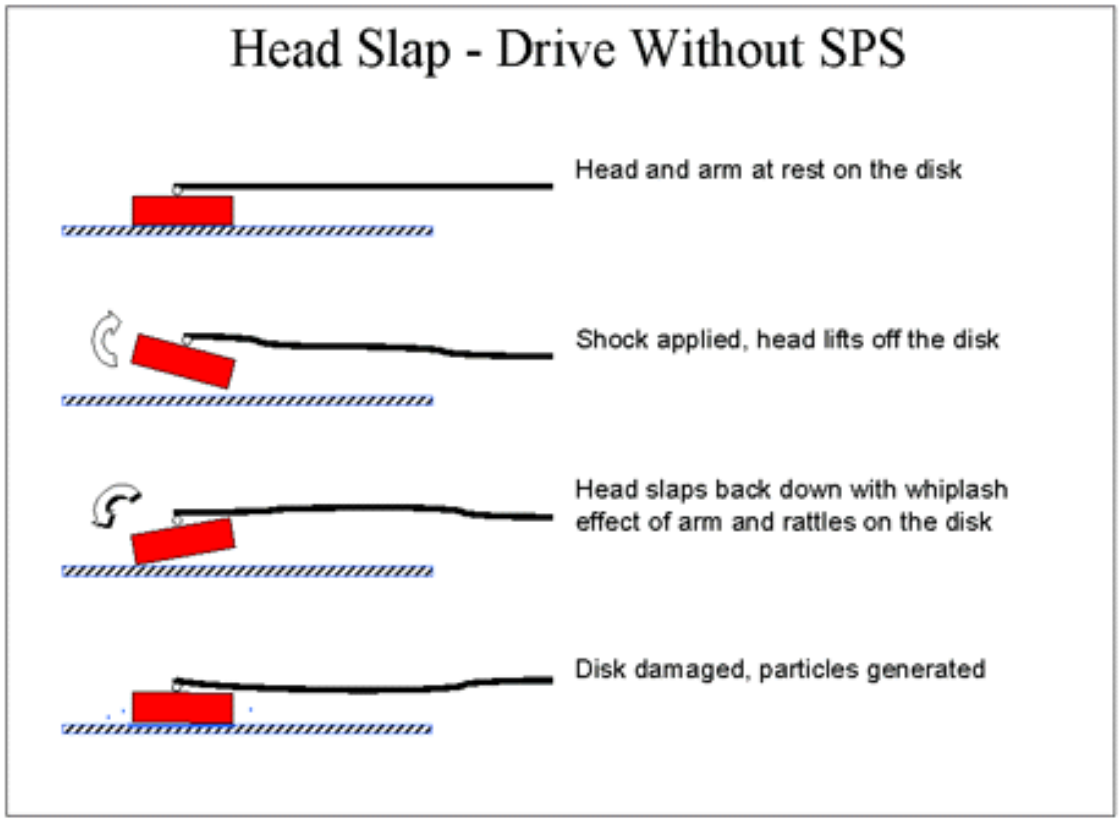

Fly Height Detection

- The **height** at which the heads fly above the hard disk platters is one of the most important design parameters of any hard disk. The reason is that this gap represents a critical design trade-off: if the heads are too high above the surface of the disk the heads then data errors can occur, but if they are too low, the risk of a head **crash dramatically increases**. The fly height of the heads must be maintained within tight parameters to ensure the reliability of the drive.
- Many manufacturers are **now incorporating sensors** into their drives to **monitor the height at which the heads** float over the surface of the platters. This is usually used as one of the data inputs to the drive's SMART feature, as any trend downward in fly height can signal the possibility of a head crash in the near future.

Wear Leveling

- Under normal circumstances--meaning, unless special reliability features are implemented--when a hard disk **drive is idle**, the head-actuator assembly just "sits there" wherever it happened to be when the **last read or write occurred**. In theory, there should not be any problem that results from this since the heads are floating on air and not touching the surface. However, at least one manufacturer (**Western Digital**) theorized that it might not be a good idea to keep the heads sitting in one spot for long periods of time even if in theory this was OK.
- To eliminate what they considered the potential for wear, Western Digital incorporated onto its drives a feature called *wear leveling*. Every fifteen seconds of idle time, the drive controller moves the heads by one track, traversing the entire disk area, so that no area of the disk has the heads sitting over it for very long. It takes about three hours for the entire surface of the disk to be traversed.
- The importance of this feature seems to still be under debate. In theory it improves drive reliability, but it also causes the drive to make more noise. It also makes some people think that there is something wrong with their drive.

Head Load/Unload Technology

- Each time the heads contact the surface of the platters, there is the potential for damage.
- Knowing that this will occur, manufacturers plan for it by incorporating special lubricants, and by ensuring that the heads normally contact the platters in a special landing zone away from user data areas. While obviously good planning, the entire way that starts and stops are handled--allowing contact with the platters--is not a great design. After several decades of handling this the same way, IBM engineers came up with a better solution.
- Instead of letting the heads fall down to the surface of the disk when the disk's motor is stopped, the heads are lifted completely off the surface of the disk while the drive is still spinning, using special ramps. Only then are the disks allowed to spin down. When the power is reapplied to the spindle motor, the process is reversed: the disks spin up, and once they are going fast enough to let the heads fly without contacting the disk surface, the heads are moved off the "ramps" and back onto the surface of the platters. IBM calls this *load/unload technology*. In theory it should improve the reliability of the hard disk as a whole.

Head Load/Unload Technology

- Diagram showing how IBM's load/unload ramp technology functions. (One head and surface is shown; there would be a different ramp for each head and surface the disk has.)

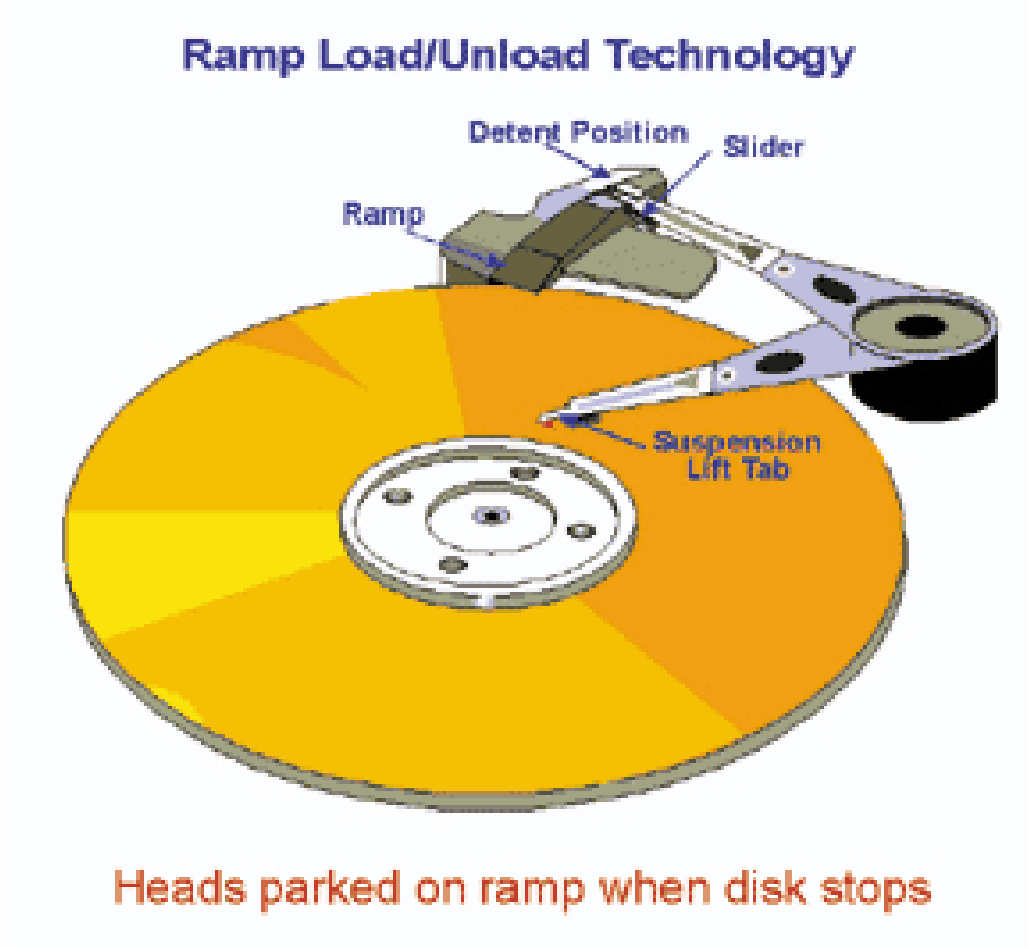

Temperature Monitoring

- Since heat has become much more of a concern for newer drives, particularly high-end ones, some manufacturers have added a very good reliability feature to their drives: *thermal monitoring*. **IBM calls this feature the *Drive Temperature Indicator Processor* or *Drive-TIP***. It has also been implemented by at least one other manufacturer (Western Digital).
- The idea behind this feature is very simple: a temperature sensor is mounted on the drive, usually on the logic board, and it records the temperature of the drive periodically over time. One or more *trip points* are set within the drive's control logic, and status notifications are sent by the drive back to the system if they are exceeded.

- **Normally, thermal monitoring is integrated with the drive's SMART feature for reporting**. A typical setup is two trip points, one at 60°C and another at 65°C; the first trip point can be changed by the user of the drive while the second often cannot. The controller may also keep track of the highest temperature ever recorded by the drive, and may also take action on its own accord if the trip point(s) are exceeded, such as slowing down the drive's activity or even shutting it down completely. The exact implementation depends on the manufacturer.

Expanded Remapping and Spare Sectoring

- When any hard disk detects a problem with a particular sector on the surface of the drive, it will remap that sector and mark it as "bad" so that it will not be written to again in the future. This ensures that repeat problems with this bad spot will not recur. This process is called *remapping and spare sectoring* and is described in this section on hard disk formatting.

- Some drives go a step beyond ordinary remapping. Instead of just remapping the sector where an error is encountered, or where a number of retries were necessary to get the data from the disk, the controller remaps a **zone of sectors around the defective location**. The logic behind this feature is that if an area of the disk is damaged badly enough to create a bad sector, the problem might not be limited to just that sector, even if the errors are only showing up there *right now*. After all, a single sector on a hard disk is *very* small; it stands to reason that there is a much higher chance that a sector near the defective one will go bad than a random sector somewhere else on the disk.

Expanded Remapping and Spare Sectoring

- The "**buffer area**" that is remapped can include sectors before and after the bad sector, as well as adjacent tracks, thus covering a two-dimensional space centered on the bad disk location. This process occurs transparently to the user. It makes use of spare sectors allocated on the drive for the purpose of remapping.

- **Western Digital** calls this feature ***defect margining***.

- As with the wear leveling feature, defect margining does not appear to have been universally adopted by other hard disk manufacturers.

Redundant Arrays of Inexpensive Disks (RAID)

- Most of the reliability features and issues discussed in this part of the site relate to making drives themselves more reliable. However, there is only so much you can do to improve the reliability of a single drive without the cost becoming exorbitant. Furthermore, since most people aren't willing to pay for ultra-reliable drives, manufacturers have little incentive to develop them. For those applications where reliability is *paramount*, the quality of no single-drive solution is sufficient. For these situations, many businesses and power users are increasingly turning to the use of multiple drives in a redundant or partially-redundant array configuration. The common term that refers to this technology is *Redundant Arrays of Inexpensive* (or *Independent*) *Disks*, abbreviated *RAID*.
- The principle behind RAID is "belt and suspenders": if you store redundant information across multiple disks, then you insulate yourself from disaster in the event that one of the disks fails. If done properly, you also improve performance--sometimes in a substantial way--by allowing the drives to be accessed in parallel. And you can make it so bad drives can be replaced without the system even being taken down.
- RAID is a big topic unto itself; there are many different ways that RAID can be implemented; various hardware and software considerations; and many tradeoffs to be considered when implementing a system.

Data Security

- Most people have personal information on their hard disks, and some worry about **what happens to this data when they send a drive in to a manufacturer for warranty replacement** (or to a data recovery company for that matter). If your hard disk fails, most of the time almost all of the data is still intact on the platters. A burned-out spindle motor or flaky actuator has no impact on the data stored on the disk. Even a head crash or other defect related to the data surfaces will result in little of the actual data being removed from the drive.
- If you are worried about the data on a disk that is totally dead, there is really nothing you can do to get the data off the drive that won't also void the warranty of the drive. If you are worried about the data on a drive that is acting "funny" but has not totally failed, you can wipe the data off the drive before sending it back to the manufacturer by using a **zero-fill utility**. This will for all intents and purposes eliminate all data from the drive. There have been "James Bond" like reports of experts retrieving data after a zero-fill has been performed, due to trace magnetic fields on the disk; even if true, someone would have to be *remarkably* motivated to even bother with such an exercise. Unless you're a master spy, the contents of your hard disk probably aren't nearly interesting enough for anyone with the skill to do something like this to even be bothered spending their time. And most of the *real* "cloak and dagger" stories about genuises retrieving data from disks that have had their data overwritten with various data patterns a dozen times or been blasted with shotguns are probably apocryphal. It is true that data recovery companies can do amazing things, but they aren't miracle workers. :^)

Data recovery

- You should keep in mind that hard disk technicians are professionals. A hard disk company cannot afford to have a reputation for snooping around in their customers' files--and a data recovery company can afford such a black eye even less. I personally have never heard of even one instance of a hard disk manufacturer or data recovery company going through someone's "stuff" and causing problems, so either it isn't happening, or nobody is talking about it. Remember that most of the data that is so important to you, is important *only* to you. This is especially true of most PCs used for personal purposes, as opposed to businesses.
- All of the above said, the bottom line of data security is that it is only truly secure if it never leaves your hands. If you're really worried about data security, if you truly can't afford *any* chance of the data on the drive falling into the wrong hands because it contains the entire company's payroll information or your almost-complete plans for a cold fusion reactor, then you have an alternative: don't send it in for warranty service. Just buy a new drive. After all, hard disks only cost a few hundred dollars, and if the data is *that* important, it's worth eating the cost of a hard disk and just buying a new one. This is in fact exactly what some businesses do. The old drive should be destroyed or just put in the back of a safe somewhere. If you have access to a smelter, well, that even worked on the Terminator. :^)

Data Recovery

- Hard disks differ from virtually every other component in the PC in one very critical respect: when you lose a hard disk you lose more than just hardware--you lose software and data as well. The hard disk can be replaced, but the data on it often cannot (and it certainly isn't covered by your warranty). This is why annoying people like me always harp on users to back up their data.
- If for some reason you fail to heed my sage advice :^) and do not back up your important data, it is still sometimes possible to recover it in the event of a hard disk failure. There are data recovery services that specialize in restoring lost data from hard disks that have either failed or been corrupted due to software problems (accidental formatting, viruses, etc.)
- Some of these services are pretty amazing--many of them have a very good success rate and can resurrect data even when it seems like it would be impossible to the casual observer. They use specialized equipment and software, and typically even have mini clean rooms that enable them to take the cover off the drive and access the platters directly. All this wizardry comes at a cost, of course. A full recovery usually starts at a few hundred dollars and proceeds from there. The cost is high for three reasons: the equipment is expensive, the people are highly skilled, and the company knows how valuable the data is or you wouldn't be talking to them. Compared to recreating a year's worth of data, $2,000 for a recovery is a bargain. Compared to doing routine backups, it's a colossal waste of money.
- One company well-known for performing data recovery is Ontrack.

Data Recovery

- This graph shows the most common causes of data loss requiring the use of data recovery services, based on actual Ontrack recoveries in the 1995-1996 timeframe.

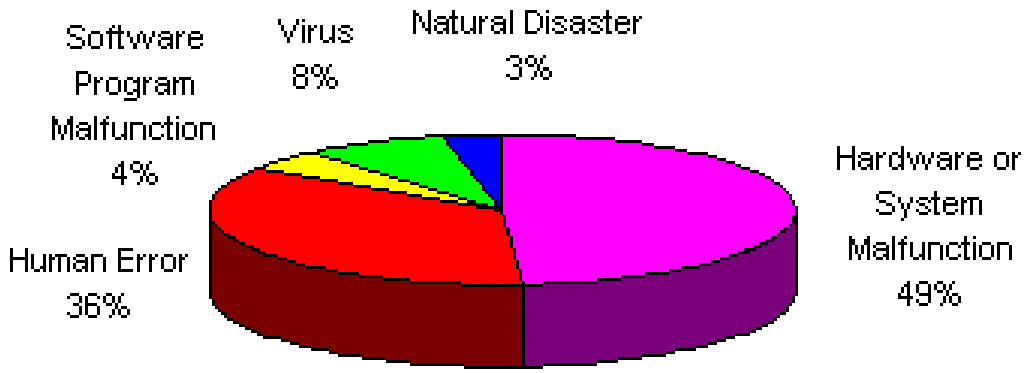

- **Note:** Even though drives have labels and stickers on them saying that if the drive is opened the warranty will be voided, this does not apply to professional data recovery. These companies have arrangements with the major hard disk manufacturers to allow them to work on the drives without affecting the drive's warranty status. In some cases the recovery company may be required to send a letter to the drive manufacturer certifying that your particular drive had recovery performed upon it.