How Solid State Drives Work - Research and Information

mhdd My Hard Drive Died

By Scott A. Moulton / www.MyHardDriveDied.com

How Solid State Hard Drives Work – Research and Information

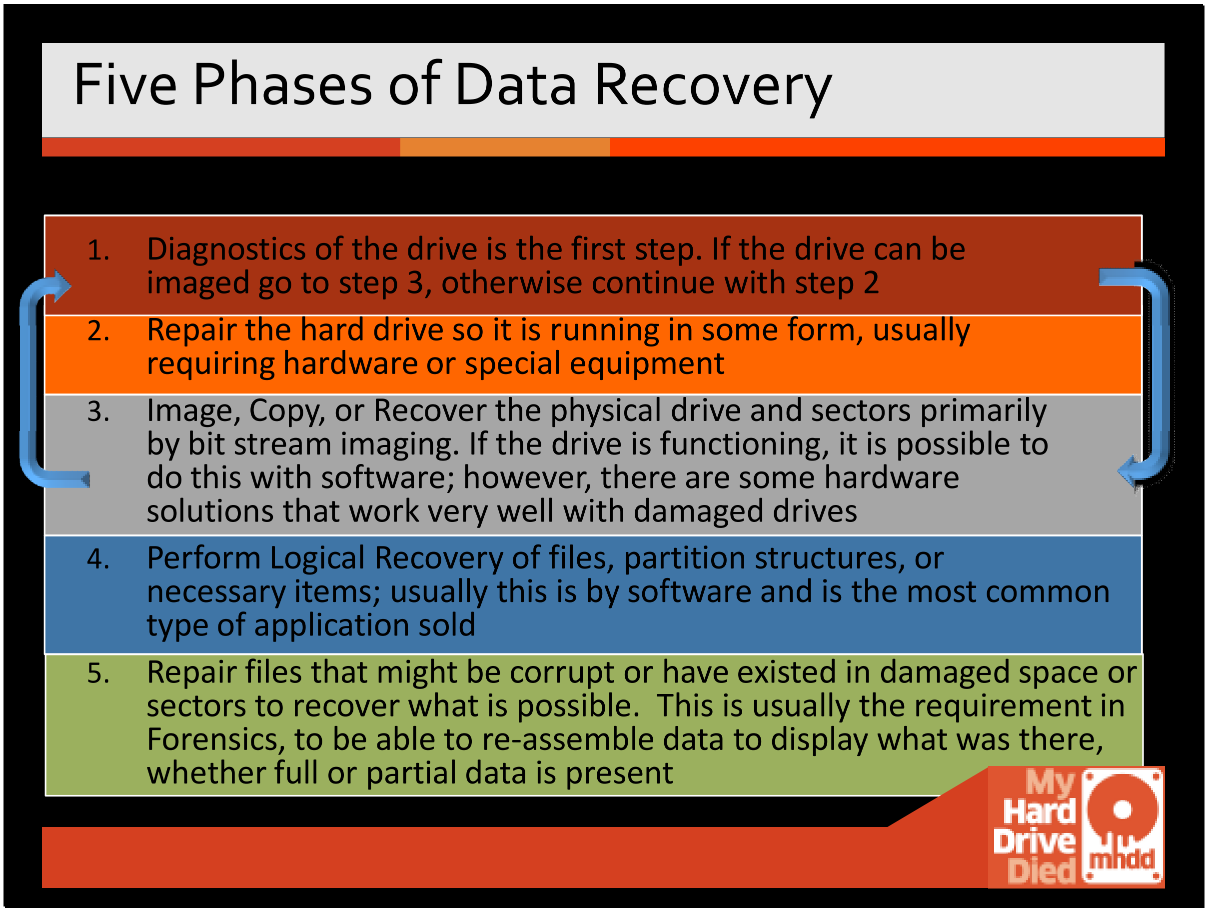

Depending on the type of data recovery, you will have five phases. This is the same regardless of Hard Drives or Solid State Drives. The Difference is what you have to do to fix the Solid State Drives compared to the repair process for mechanical hard drives. You can diagnose the problem many ways, but one of the easiest is to attempt to image the drive using hardware or software. There are some pieces of software that can talk to the drive and help you diagnose the type of problem before continuing. About 15% of the time the second step is going to be repairing a damaged hard drive. The other 85% will be imaging the drive in step three in some capacity. Until you have done this you generally will not see any data and that prevents you from doing the other steps. In essence, if you cannot repair the hard drive, that is the end of your job.

The third phase of data recovery is imaging. Not all data recovery companies image their drives. I generally find it to be very valuable to image the drive and work from the working copy. This allows me to maintain the state of the original drive without making changes to it, because as most of you know, making changes to the original drive is a bad thing. Let's say I have a bad drive that came from a Mac, and I hook it up to the Macintosh, and run a tool like disk warrior against it. Now, if disk warrior tries to make changes to catalog, you will have no resource to go back to, and in many cases you will not know what type of damage you have to the drive until it is too late. In some cases, if the drive is damaged you might just get the one shot, so you have to be very cautious not to waste your one shot.

I believe that you should make a physical image or clone of the drive before you proceed to the third step, logical recovery. There are several tools specifically for dealing with damaged media. You might consider using software-based tools on Linux such as dd_rescue (use with dd_rhelp script) or ddrescue (dd_rescue and ddrescue are two different tools). These tools have a special feature that allows them to image backward (understanding why you need to image backwards is very important in data recovery). There is also some excellent hardware for doing imaging from physically damaged hard drives ,such as the DeepSpar Disk Imager.

And then the third step is the logical recovery portion where you're repairing partition structures or corruption and software. This is what the most common types of data recovery software do. This is what the most common types of data recovery software do. They work by talking to the operating system ,making requests and expecting a response. In most cases, they have no idea how to talk to damaged media or corrupted files and can cause them to fail. Most of the software in this category is defined by how good the headers are in the application. Most of them began by scanning the hard drive from beginning to end, examining header information and trying to determine what applications or data files are on the drive. Some can parse the MFT or fat table or whatever catalog the operating system of the damaged drive is, but others only scans for file headers or partitions structures. We will cover some of the most common applications and what your options are for repairing some of this corruption. Most of this focus will be on automation primarily, because in a data recovery lab you're handling dozens of drives a week and just would not have the time to do everyone by hand. So this is about performance. And again keep in mind, they could still take weeks to run through just this portion of the recovery and validate the results.

So the fourth phase to data recovery is after you have recovered the files. In many cases, especially if you had to repair the media, or there were bad sectors and damage to the drive, there may be corrupt files after you have recovered them. This is typically where I look for tools which I call one offs. Generally that means for whatever type of file it is , as time goes on particular tools get better at recovery or new tools appear on the scene. In most cases this means that the same tool that I used six month ago may not currently be the best tool to use to recover that file. In some cases it may be even possible, just use a hex editor and cut out the information that you want. But generally speaking, this is just whatever the tool of the day is. We will cover much of what that is in this class because our focuses on the media can image and doing logical recovery (reword for clarity. The confusion starts at the "…focuses on the media" and continues through the remainder of the sentence. Repairing individual files is something that you have to take on one at a time.

This is also applied in data recovery for corrupt Word and Excel documents.

This is all very confusing, Scott. You’ve already covered the third step on the previous page, but then you go into the third and fourth on this page, but they seem more related to steps 4 and 5. Is this simply an error? Should they in fact be steps 4 and 5?If not, where is the explanation for step 5?

Quoting Microsofts Blog on the advantages of TRIM.
http://blogs.msdn.com/e7/archive/2009/05/05/support-and-q-a-for-solid-state-drives-and.aspx

"Microsoft and SSD manufacturers are adopting the Trim operation. In Windows 7, if an SSD reports it supports the Trim attribute of the ATA protocol's Data Set Management command, the NTFS file system will request the ATA driver to issue the new operation to the device when files are deleted and it is safe to erase the SSD pages backing the files. With this information, an SSD can plan to erase the relevant blocks opportunistically (and lazily) in the hope that subsequent writes will not require a blocking erase operation, since erased pages are available for reuse. As an added benefit, the Trim operation can help SSDs reduce wear by eliminating the need for many merge operations to occur. As an example, consider a single 128 KB SSD block that contained a 128 KB file. If the file is deleted and a Trim operation is requested, then the SSD can avoid having to mix bytes from the SSD block with any other bytes that are subsequently written to that block. This reduces wear. Windows 7 requests the Trim operation for more than just file delete operations. The Trim operation is fully integrated with partition- and volume-level commands like Format and Delete, with file system commands relating to truncate and compression, and with the System Restore (aka Volume Snapshot) feature."

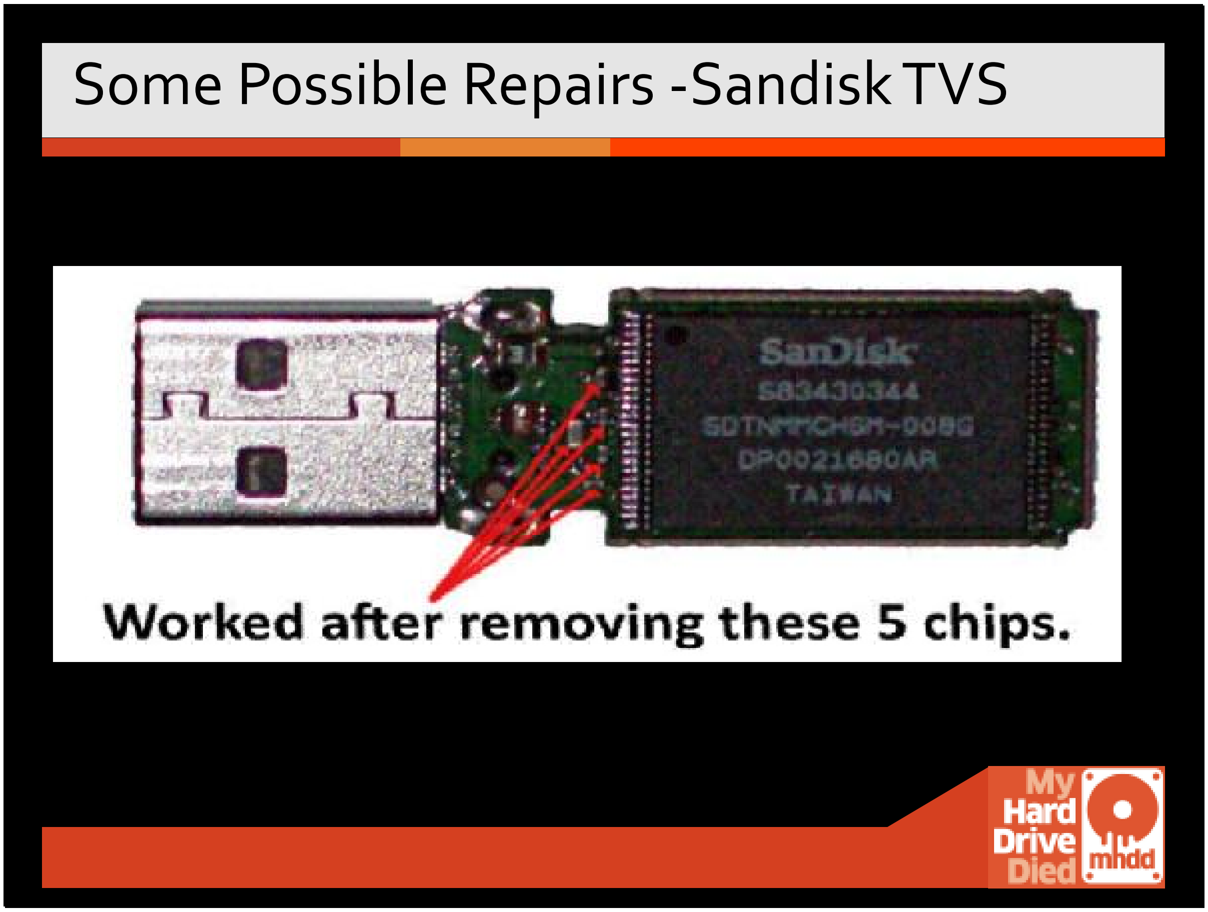

On a hard drive there are chips called TVS Chips (transient voltage sensors) and when they trip the drive stops working to protect itself from power overloads. Same is true of Solid State Drives and USB Thumb Drives. The above is an example.

Quote from a User

Basically the USB Memory stick in question just got hot when plugged in. The PC did nothing.. so basically it was a nice heater, but not good for storing data. I bought 2 new Sandisk Cruzers of the same model number at $15.00 each.

1. I measured those tiny little resistors that were present on the board, and compared with the ones on the new unit
2. All of the resistors on the bad unit were measuring as a completely closed circuit.
3. All of the resistors on the good unit were reading OPEN for resistance!

4. Oddly enough, the new unit - despite having the exact same model number - was missing 3 of
 the 5 resistors.
5. Conclusion... SNAP OFF ALL THE RESISTORS!!

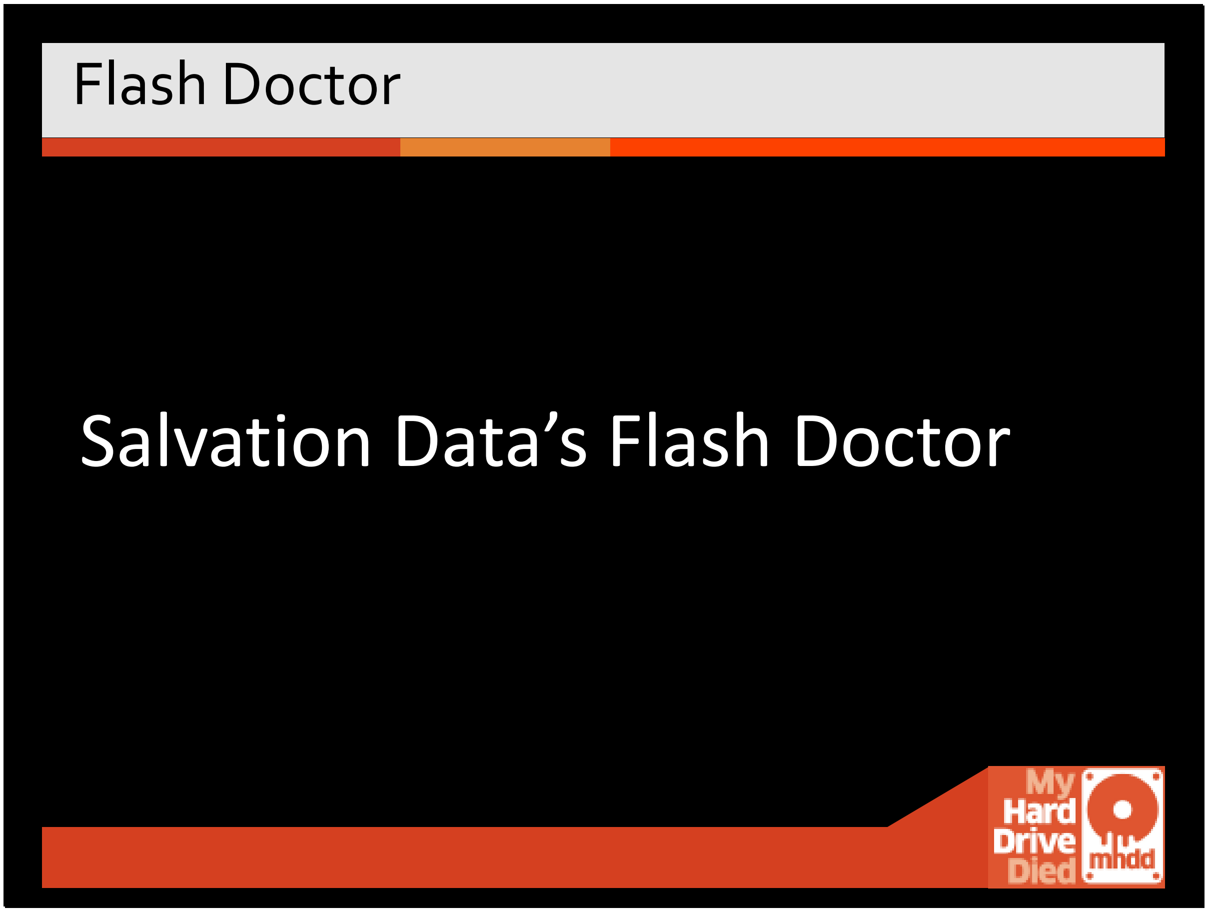

These are going to be examples of what software looks like when you are repairing Chips that were imaged from Drives.

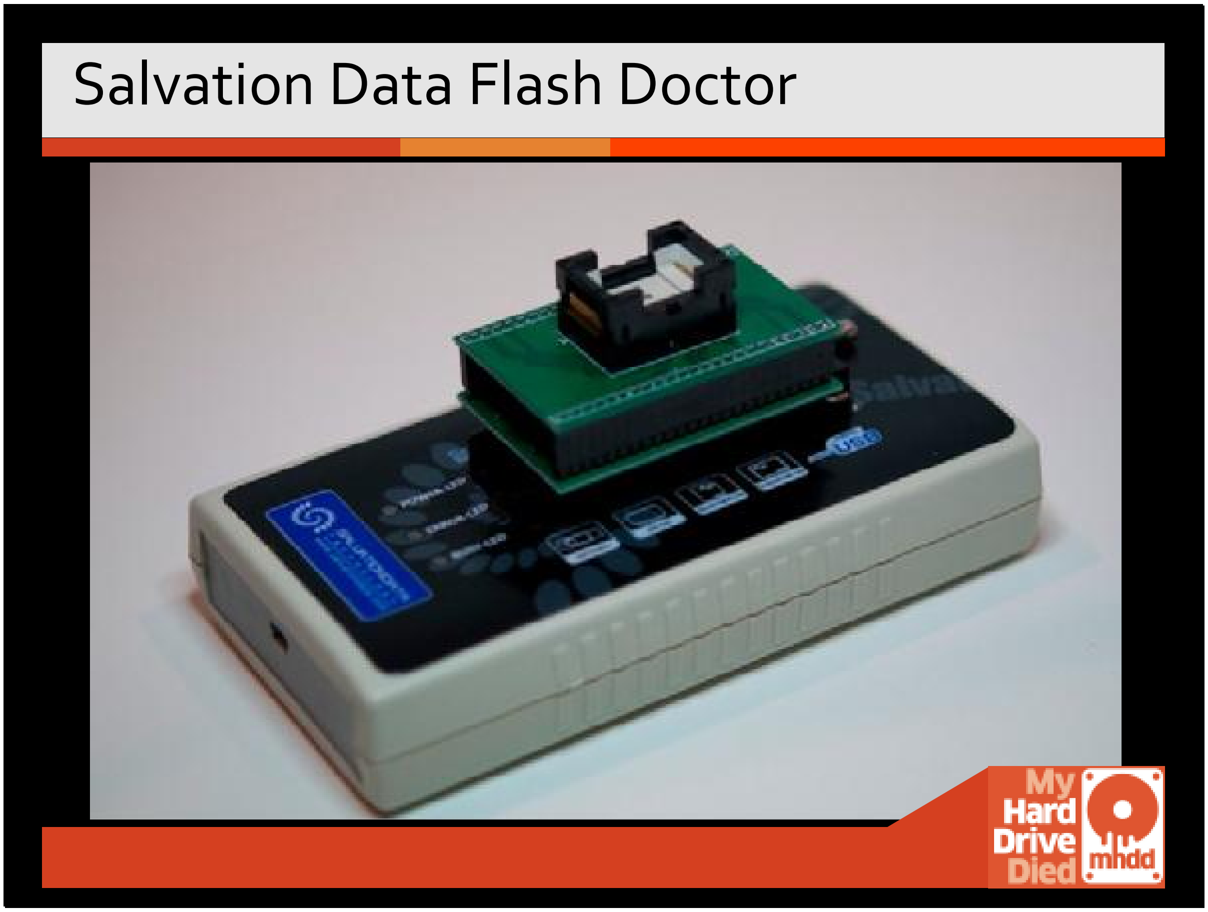

This is Salvation Data's Flash Doctor. This is used for Solid State Repair. You remove the chips, and then image them, and use software to reinterpret what the content was on the device.

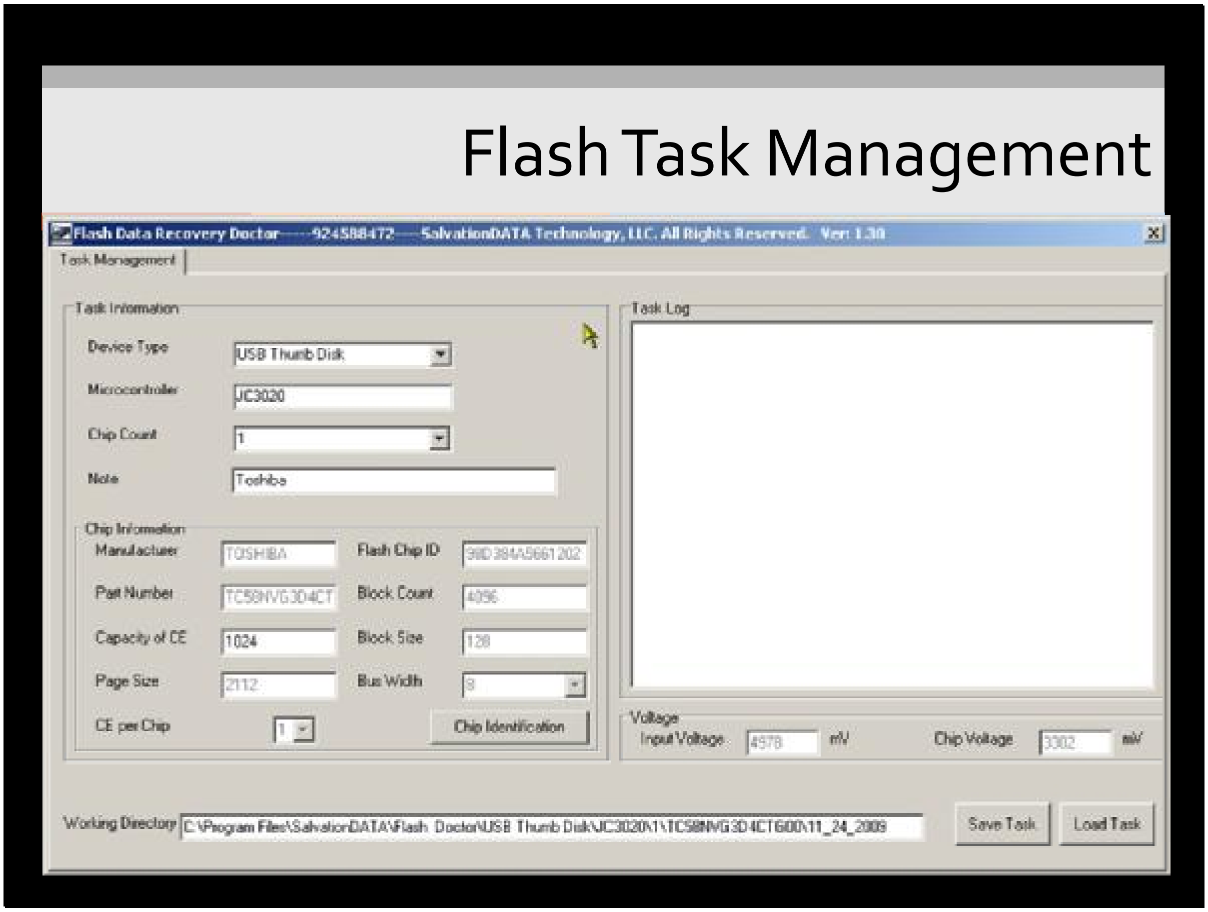

These are going to be examples of what software looks like when you are repairing Chips that were imaged from Drives.
Salvation Data's Flash Doctor Flash Management Task Screen.

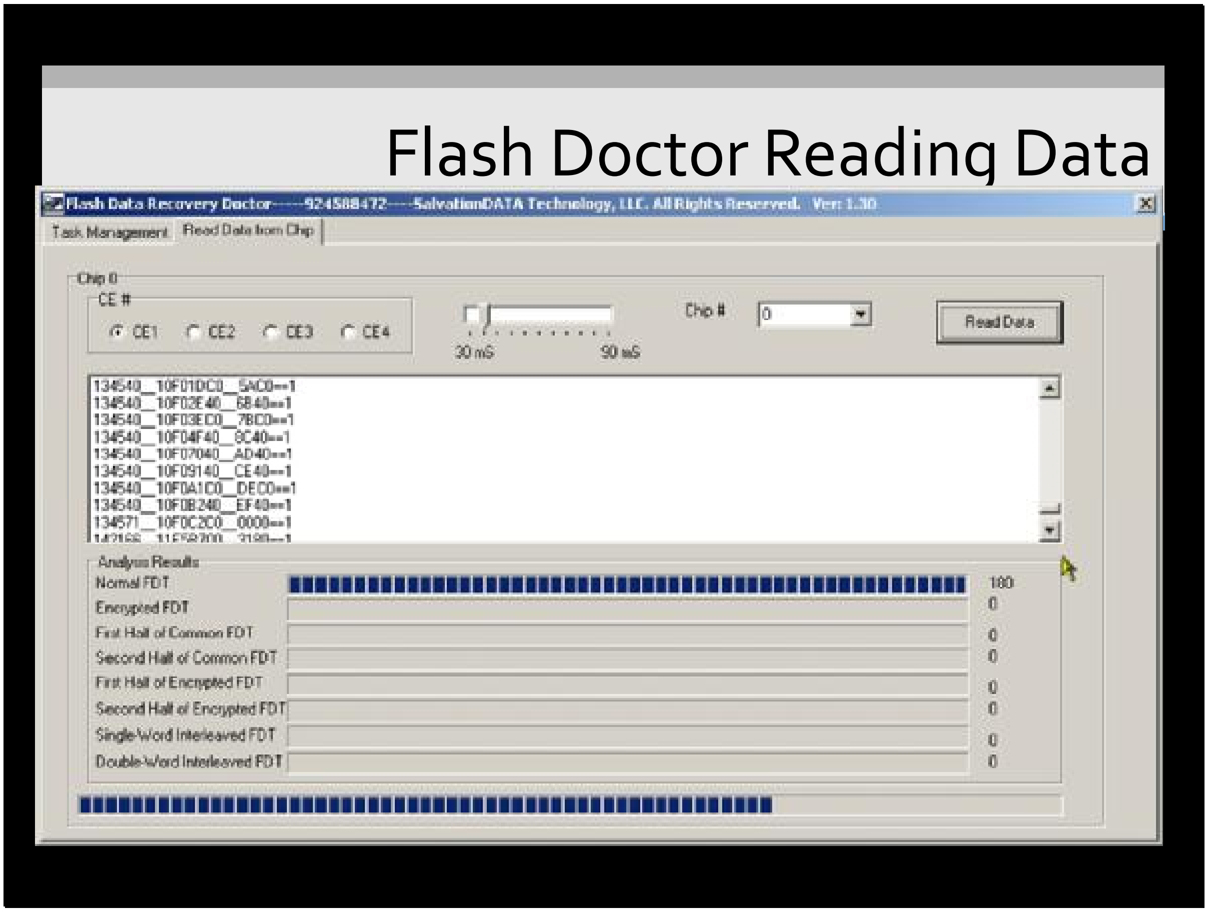

These are going to be examples of what software looks like when you are repairing Chips that were imaged from Drives.

Salvation Data's Flash Doctor Flash Reading Data from Screen.

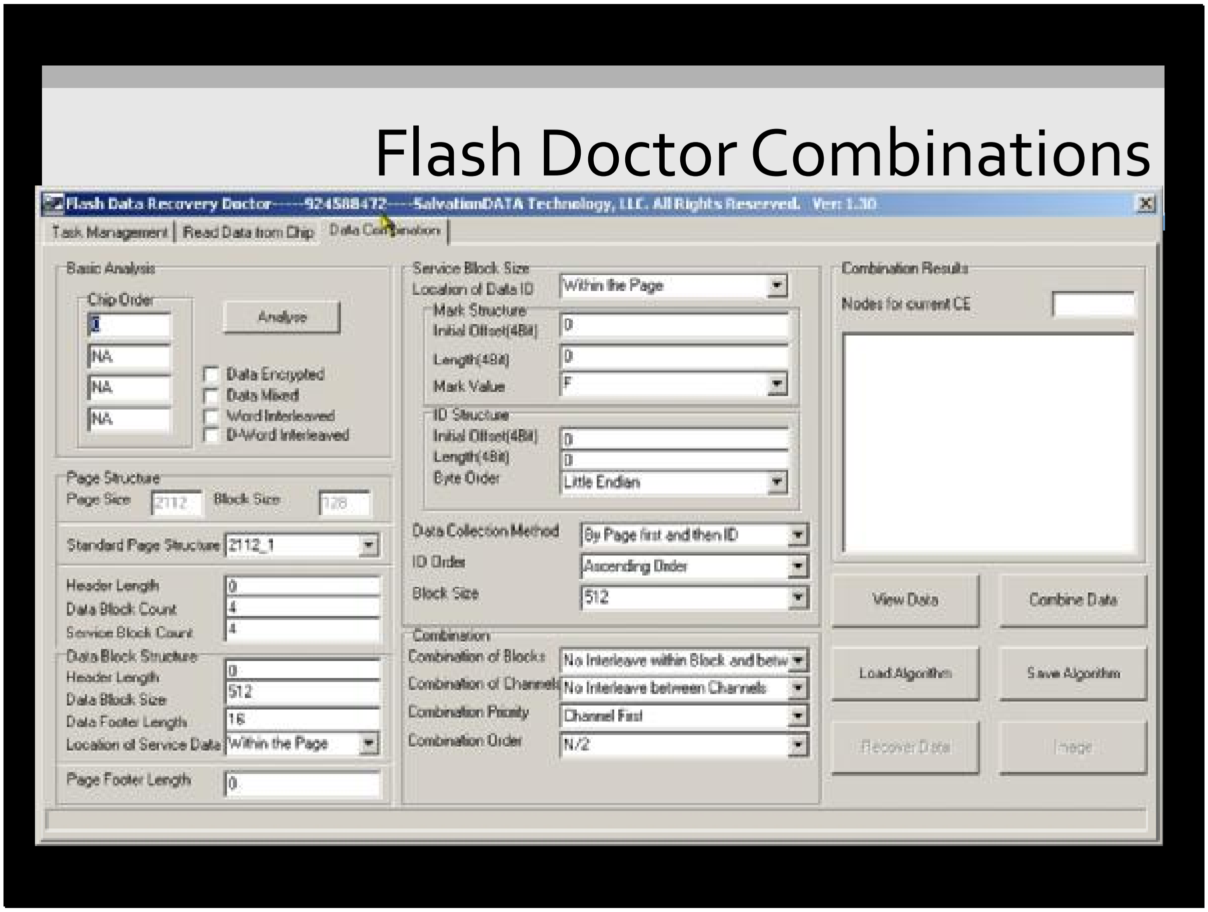

These are going to be examples of what software looks like when you are repairing Chips that were imaged from Drives.
Salvation Data's Flash Doctor Flash Management Data Combinations.

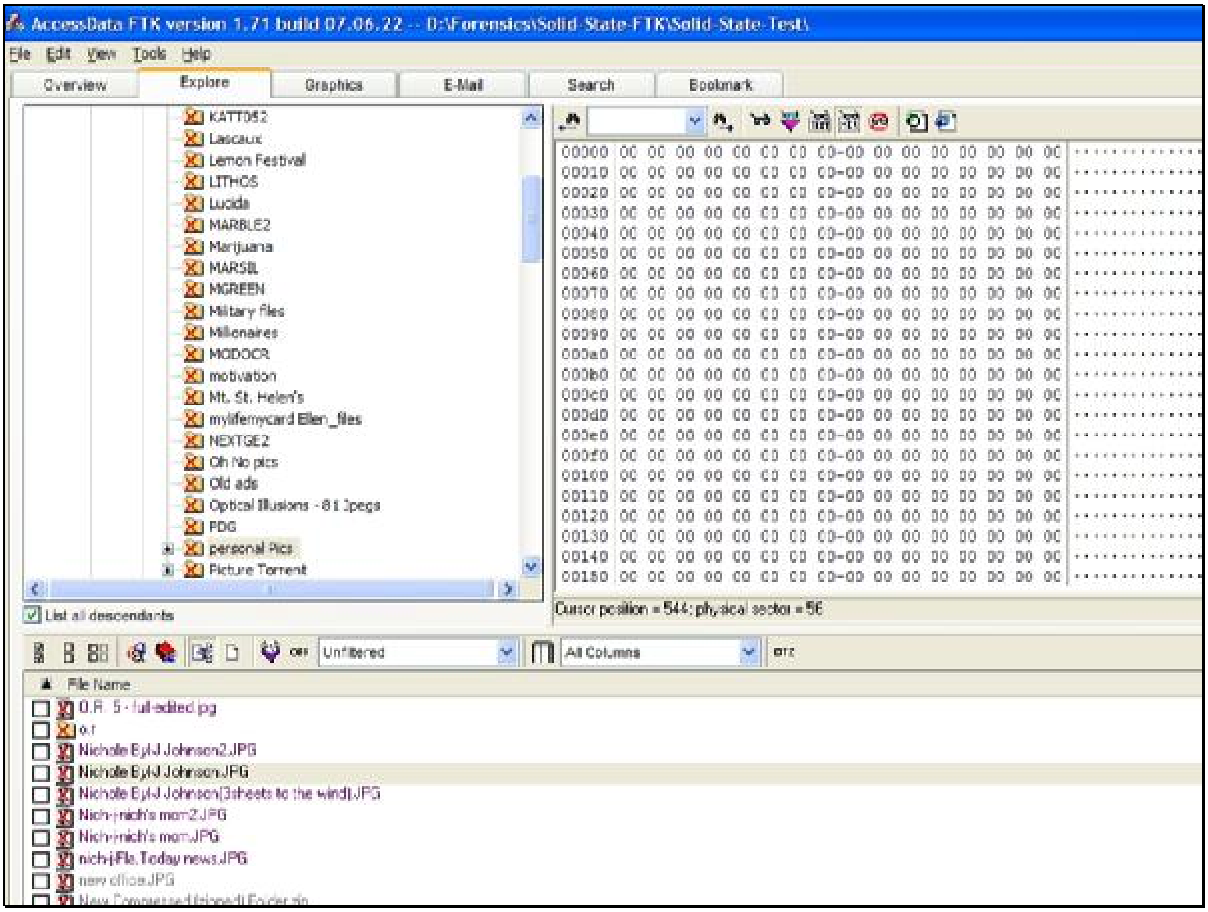

One of the items that got me started looking at solid state drives was when I had a forensic case in 2007 that had the first solid state drive I had ever seen. The first thing I noticed was that the drive looked wiped. I ran a number of tests and found out that the data in the slack space and in the unallocated space went to zero, even though the MFT contained data that looked like there should be something. This was a sample of some tests I ran over several months where the data would appear to be wiped.

All vendors are trying to make the Solid State Disk the standard, and they want it to be their own.

- Ex: When first came out a 32 gig was $2,000.
- Year ago $1200.
- 64 gigs for $1200. one year ago
- 128 gigs for $500.

The first SSD's that came out tend to have problems not very long afterwards. Example: When you are watching a movie certain drives may stutter.

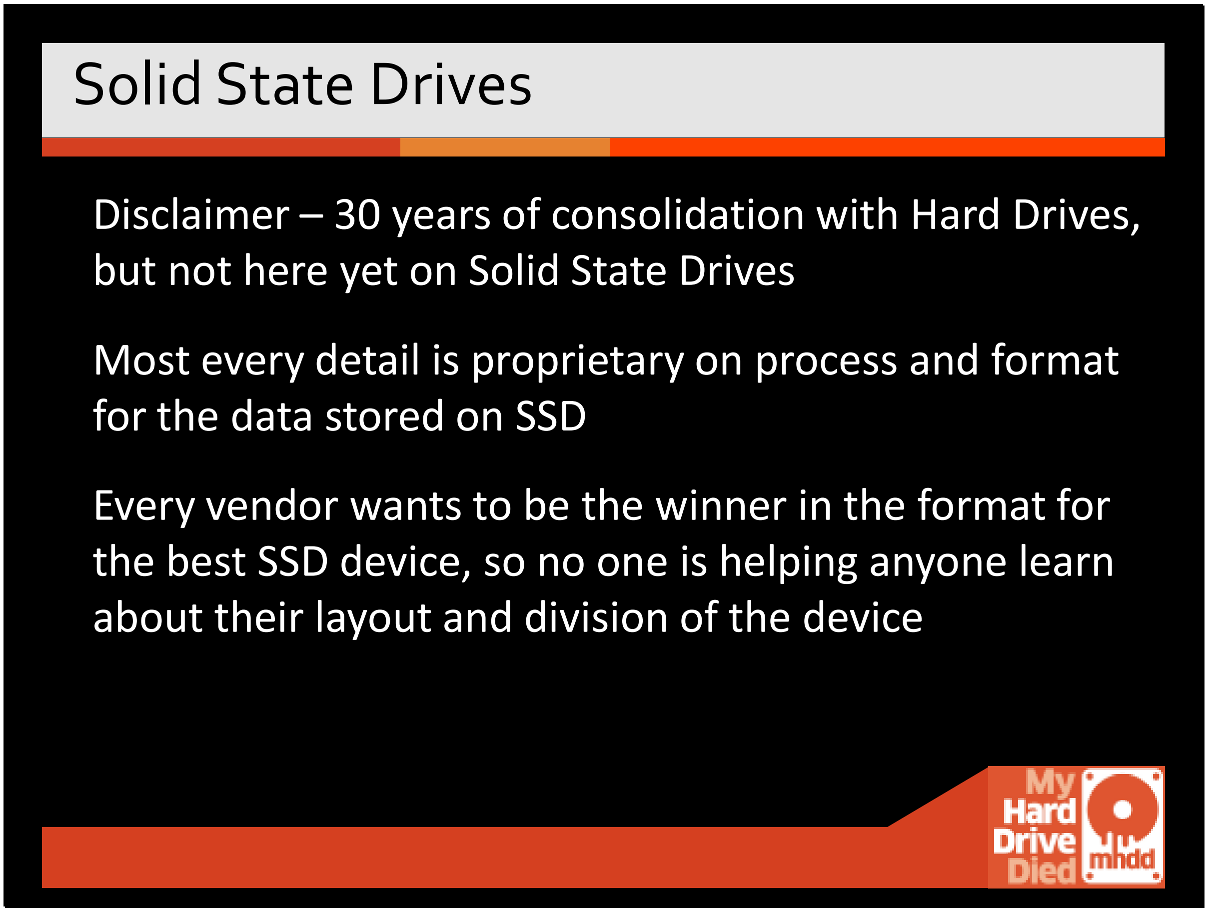

Solid State Drives emulate what the IDE does and plugs into an IDE/SATA/ZIF sockets. Zif sockets are very fragile and can only be used about 20 plug-ins before they stop working.

Most SSD manufacturers use non-volatile flash memory to create more rugged and compact devices for the consumer market.
SSD's do not require batteries. Often packaged in standard disk drive form factors (1.8-inch, 2.5-inch, and 3.5-inch).

Non-volatility allows flash SSDs to retain memory even during sudden power outages, enduring data persistence.

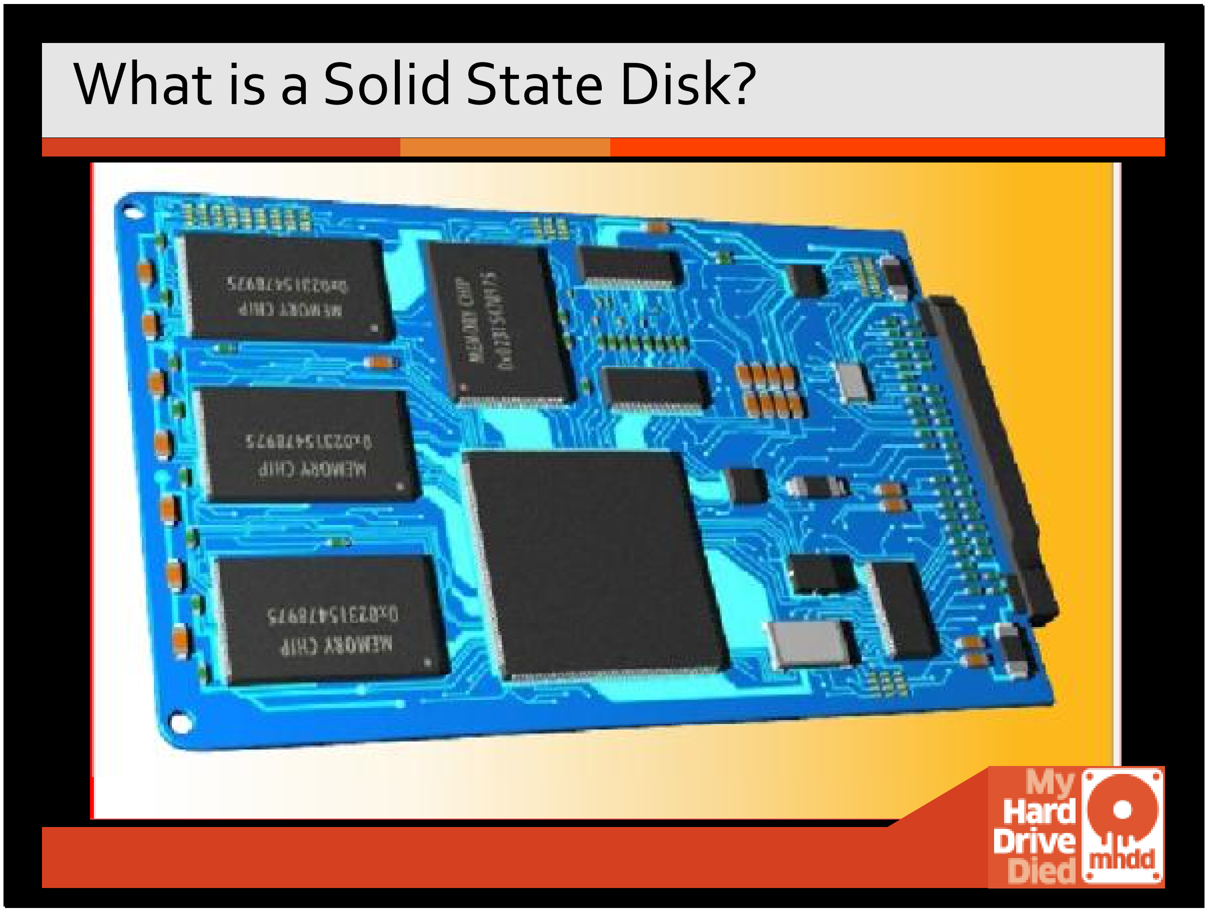

Some vendors have been using the names of both Solid State Drive and Solid State Disk or just SSD. For today, the term SSD has come to mean that it has emulation of an IDE. A memory stick does not emulate IDE because it does not have an IDE connector, however, it has no moving parts so still qualifies as an SSD.

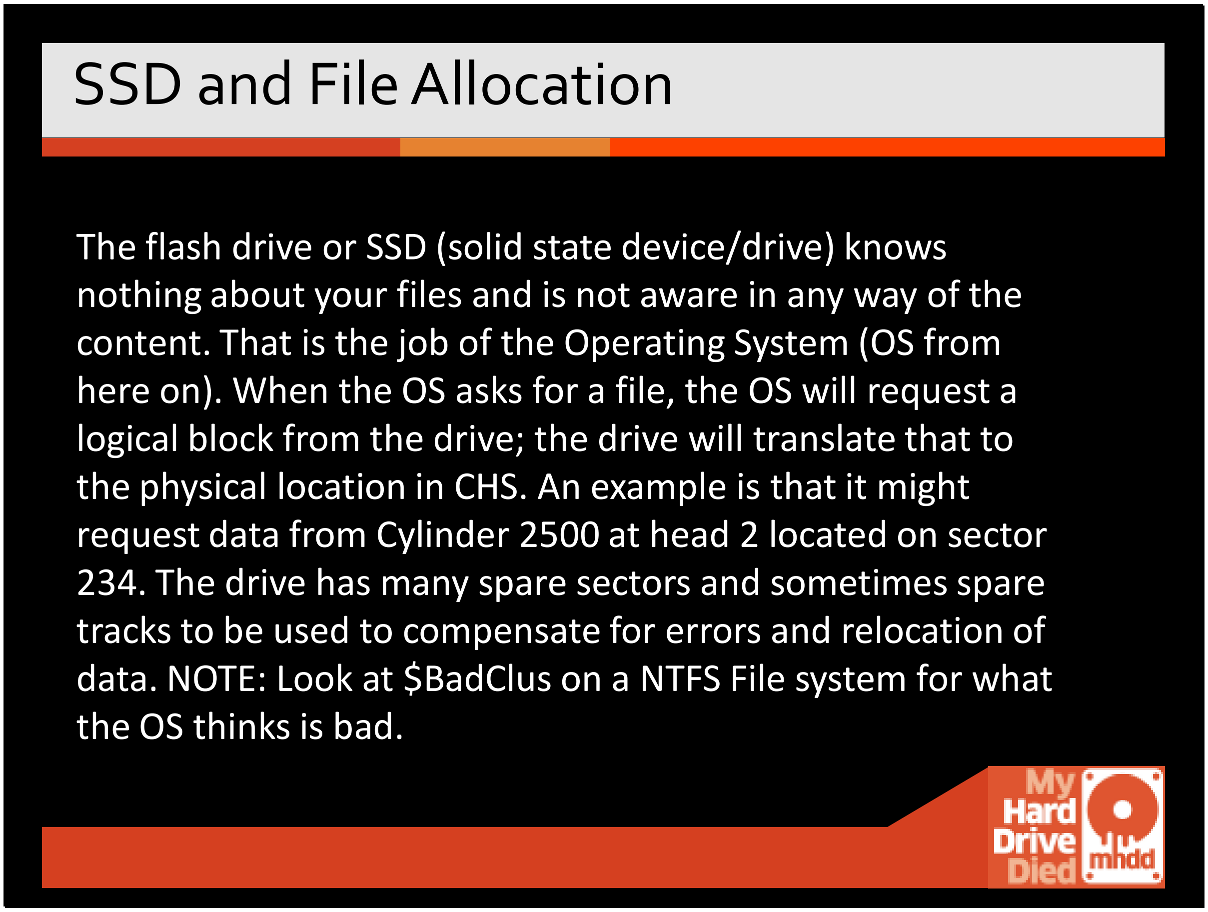

The flash drive or SSD (solid state device/drive) knows nothing about your files and is not aware in any way of the content. That is the job of the Operating System (OS from here on). When the OS asks for a file, the OS will request a logical block from the drive; the drive will translate that to the physical location in CHS.

An example is that it might request data from Cylinder 2500 at head 2 located on sector 234. The drive has many spare sectors and sometimes spare tracks to be used to compensate for errors and relocation of data. NOTE: Look at $BadClus on a NTFS File system for what the OS thinks is bad.

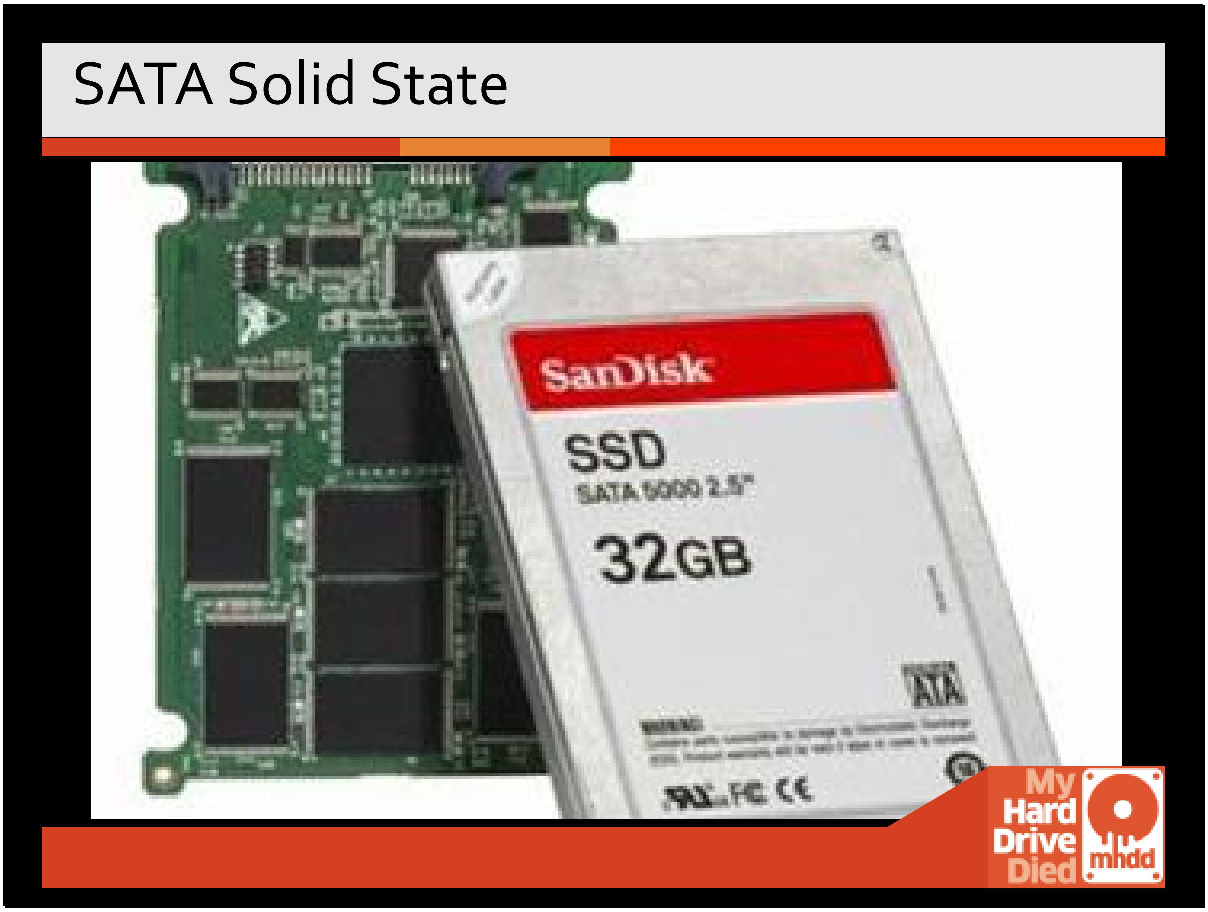

This is what a traditional solid-state disk looks like today. This disk has an SATA connection

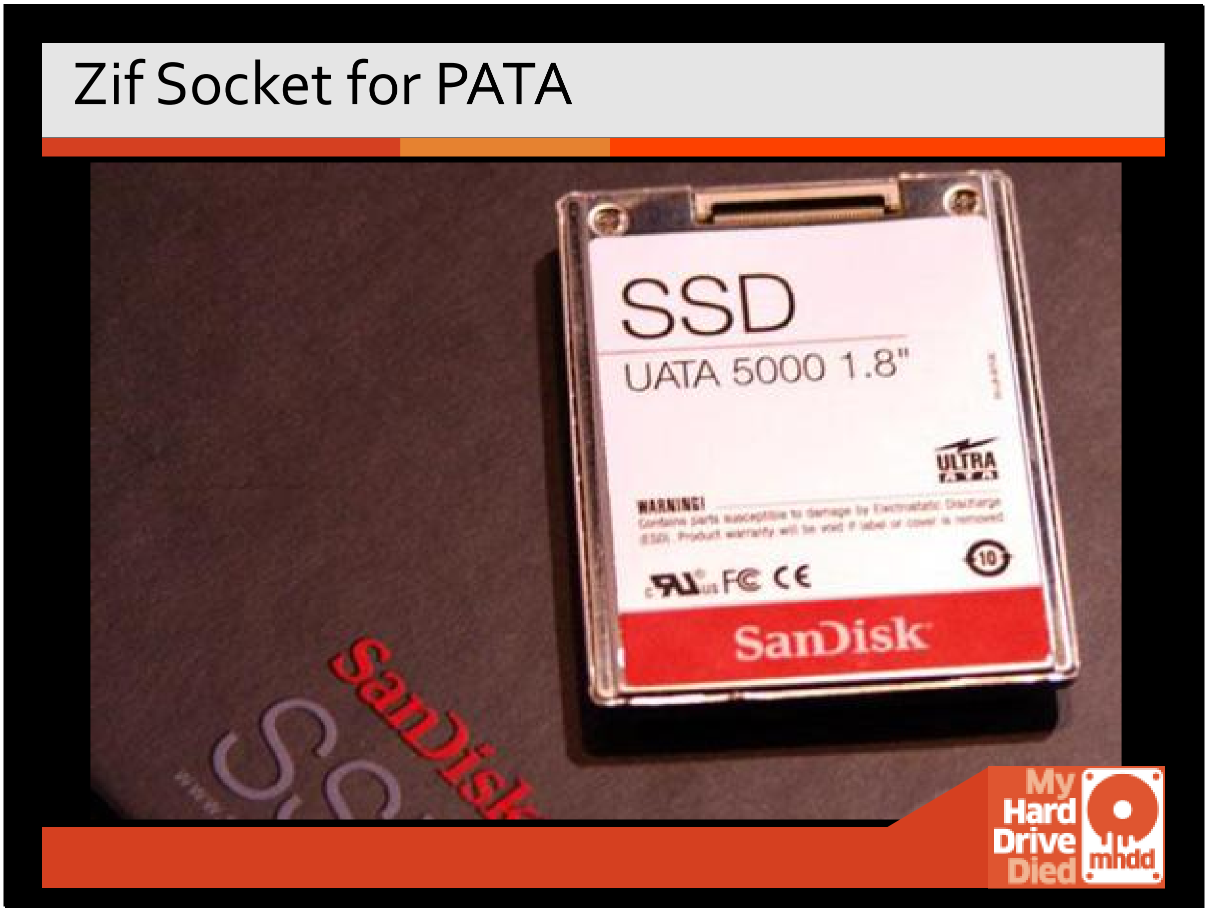

This is a special ZIF socket, which is still a PATA format used by many solid-state disks. The ZIF sockets are very touchy and do not last very long. They are very difficult to connect and generally will only last about 20 connections before failing.

So what if it was a hard drive that had fallen off of an escalator or out of the window?

The damage done to a hard drive by dropping it or any kind of fall, is still very destructive for a traditional hard drive.

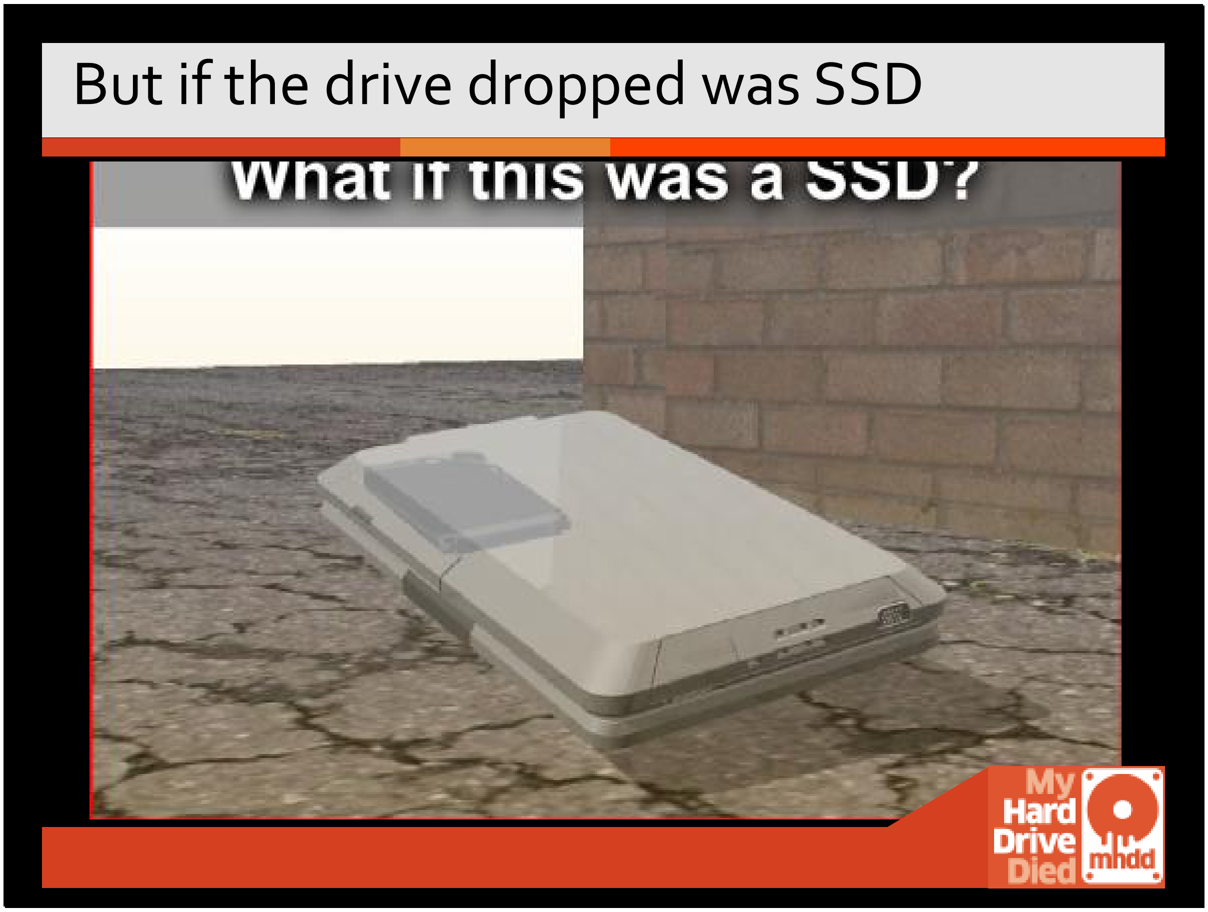

But, what if the drive that fell was an SSD?

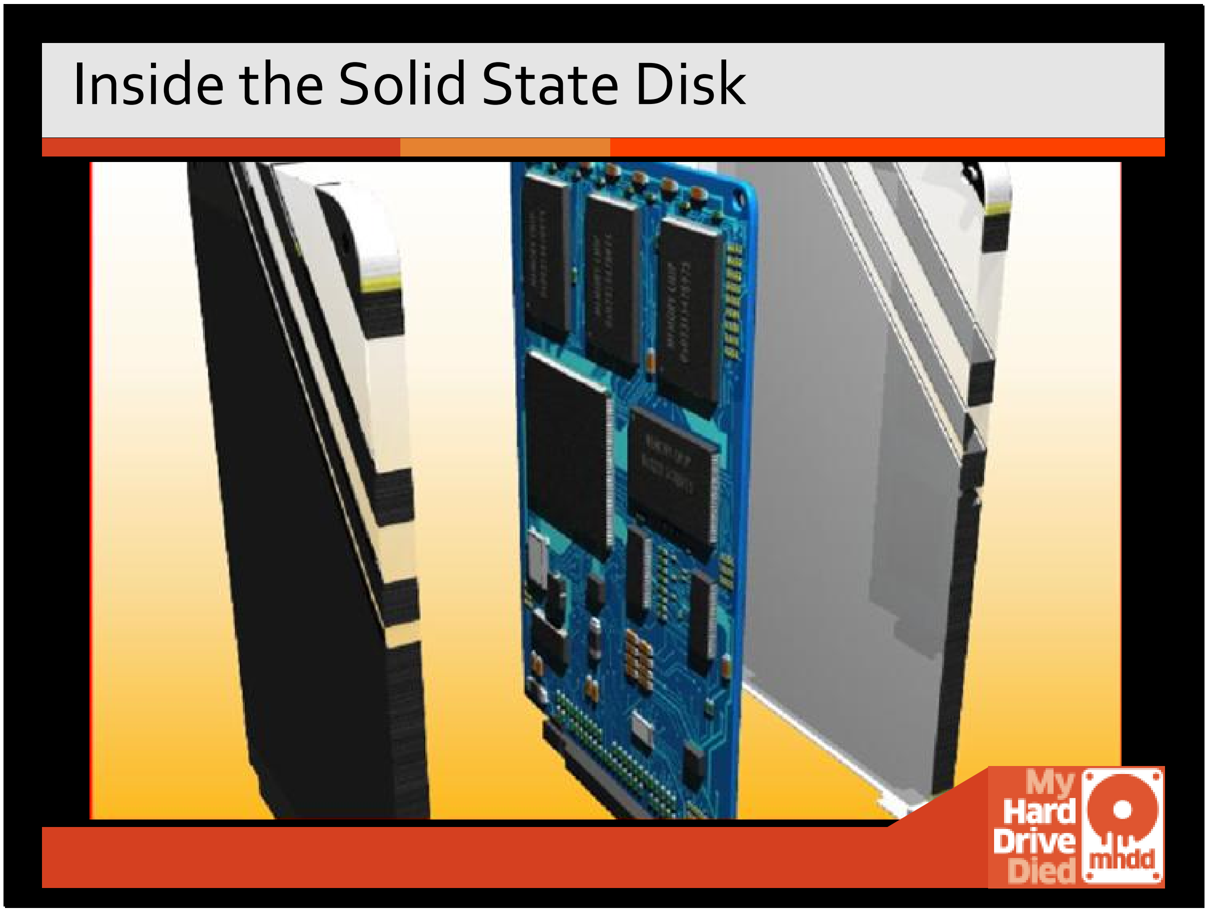

SSD has no moving parts. With no moving parts, solid-state drives are less fragile than hard disks and are also silent (unless a cooling fan is used); as there are no mechanical delays, they usually employ low access time and latency.

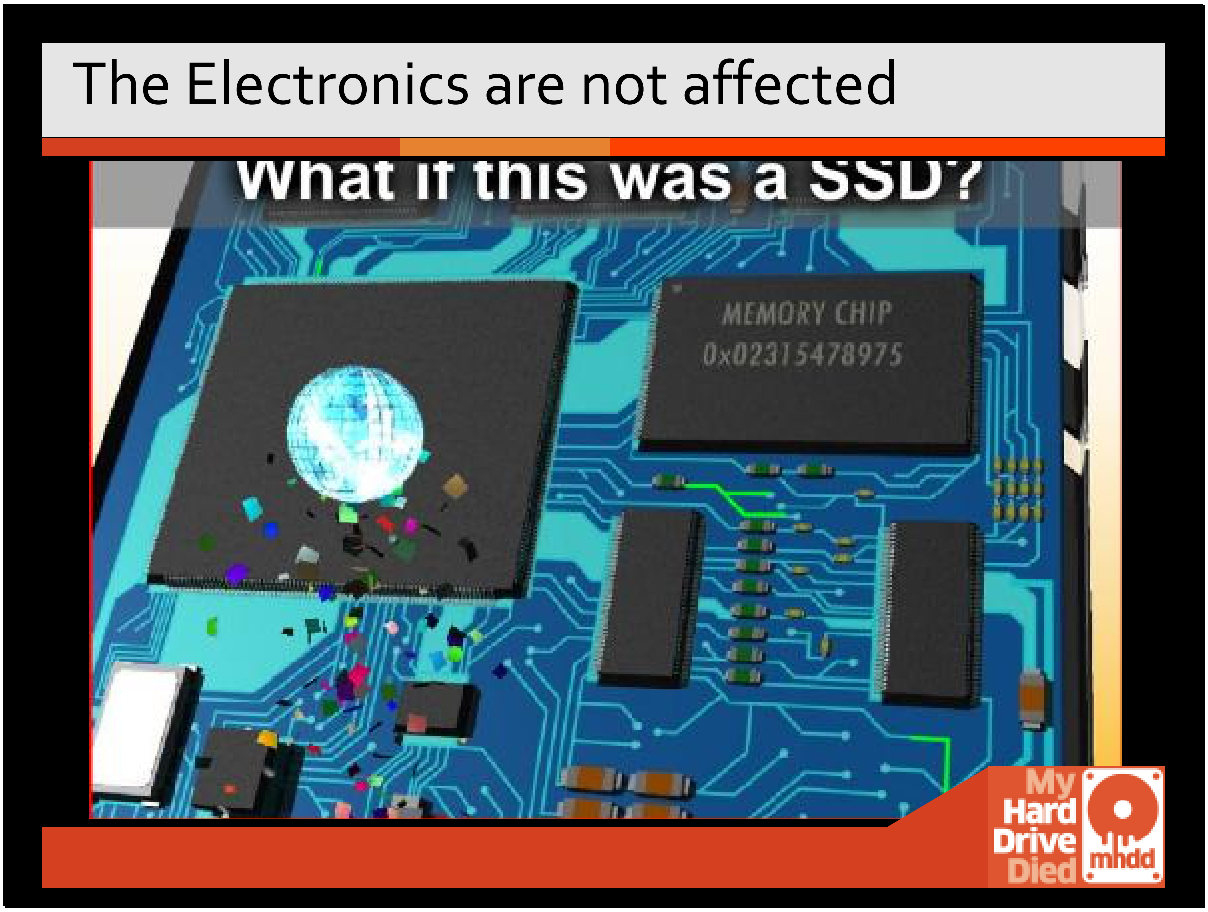

A crack on a board would be major damage for an SSD.

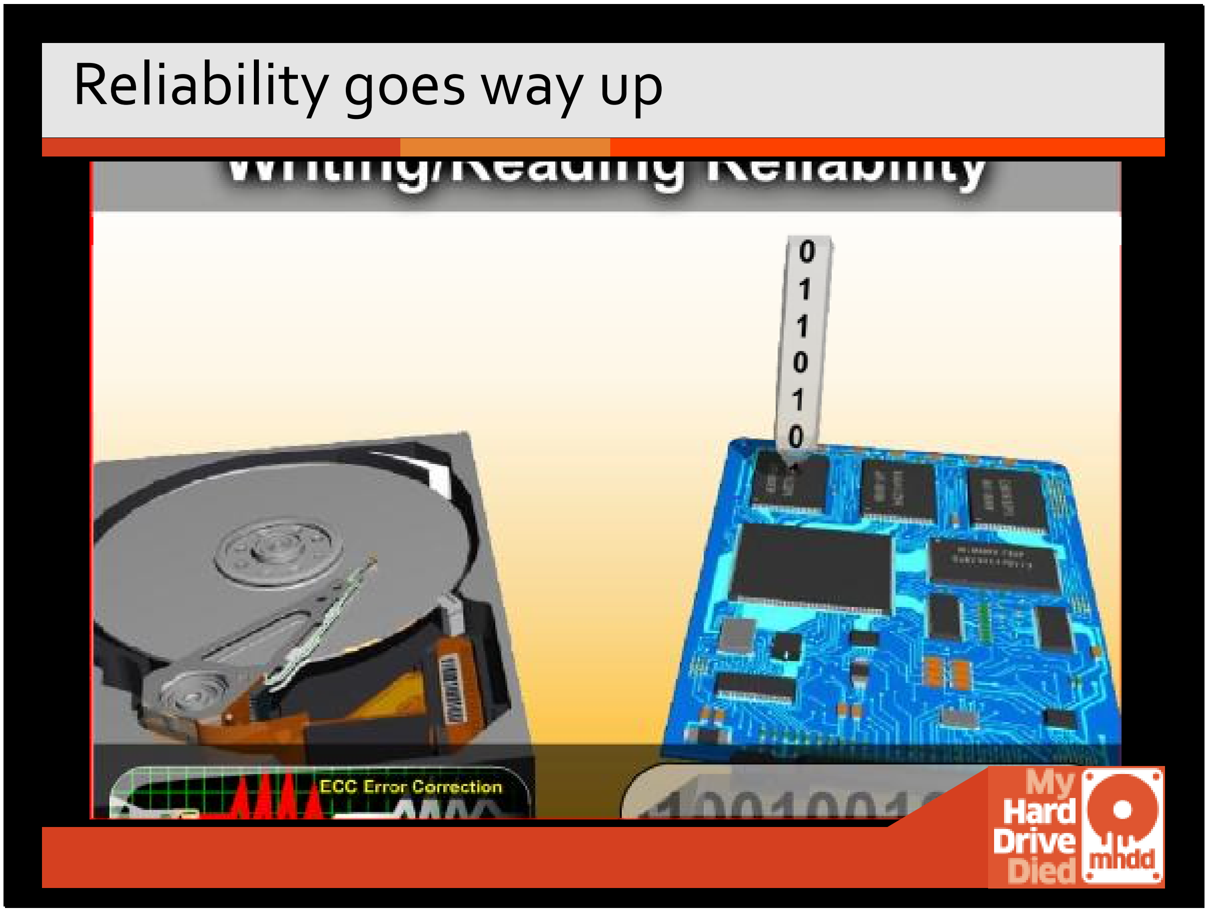

When data is stored as zero's and ones, that means the drive is a lot more reliable.

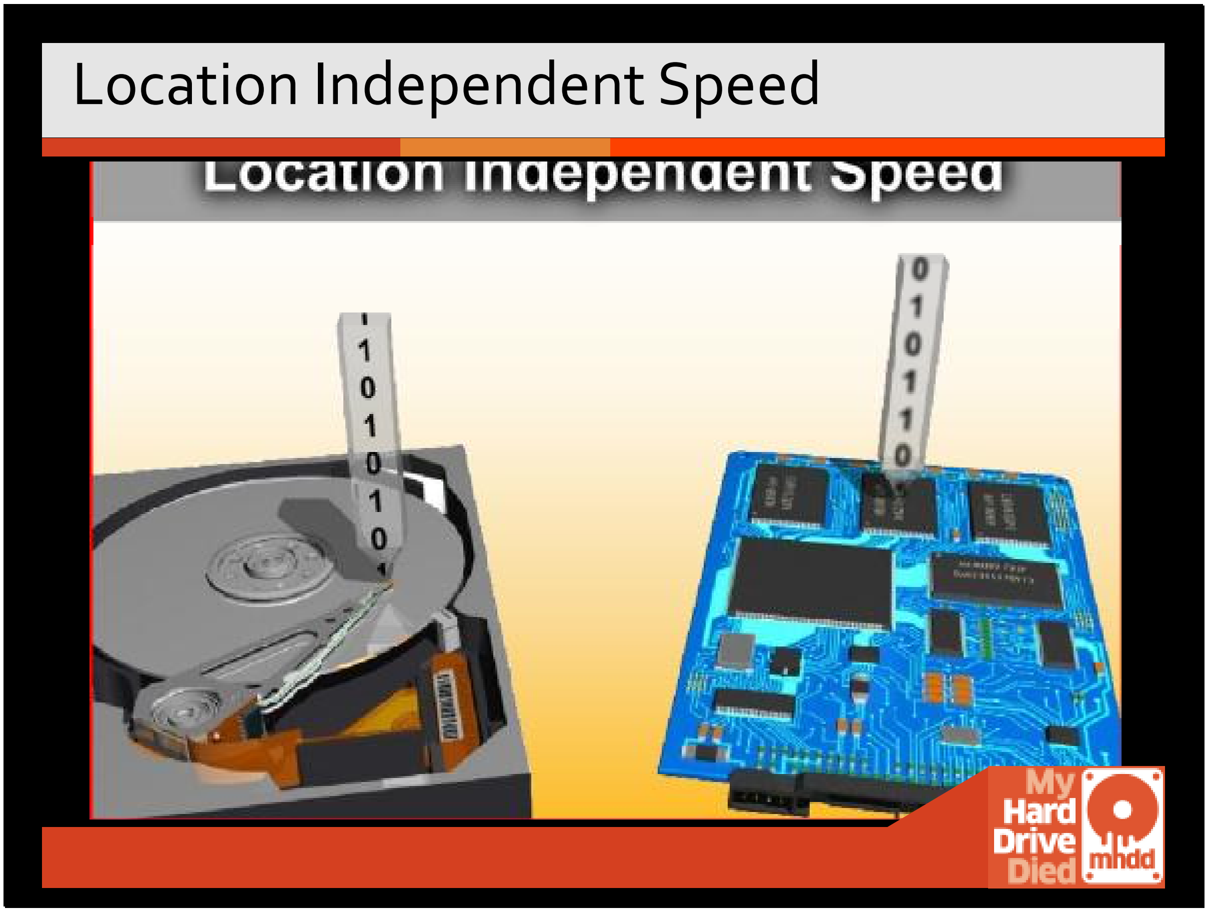

On the hard drive: You have wave forms as disks are being encoded. On solid state disks: Data is stored as 0's and 1's (a lot fewer errors). Hard drives use Constant angular velocity.

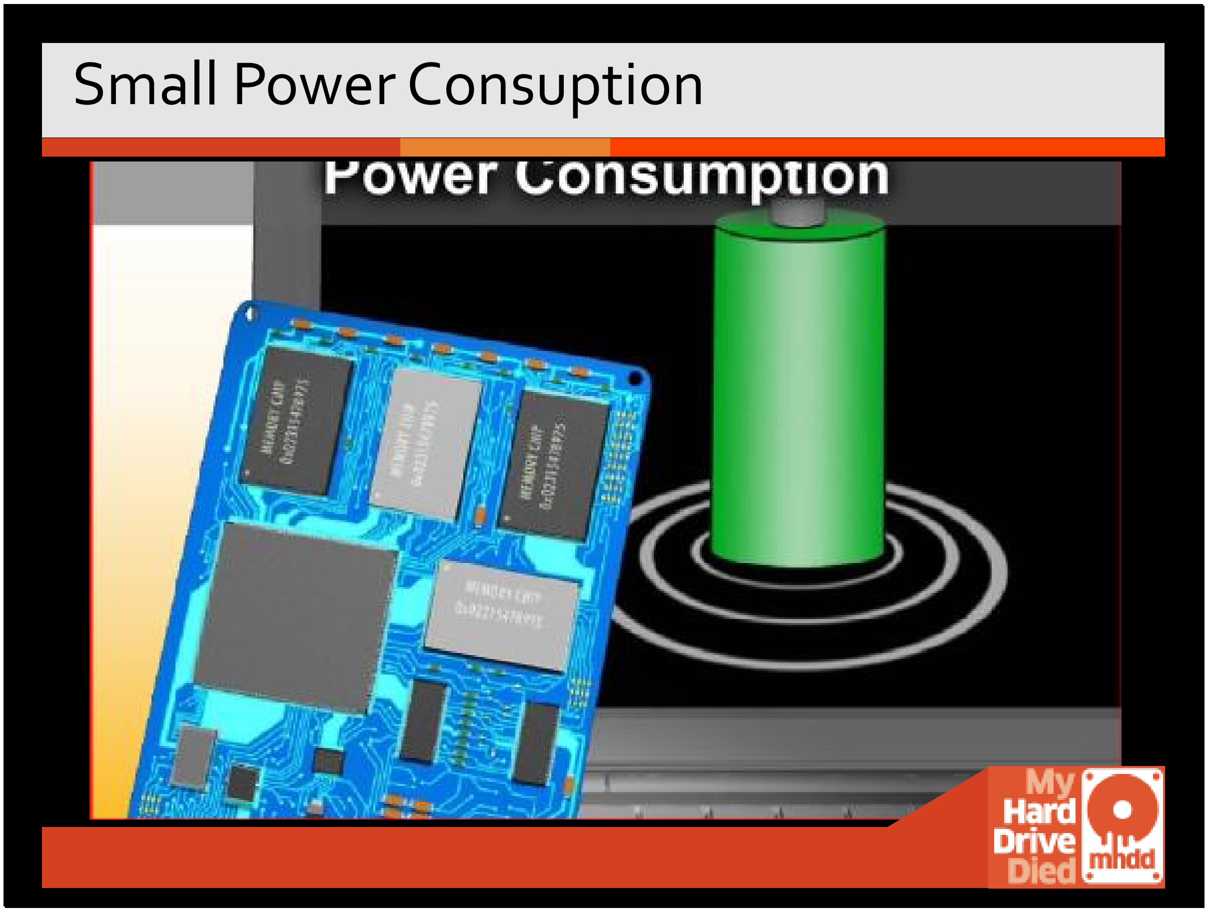

There is a huge voltage differential. An SSD is also faster because there is no spin-up time and it goes straight from one state to another. It's either on or off, inactive or active.

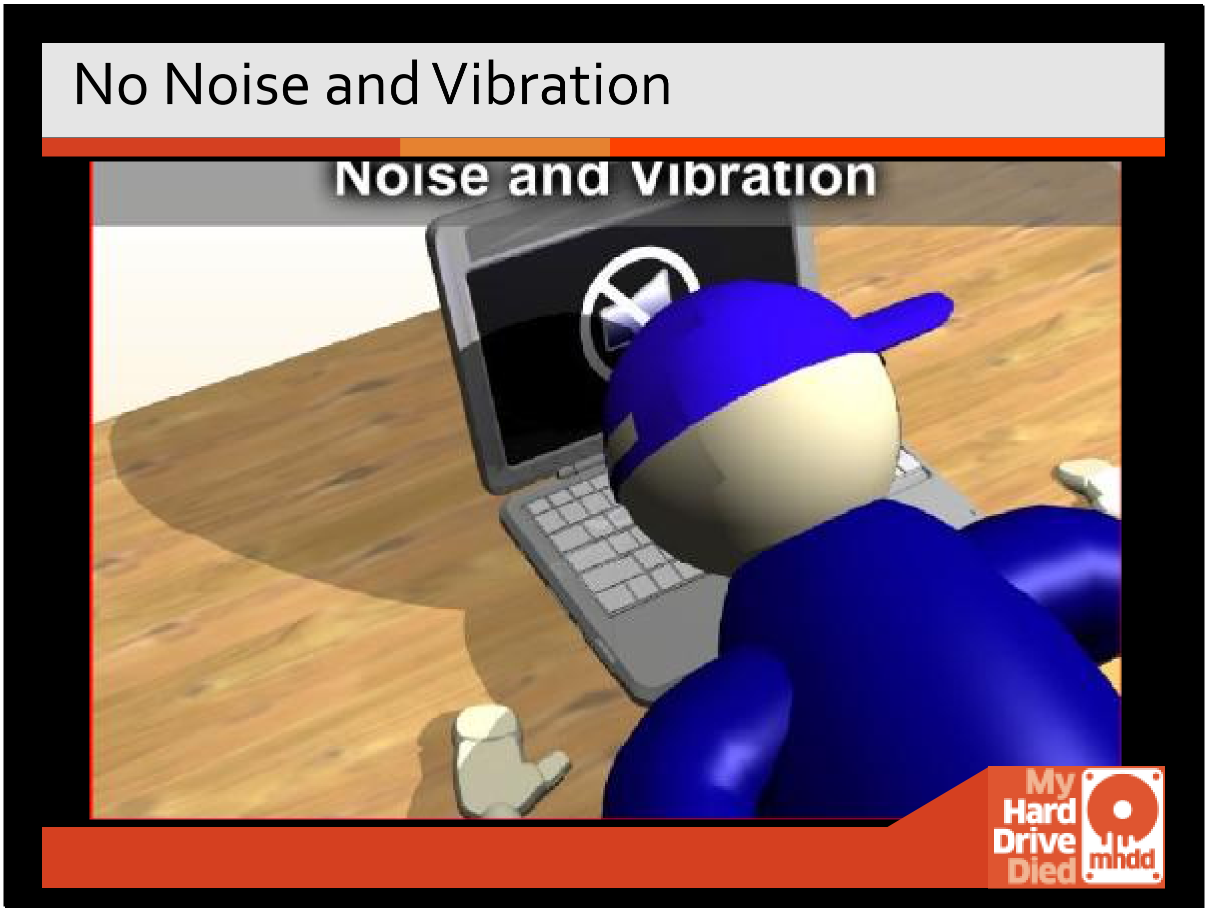

There is no noise or vibration from the motion of the drive. No need to wait for a laptop to go into hibernation mode like you have with spinning drives, you can just put it away.

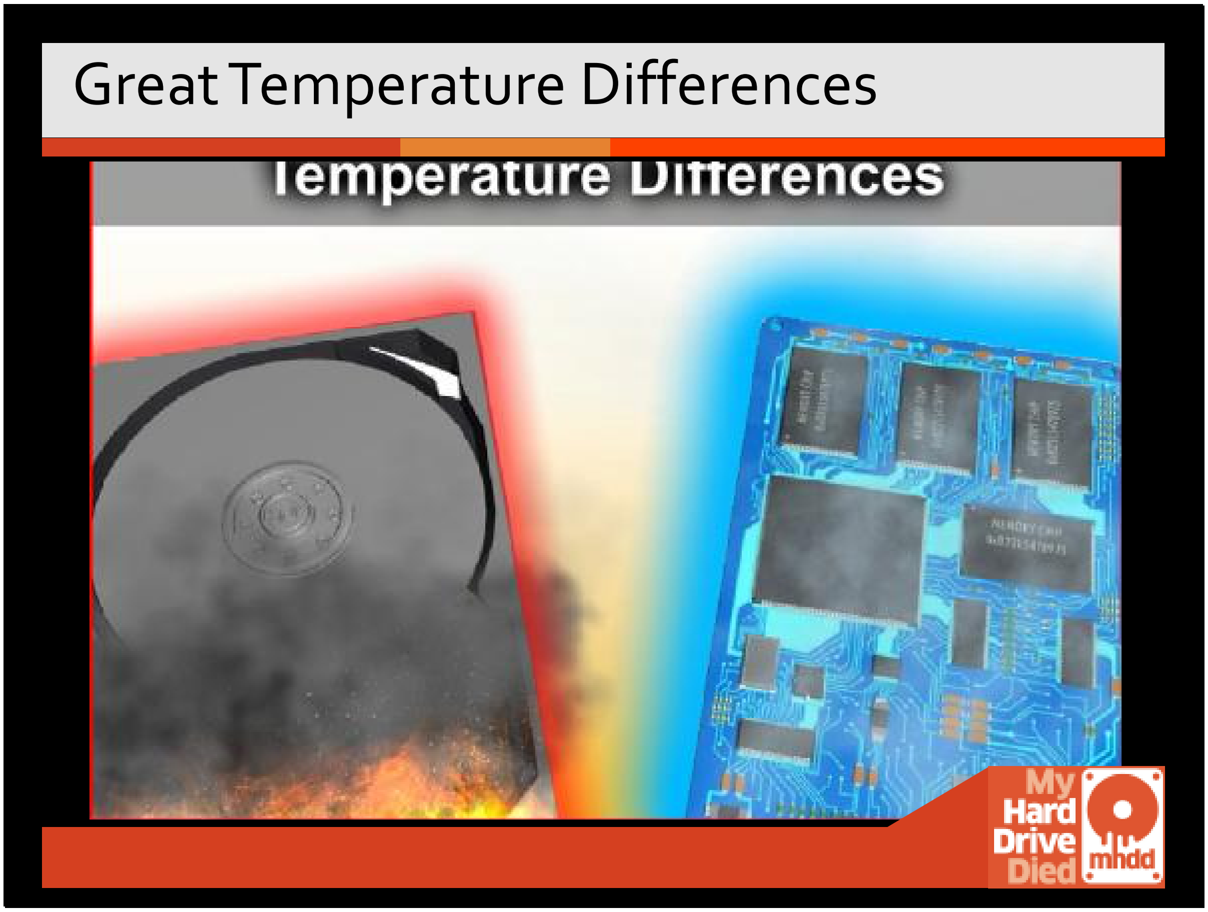

Old hard drives can get up to 179 degrees on older Mac Books. Hot enough to scorch your leg. SSDs are relatively cool compared to hard drives.

SSDs are known to be eco friendly., and they make laptops a lighter weight. Another plus is they don't depend on airflow, whereas, a hard drive does. Therefore, hard drives are more sensitive in higher altitudes in places with high elevation. In areas where the air is thinner the heads are closer to the spinning platters and have a greater probability of touching.

There are at least three basic types of Solid State Disks. NAND – NOR – DRAM.

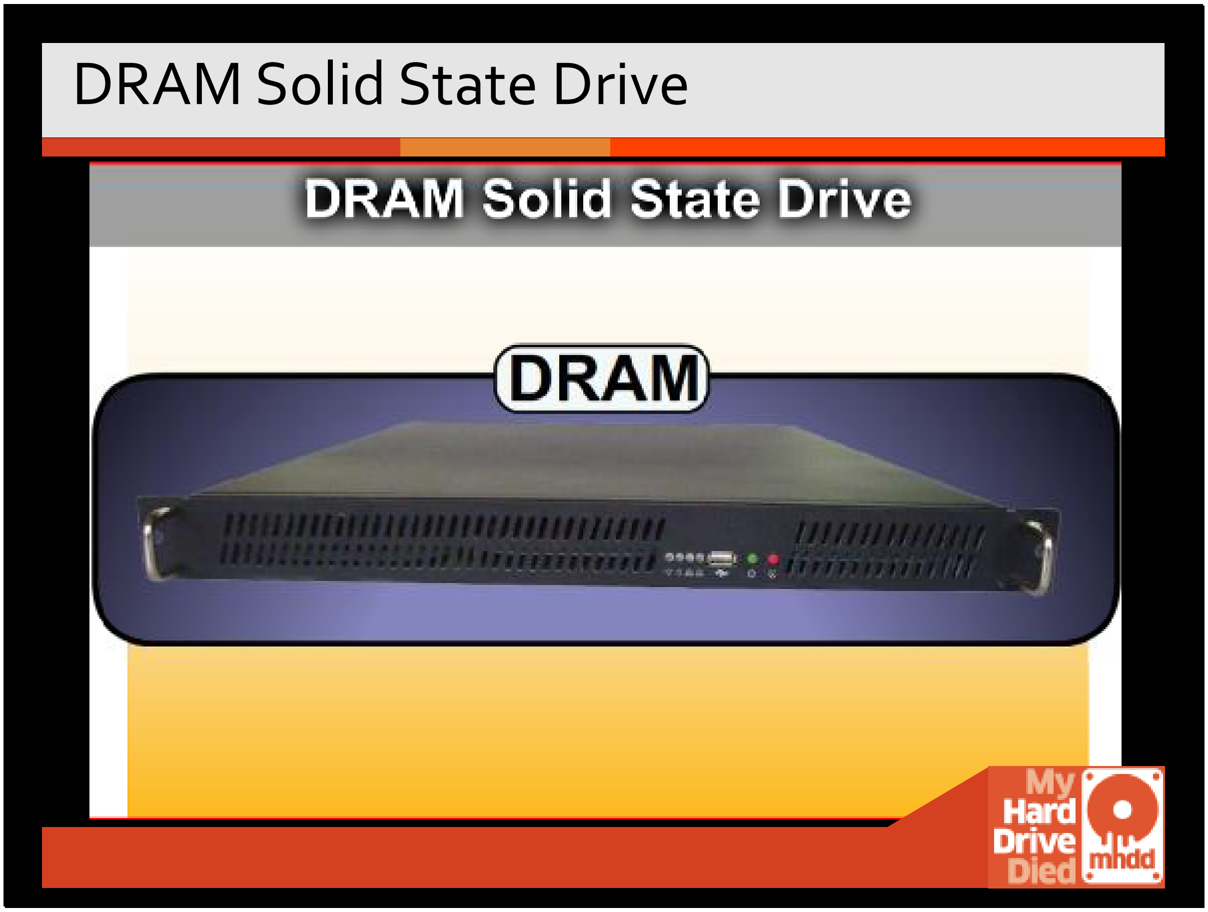

DRAM SSD's have been around for 30 years or so. They are 300-400 times faster than any current "SSD" or Hard Drive. They are common in banks, and military. They are extremely fast, but they do need battery back up and have to hibernate similar to laptops. When the device senses something is wrong it goes into hibernation mode, saves all its data, and returns back to its original state when restored.

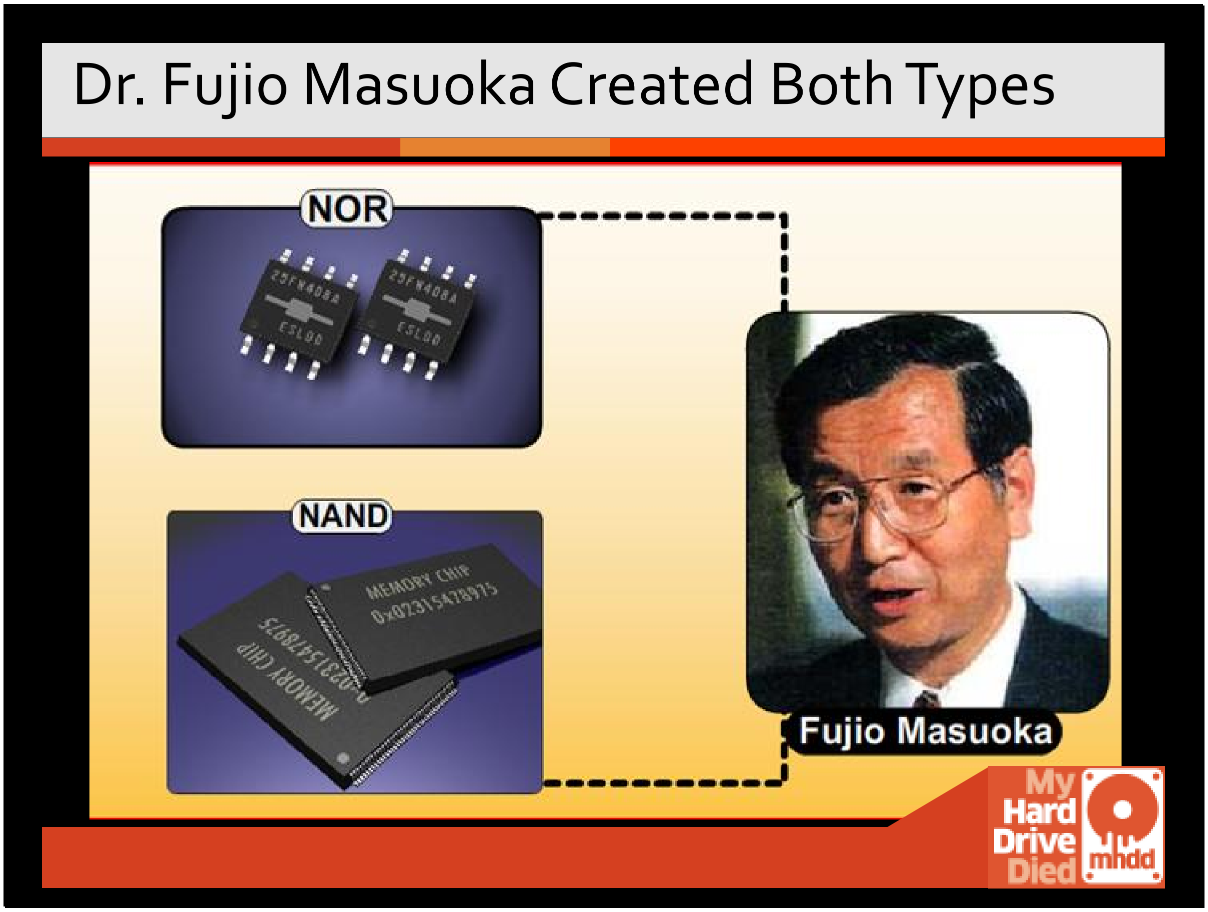

NOR and NAND are the other two types we use most for all SSDs. Dr. Fujio Masuoka was a former employee of Toshiba and is the inventor of flash memory. He also developed the SAMOS memory (stacked gate avalanche injection type MOS) structure. The SAMOS is employed as a memory IC. (http://en.wikipedia.org/wiki/Fujio_Masuoka).

This report is from Forbes. http://www.forbes.com/global/2002/0624/030.html
Unsung hero by Benjamin Fulford, 06.24.02

Fujio Masuoka says that Toshiba tried to demote him after he invented a $76 billion product. He invented flash memory, a technology used in semiconductors with sales of $76 billion in 2001. These chips went into products worth more than $3 trillion, including automobiles, computers and mobile phones.

Flash memory was the most important semiconductor innovation of the 1990s, and it should have made Masuoka very rich. But the 59-year-old inventor lives in Japan. His employer, Toshiba, recognized his efforts by awarding him a bonus worth a "few hundred dollars"--and promptly let its archrival Intel take control of the market for his invention. Subsequently, Masuoka says, Toshiba tried repeatedly to move him from his senior post to a position where he could do no further research.

Toshiba is embarrassed by all this. Its public relations department repeatedly told FORBES GLOBAL that Intel invented flash memory. But Intel says that it was Toshiba, and in 1997 the Institute of Electrical & Electronics Engineers in New York gave Masuoka its Morris N. Liebman Memorial Award in recognition of his invention of flash memory while he worked at Toshiba. When reminded of this, Toshiba admits that it did, in fact, invent flash memory but failed to capitalize on its initial lead.
Masuoka, now a professor at Tohoku University at Sendai in northern Japan, expects to have the last laugh. Since quitting Toshiba in 1994, he has been working on what he expects will be an even more important invention: a "three-dimensional silicon-based semiconductor," he says, which will increase the capacity of semiconductors by a factor of ten. If his invention works as he says it will, Intel could make a 20-gigahertz Pentium chip with the equipment it now uses to make a 2-gigahertz chip. The same would go for other semiconductors, such as DRAMs. It would also delay by 30 years, until 2040, the date when silicon semiconductors reach their theoretical limit. The cost per bit would be a tenth of current costs, he says.

This time Masuoka is applying in the U.S. for patents in his own name. He is seeking venture-capital funding so that he can reap the rewards of his creativity in a manner more in tune with Silicon Valley than Japan.

Masuoka's tale illustrates how Japan lost the semiconductor race with the U.S. in part by neglecting basic research in favor of applied work on established products. He is not the only talented Japanese to become frustrated by the lack of recognition. Shuji Nakamura invented a semiconductor-powered light bulb; in 2001 he sued his employer, Nichia, over ownership of the patents. He now works in the U.S.

Masuoka, a shy but confident-looking man, seemed destined for great things. Four months after he joined Toshiba in 1971, Masuoka, who had just received a doctorate from Tohoku University, invented a type of memory known as SAMOS. After five years at Toshiba, he invented another type and was moved to the semiconductor production division, where he developed a 1-megabit DRAM.

What fired him up, though, was an idea that came to him--yes--in a flash. One of the biggest challenges facing the semiconductor industry in the 1970s was to find a way to retain memory so that it did not vanish every time the power was turned off. Engineers found it too cumbersome to build a nonvolatile memory for each bit of information. Masuoka's insight was that information needed to be stored in big batches rather than in single bits. It was easier to engineer the retention of big batches because this could be done with simpler, more compact circuit designs.

Without permission from Toshiba, Masuoka began spending his nights and weekends working on this idea. By 1980 he had applied for the basic patents on a type of flash memory now known as NOR-type (not/or) flash memory.

It was not until four years later, after a promotion, that he was able to produce the first flash memory. "I was now senior enough that I could go to the factory without permission and order them to make me one," he said. (His promotion resulted from innovations he made in working on incremental improvements in DRAM technology.)

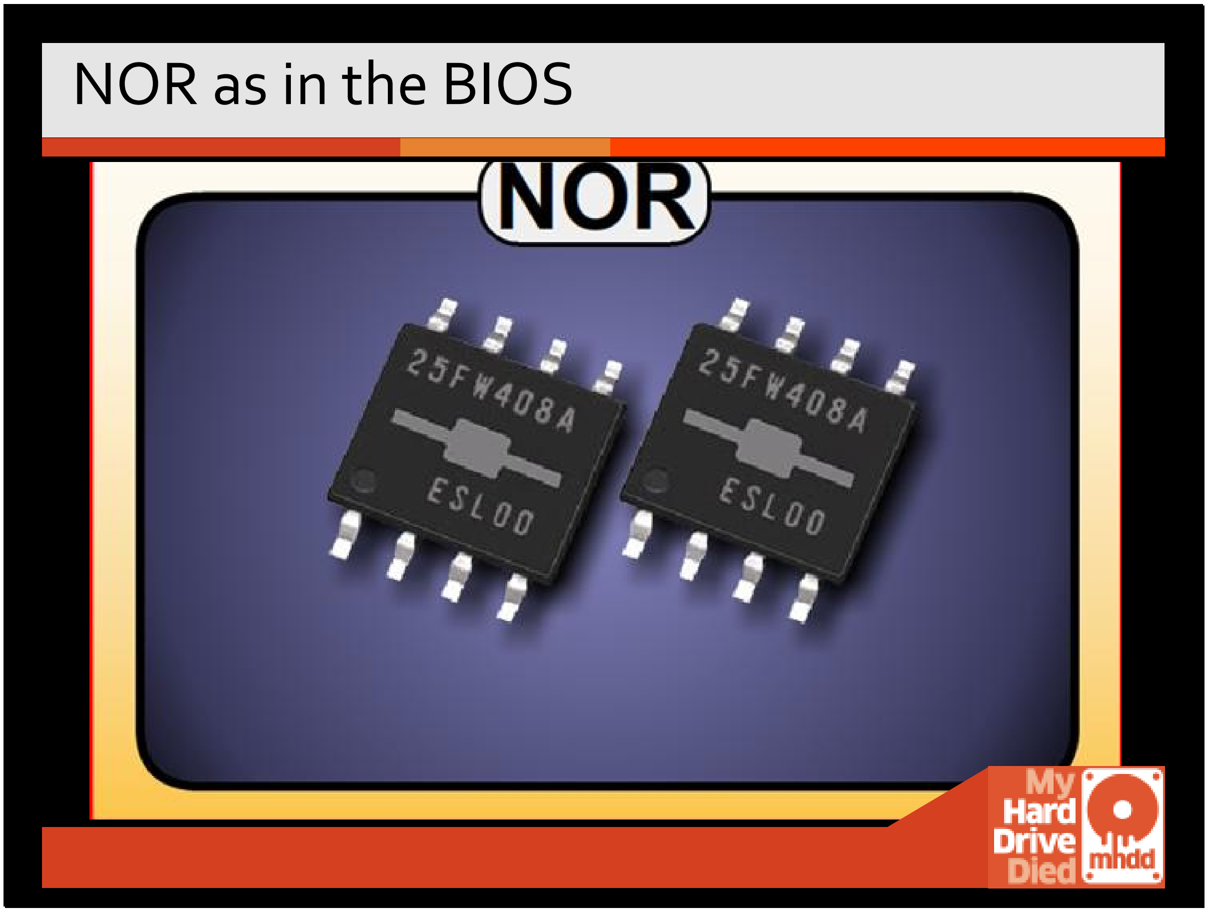

NOR has very high speed read, but slow write. NOR memory is also damaged after 10 writes or so. Intel first used it in 1986. This is great for things like the BIOS, but not great for things that have to be rewritten often.

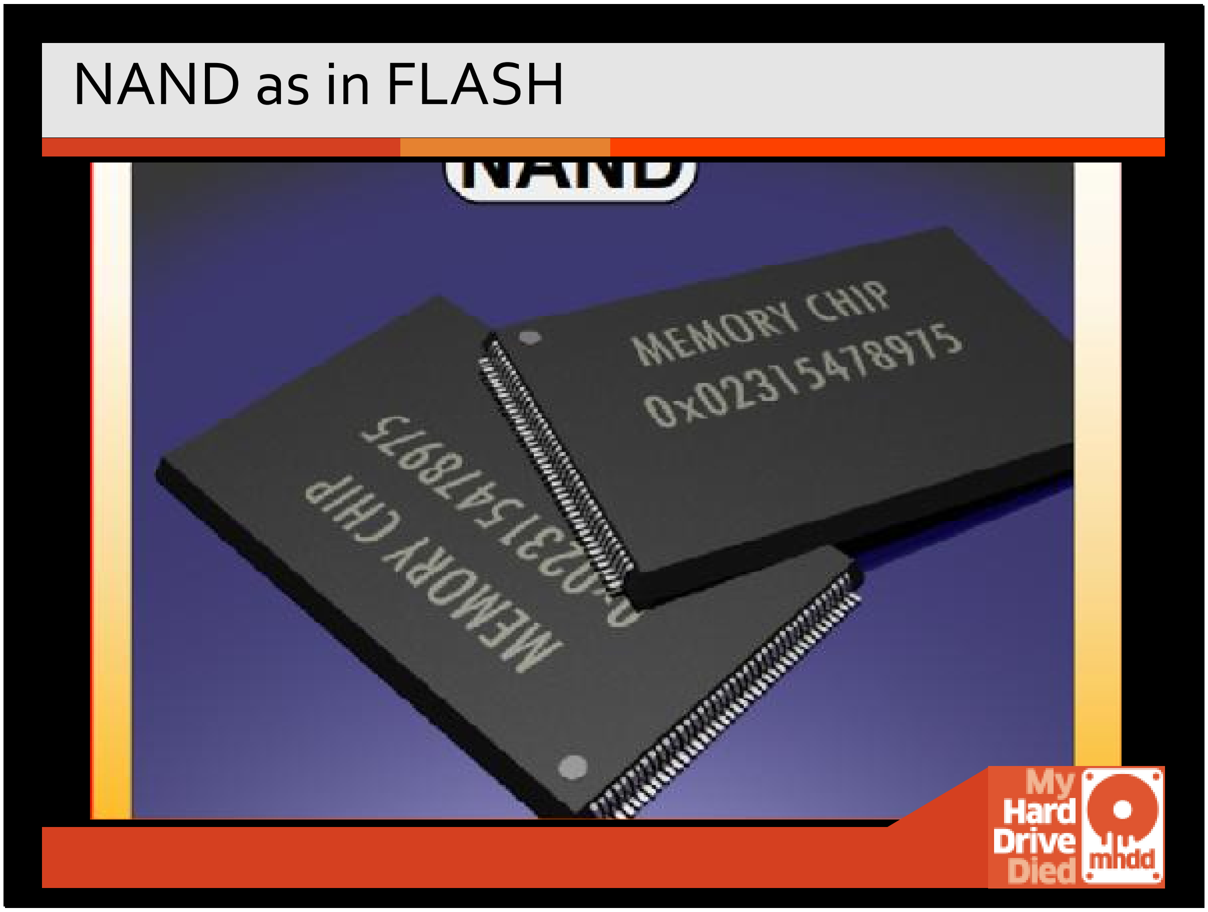

In the last 8 years, NAND has dropped greatly in price and has just become affordable. It can write many times before the cells are damaged beyond usability. The life-cycle is greatly exaggerated by marketing people stating up to 1 million writes. That is incorrect and the number still hovers around 100,000 writes per a gate before damage occurs. However, it can write safely more often than NOR.

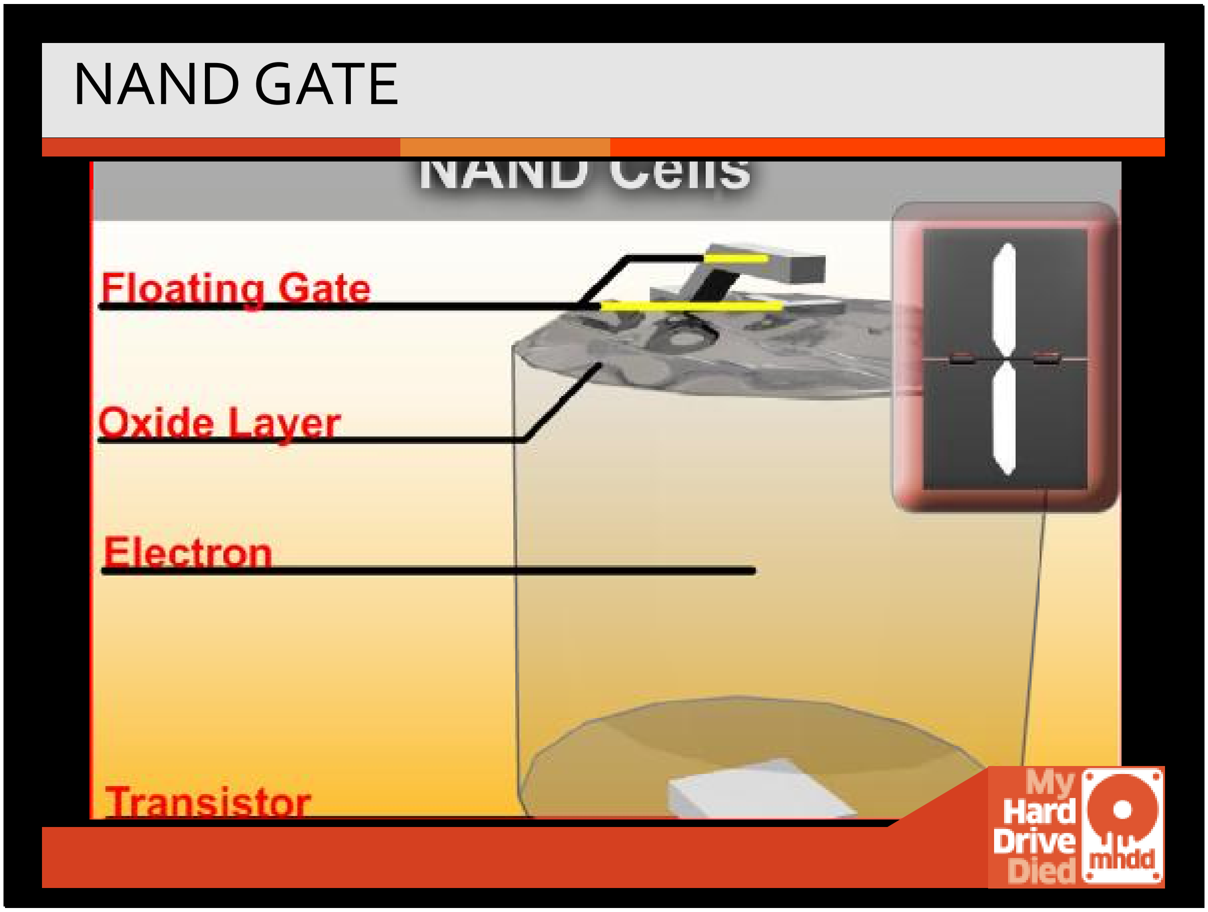

The floating gate itself will push the electron though the oxide layer using a process called "hot electron injection." It will store the electron in the cell. When no content is in the cell and it is empty, its state is One. Usually we would think of that as state zero, but in NAND it is in state One.

When the electron is stored there using hot electron injection, the state will change to zero. It changes in opposite order than what you would expect.

Note: This is not what a true gate looks like. This is my artistic rendition of a cell. If you would like to understand exactly how a gate works you need to refer to source material for engineering purposes.

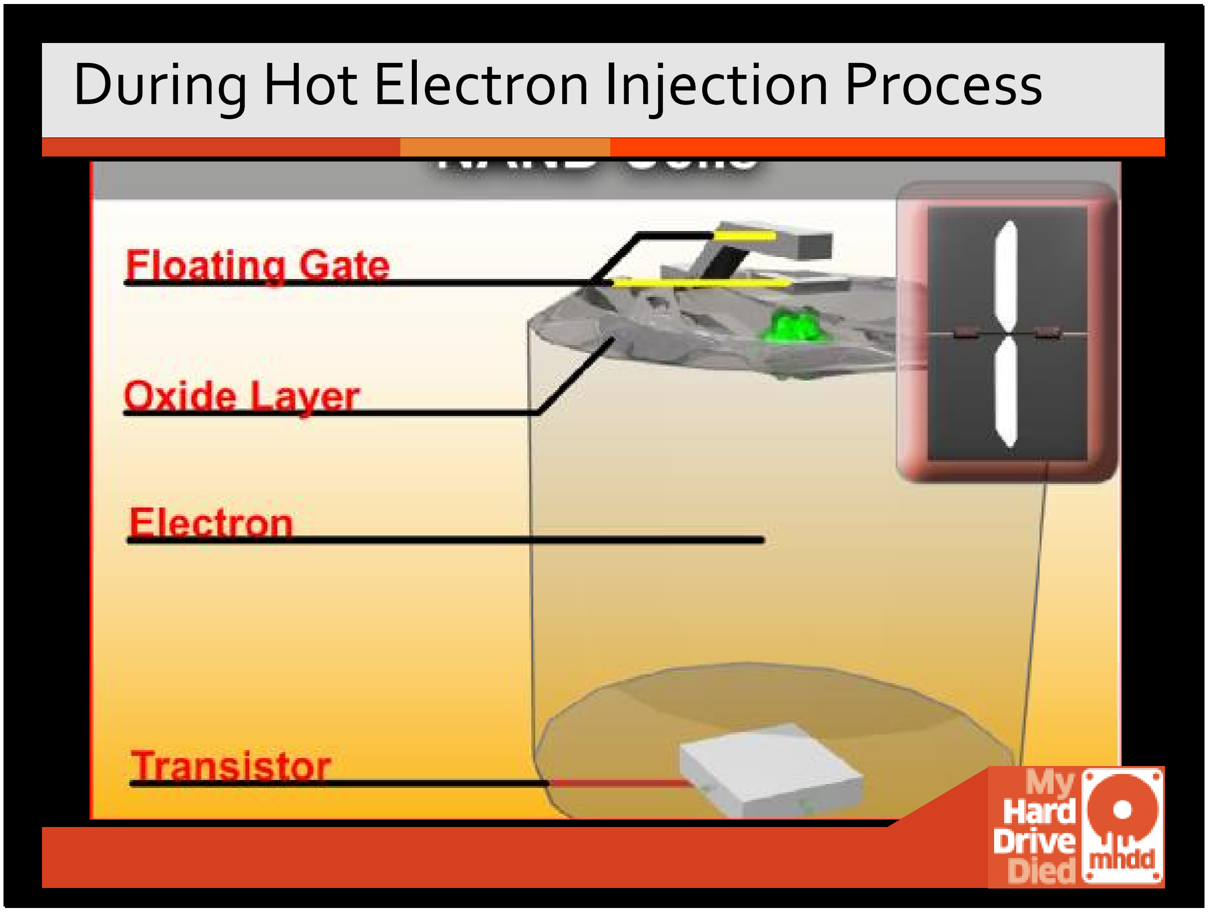

Currently the floating gate is pushing the electron through the oxide layer (hot electron injection), and that will take the state from 1 to 0.

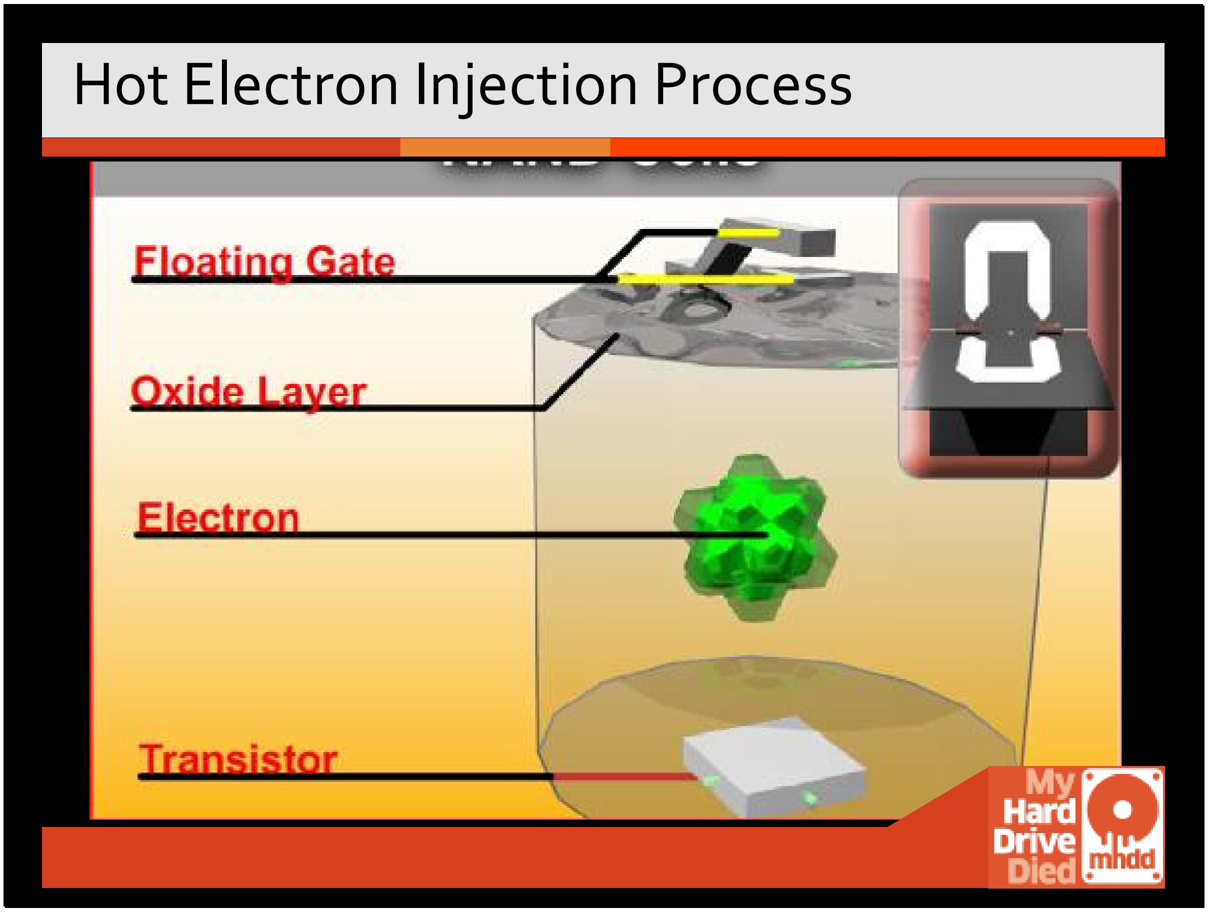

Now the "Hot electron injection" process is complete.

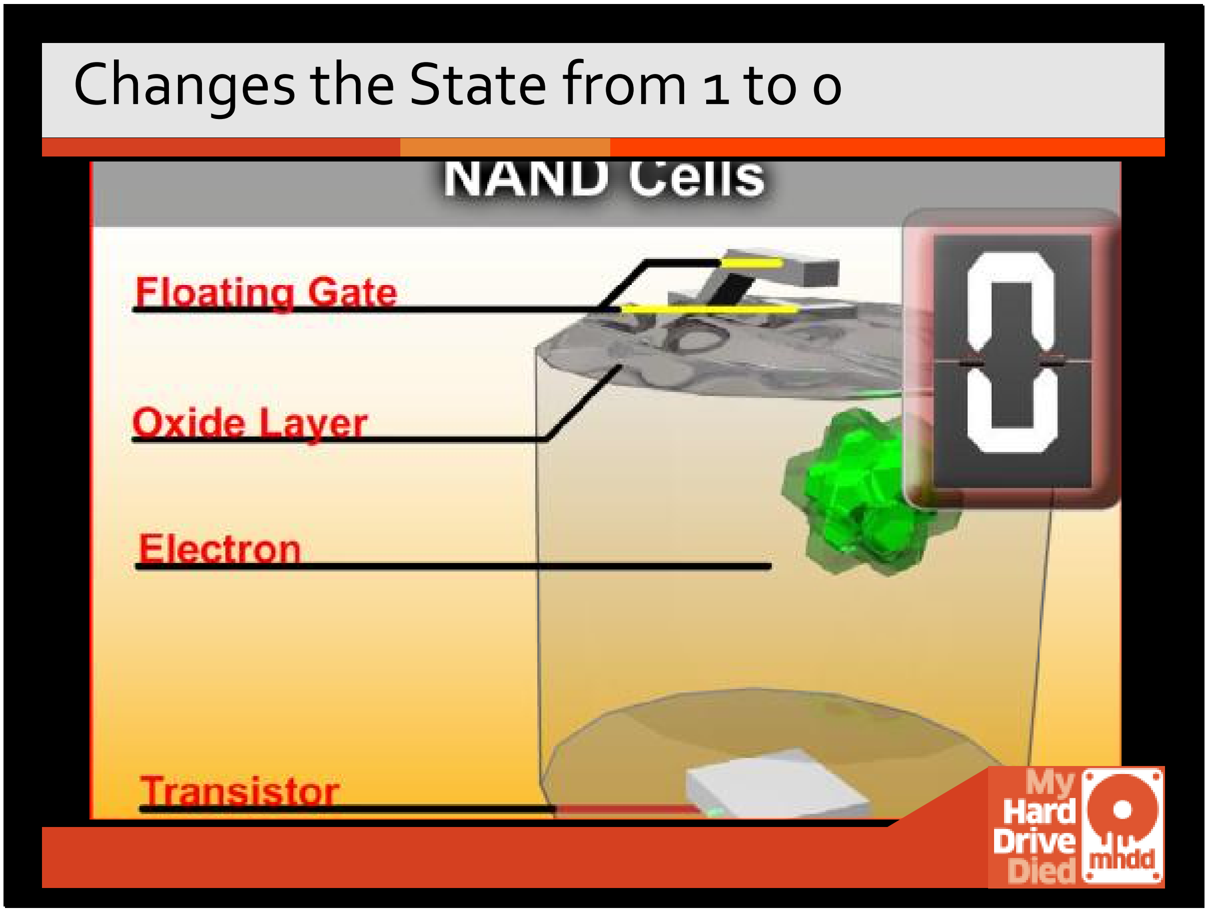

You will notice that the state has not changed to 0. The transistor now will be able to pick up the stored electron and see that there is a charge.

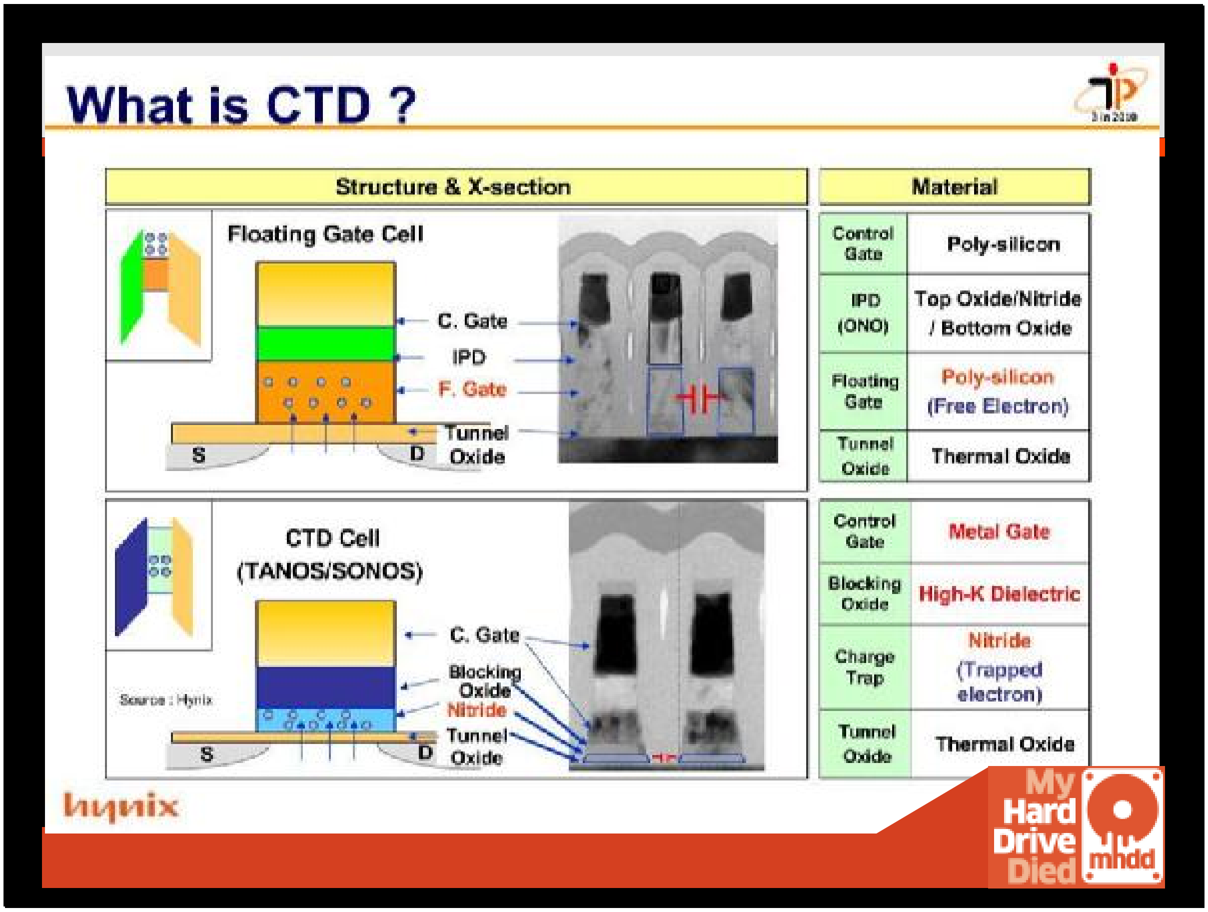

Presentation by Arun Kamat on Webcast 2008

"2012 is only 3 years away"
https://www.denali.com/ui/tv/Webcast?s=webcast2008_hynix2012
Used with Permission for Scott Moulton's Class

Charged Trap Technology, does not use a Floating Gate. It uses a charged trap to capture a charge. Floating Gates will reach their limits and will not be usable en mass for large devices in the future. The charge will be detected by a few hundred electrons and will change and have problems if you lose just a few dozen electrons.

The charged trap has the best future technology to replace the Floating Gate. The costs and power consumptions are comparable.

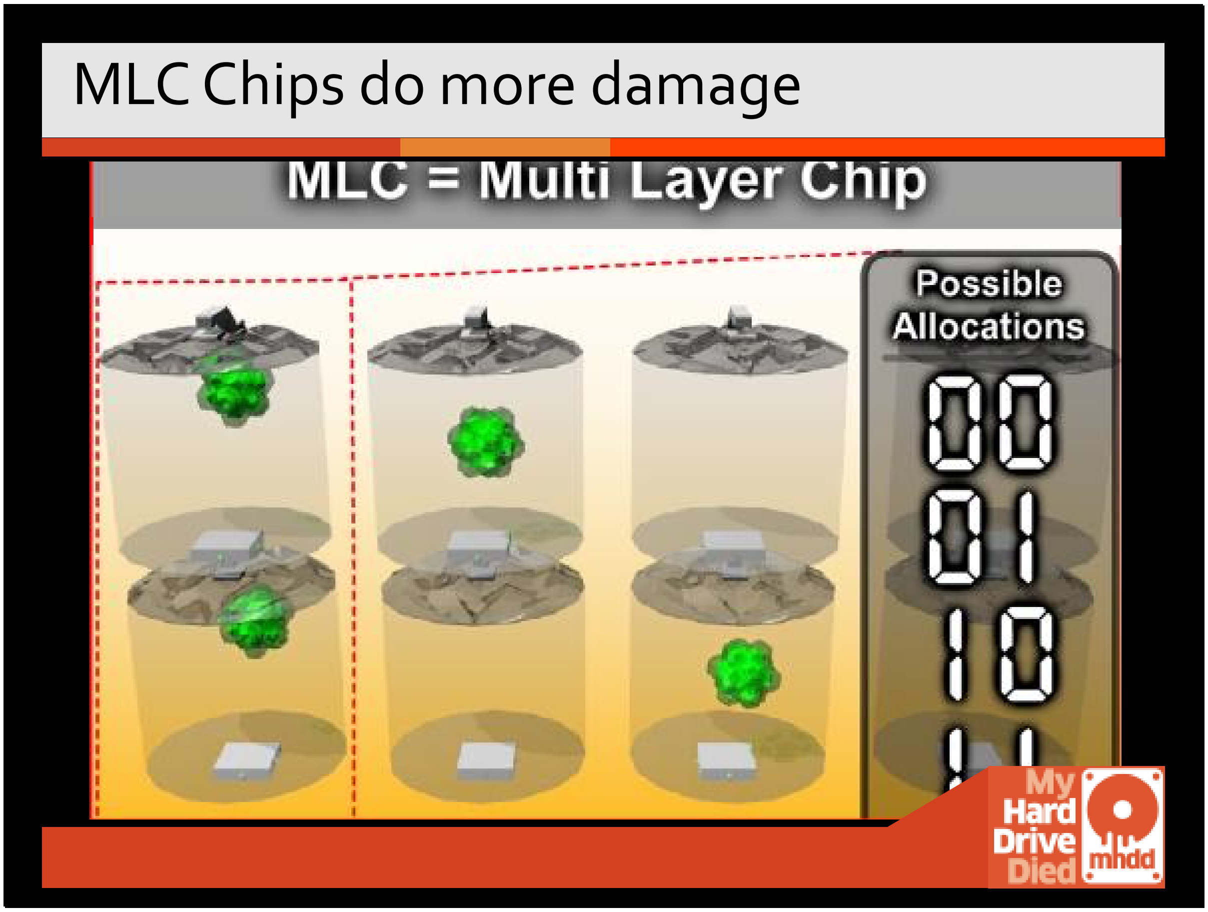

NOTE: They are not separate CELLS just cells CHARGED AT MORE LEVELS – That is why they are damaged more in MLC and TBC (triple bit cells) because they are charged more. They are also as much as 40% slower.

Layer of 1 set of cells = Single layer chips = SLC. Multi-layer chips are MLC. Some MLC have up to 8 different layers currently. MLCs have to write all of these cells at the same time. You need to know or expect to know what will be stored in there in advance, as you do not get a second chance to write to the block. MLC chips have been slower than SLC chips by 20 to 40%. Lower priced drives usually use MLC flash memory, which is slower than SLC flash memory. Flash-memory cells will often wear out after 100,000 to 1,000,000 write cycles.

Events leading up to year 2012. According to Hynix, this is where everyone will be very aggressive. Predicted to be 35% of the NAND market in 2012.

SLCs are a lot more reliable. SLC is used more in the Industrial Market and MLC in the Consumer Market. MLCs require more power to operate than SLCs.

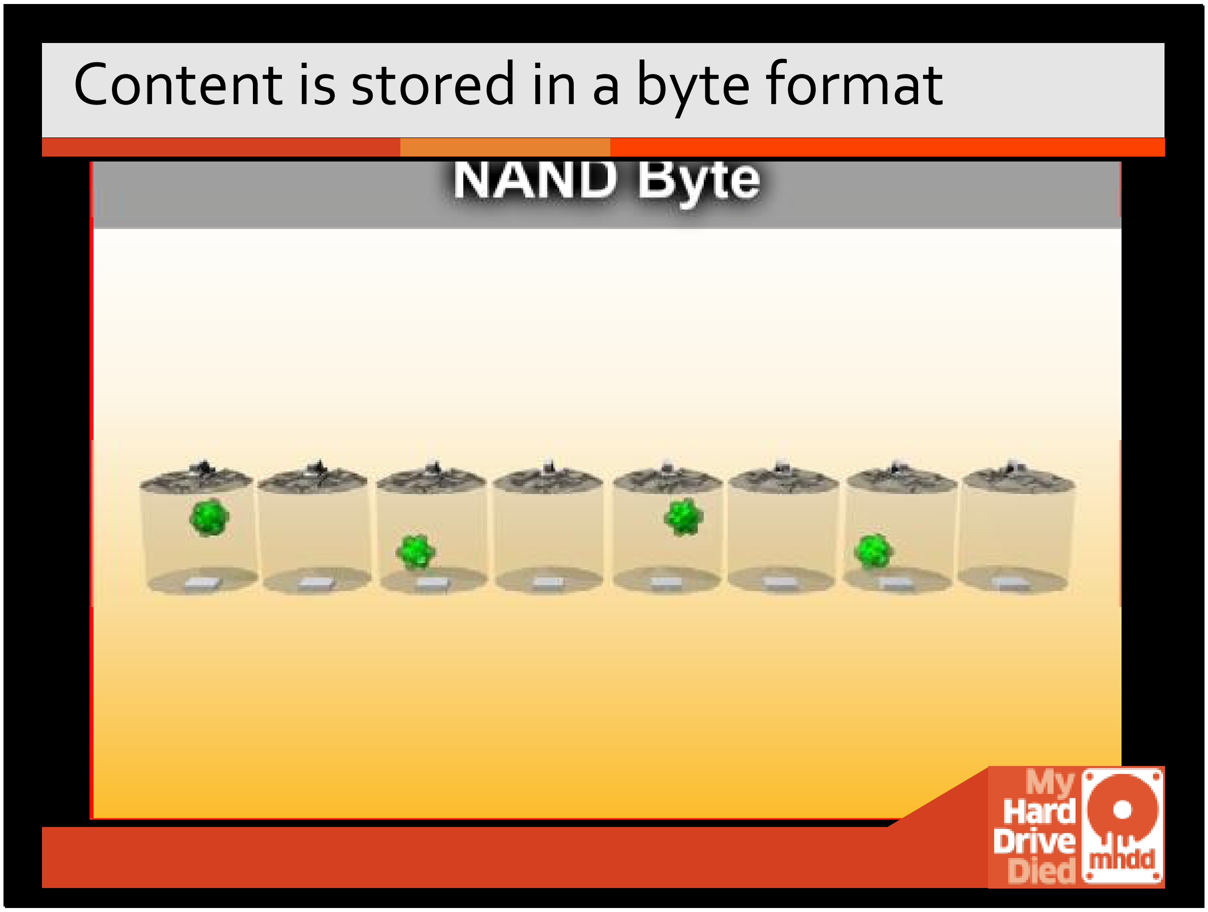

And while we still have a Byte, that is not very usable in NAND.

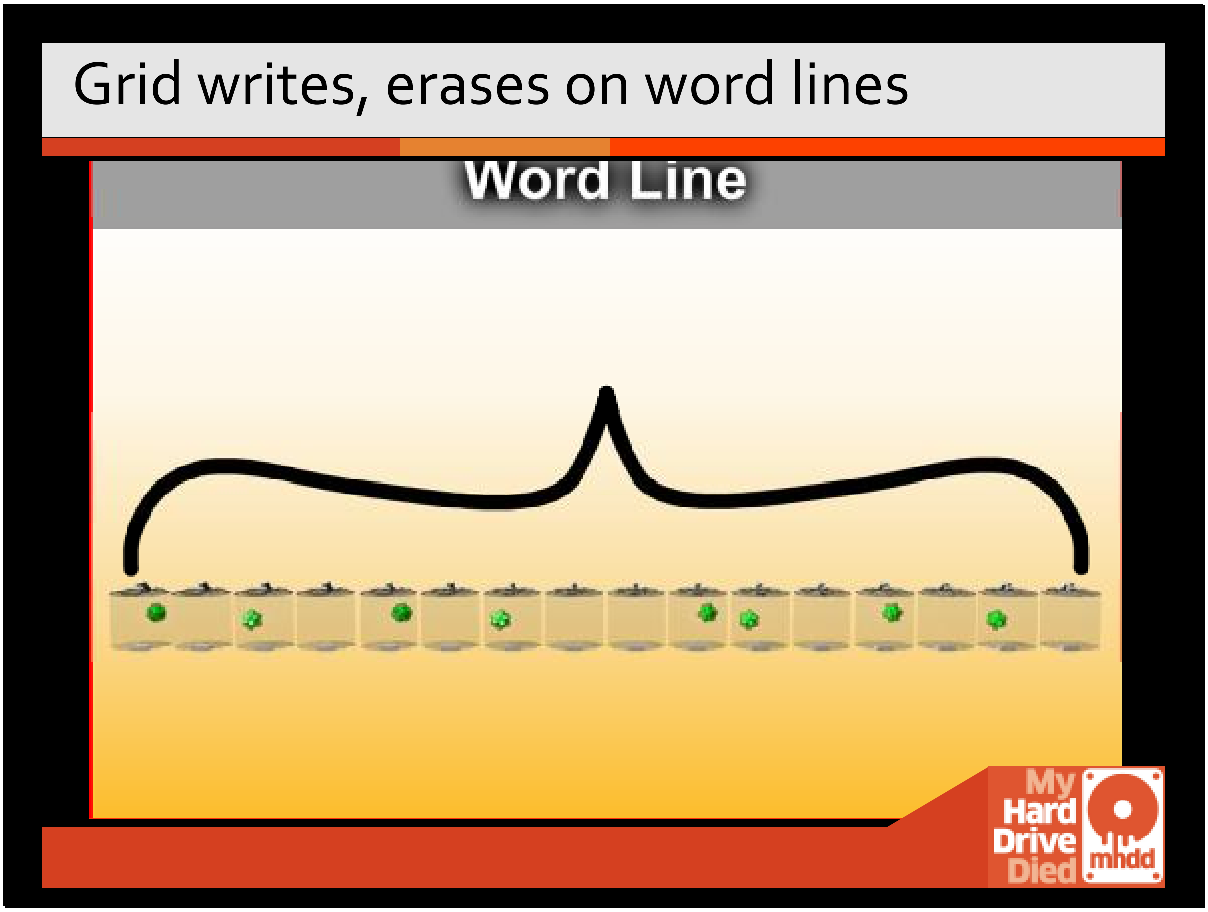

Due to the arrangement in a grid, you can't write to just one line. It must be a group of gates, according to the manufacture.

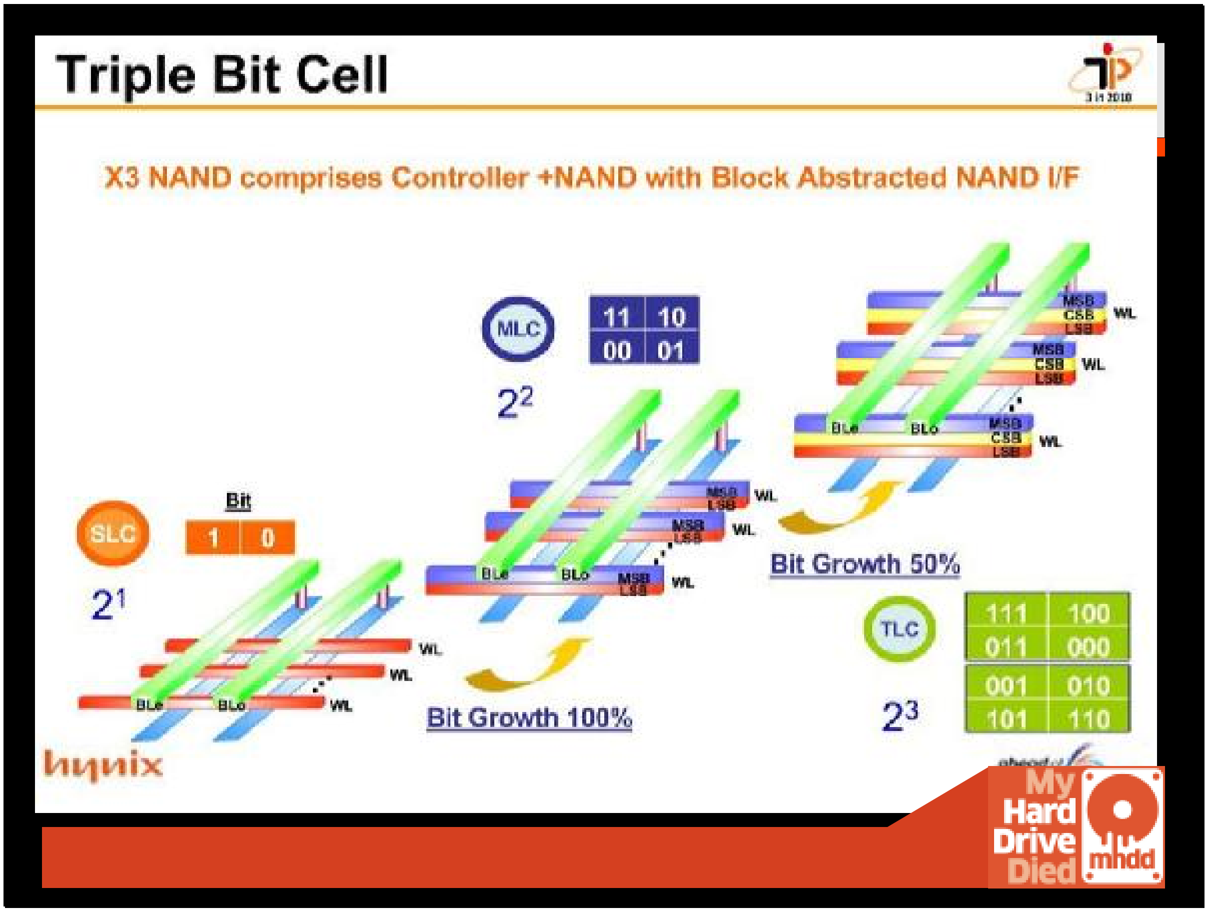

Presentation by Arun Kamat on Webcast 2008
“2012 is only 3 years away”
https://www.denali.com/ui/tv/Webcast?s=webcast2008_hynix2012
Used with Permission for Scott Moulton’s Class

Multibit cell technology will change the future for the next few years in NAND flash. We will be moving past the SLC and MLC and moving into the Triple Bit Cell. A three bit cell will be able to increase the number of bits by 50%, and you will have to have an embedded controller to make this work with this complexity.

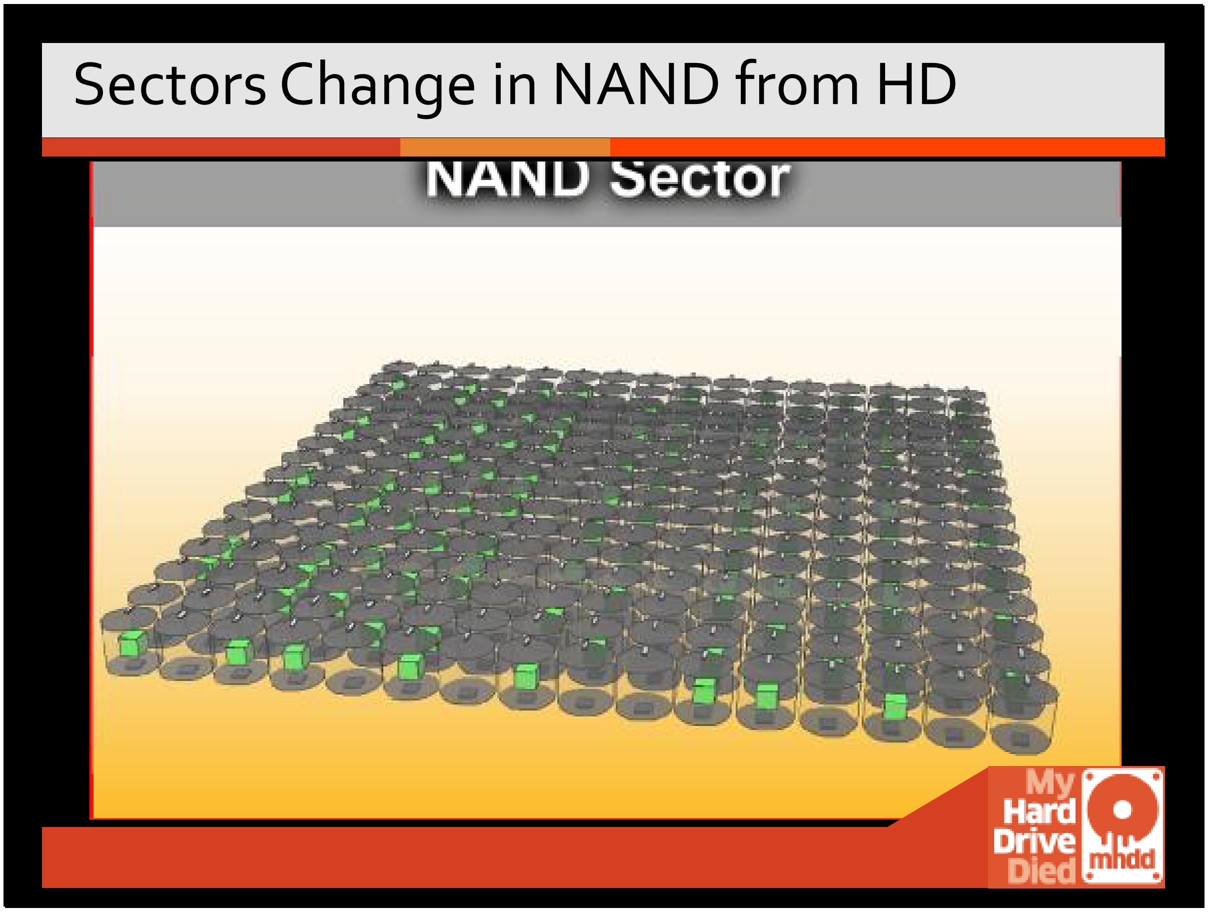

512 bytes was good until after 256 megs. Afterwards, no use for 512k sector anymore (though, vendors can vary that). Standard sector in a memory stick. Most sectors are now 2k, 4, or 8k

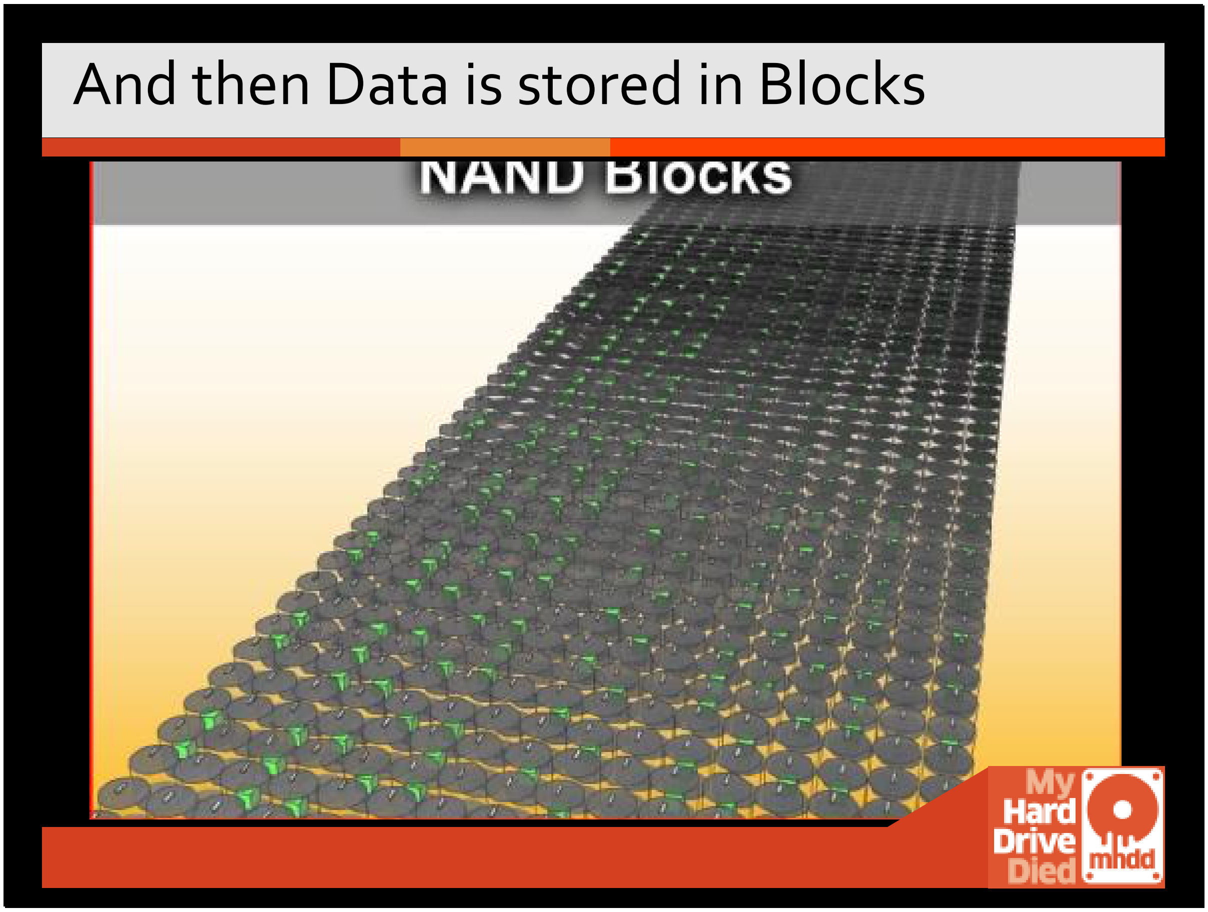

Blocks of data are written or erased in cycles: write cycles, read cycles, and erase cycles.= (they do different things.) When you release a cell (gate), it needs to be in a block size, and that is defined by NAND/or the manufacturer.

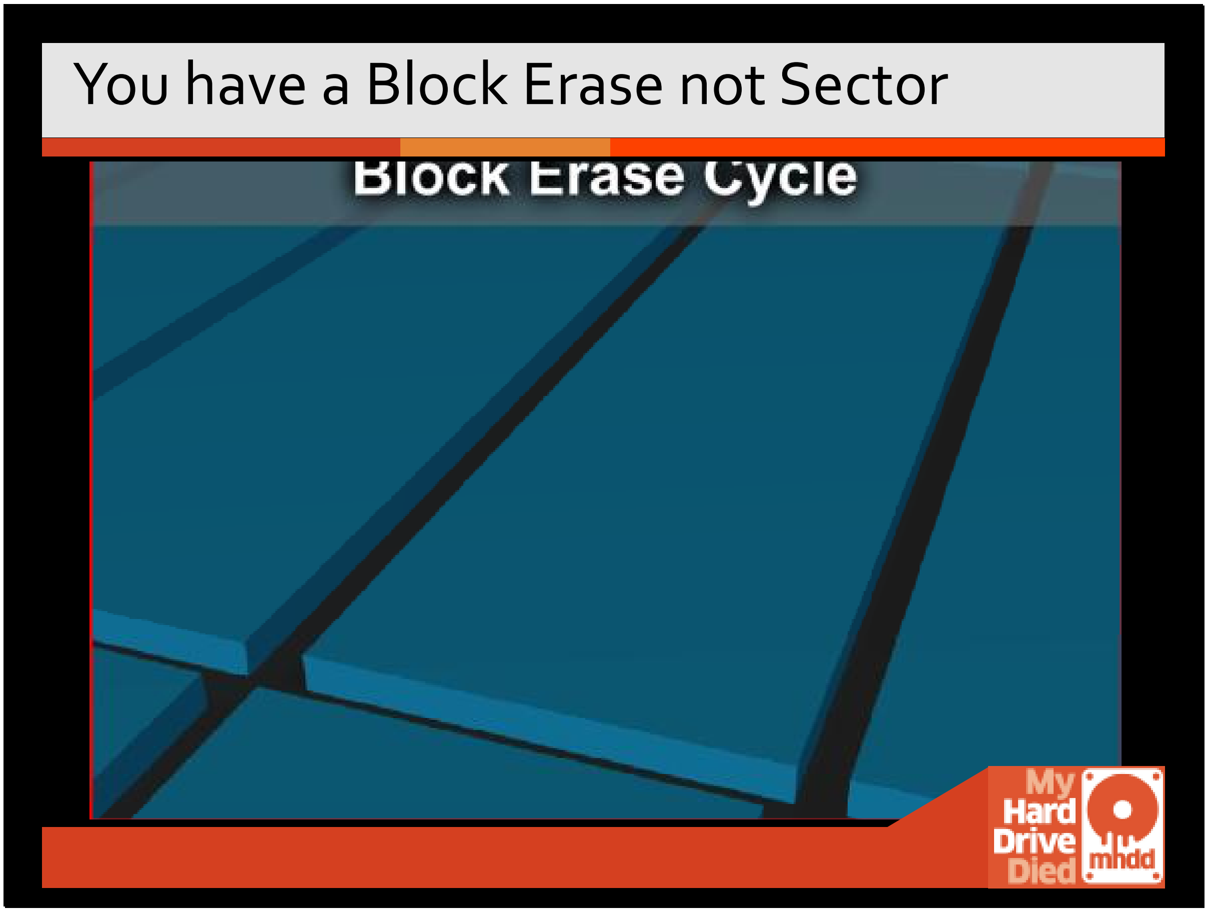

Dr. Fujio Masuoka worked together with Dr. Masahiko Suzuki and came up with the title 'Flash Memory' during the "Drain function".

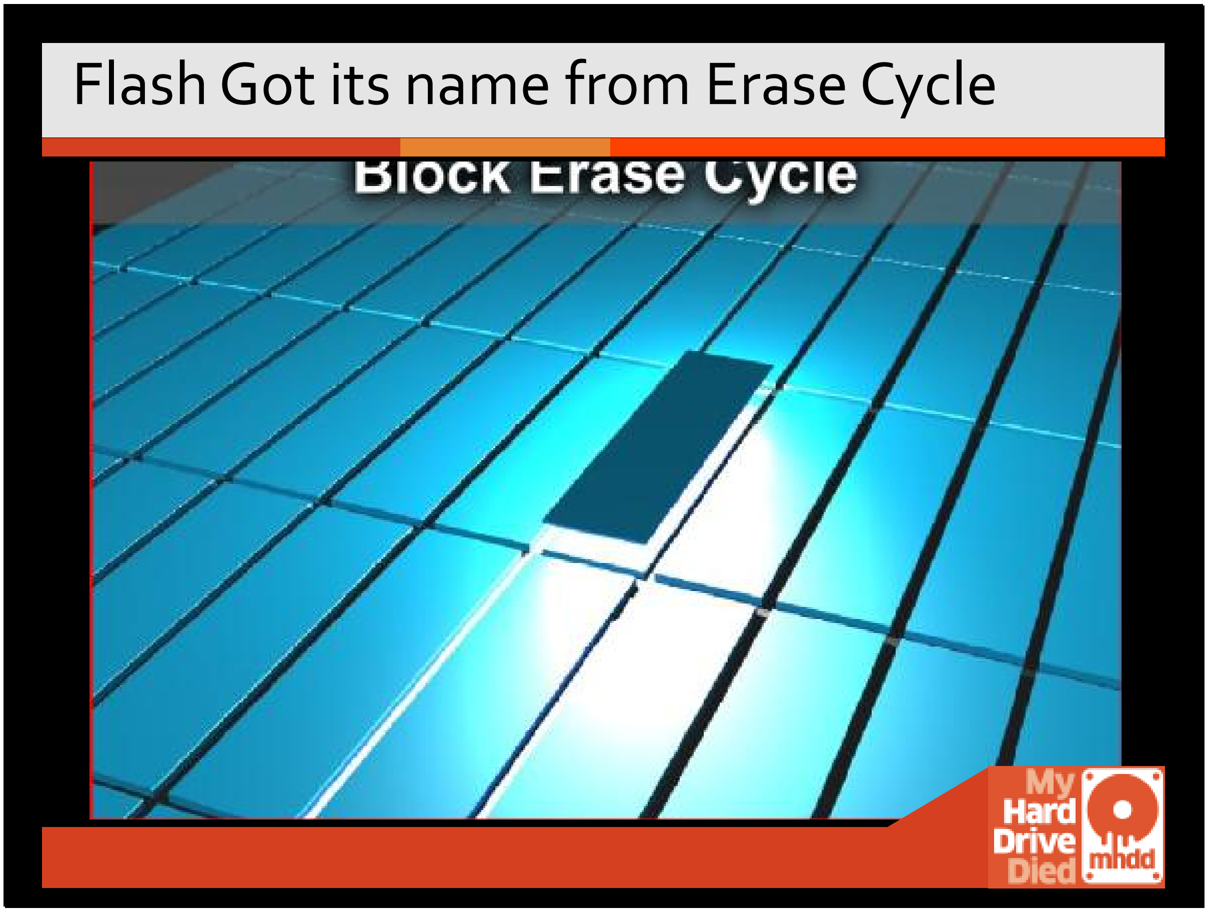

Plastic does not chip or pop off, but there is a flash as the electrons are released in an erase cycle, causing Dr. Suzuki to state it look liked the flash of a camera. The nickname stuck, and we have Flash Memory.

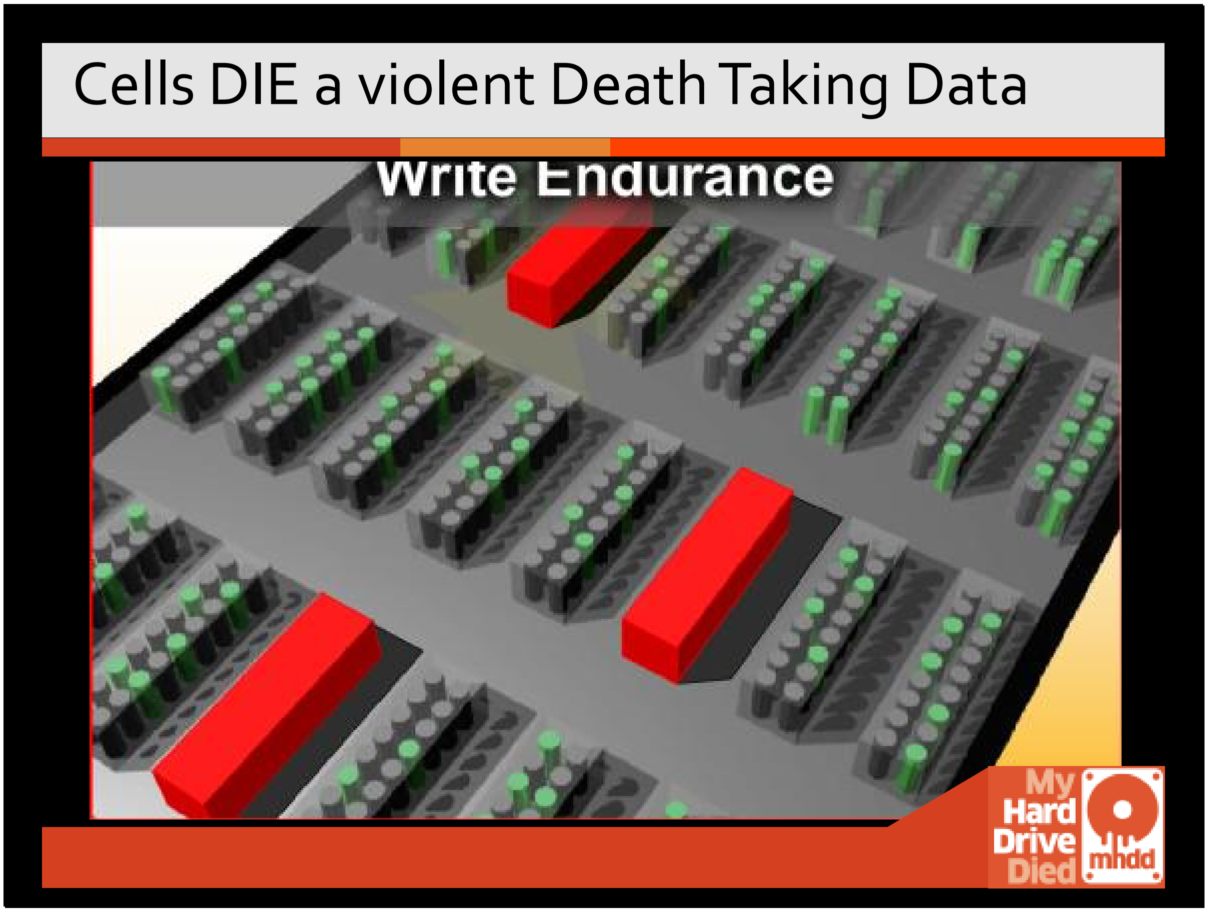

Cells will die causing blocks to be marked bad. In flash-based solid state disks the write endurance is the number of write cycles to a block of flash memory. Once you have met the write endurance limit, the disk may become unreliable or unable to use any of the cells.

In flash-based solid state disks the write endurance is the number if write cycles to a block of flash memory. Once you have met the write endurance limit, the disk may become unreliable or unable to use any of the cells.

Do not believe Marketing material that says it last 50 years, It will not last under constant USE. Destruction of the cells: As you keep stuffing them in there they will get damaged, and eventually, after a long time being in there, they escape anyway and damage the cell. It is still questionable if cells are reusable, because no one has had them for more than 10 years.

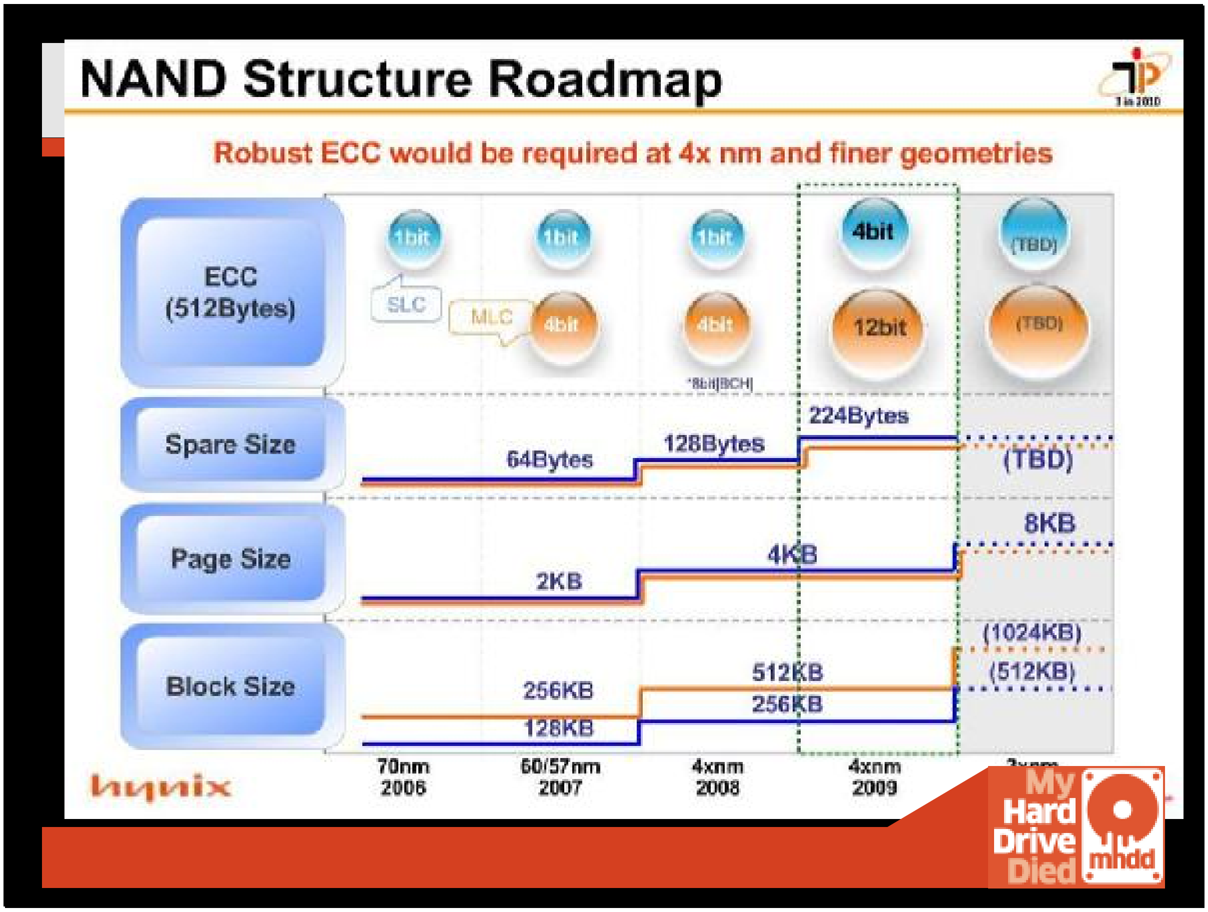

Presentation by Arun Kamat on Webcast 2008
"2012 is only 3 years away"
https://www.denali.com/ui/tv/Webcast?s=webcast2008_hynix2012
Used with Permission for Scott Moulton's Class

Currently we are seeing about a 4-bit ECC requirement for NAND. You will go from a 4-bit ECC to a 12-bit ECC for Triple Bit, and the Spare Sizes have to increase as well, so we will go from a www.MyHardDriveDied.com & Forensic Strategy Services, LLC. @ Scott A. Moulton Section size of 2k to 4k.

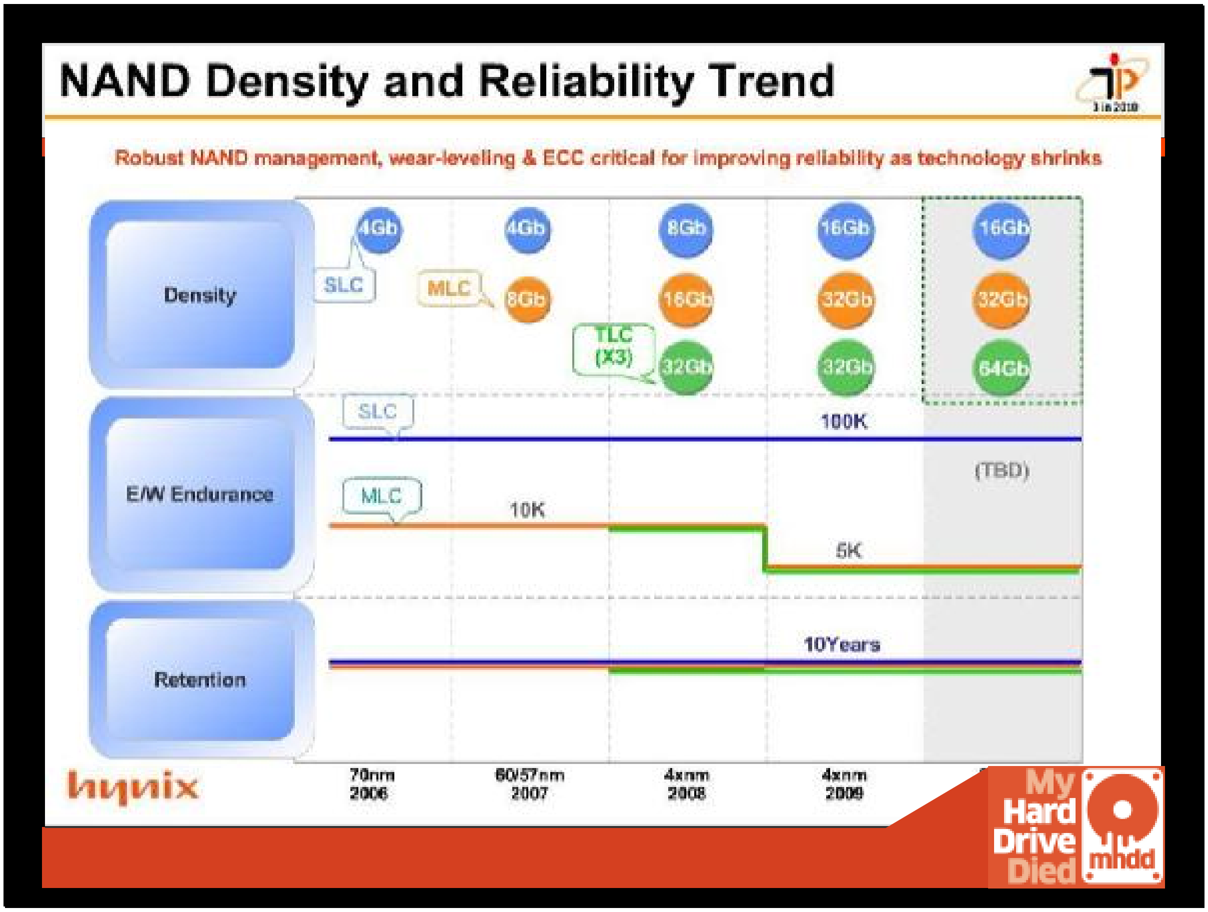

Presentation by Arun Kamat on Webcast 2008
“2012 is only 3 years away”
https://www.denali.com/ui/tv/Webcast?s=webcast2008_hynix2012
Used with Permission for Scott Moulton’s Class

By 2010 we expect to see a growth from TLC (triple layer cell) x4-bit technology. SLC has a 100k Cycle. The MLC will go to a 5k endurance cycle for the high density products.

Quoted Below. Good ones are sorted and sold to the big companies, the bad Flash is sold in Retail chains in the junk you know as Memory Sticks.

Bad ones go to the consumers - good ones goes to industrial.

This quote was on Adtron's site until I started using it in presentations, recently pulled.
http://www.adtron.com/products/flash-disk.html

Archive.org
http://web.archive.org/web/20070115073812/http://adtron.com/products/flash-disk.html

"The first step is to evaluate only the highest quality NAND flash available on the market today. Changes in flash die geometries, cell structure and address control functions impact yield rates from the wafer during the fabrication process. The concept that certain spots on a wafer are more likely to produce a higher quality chip is prevalent throughout the semiconductor industry. When flash memory wafers are tested and probed, a distribution of bad, weak, and strong die are identified across the wafer. Bad cells are marked, and the remaining cells are sorted into consumer and industrial quality flash. Adtron selects only industrial quality flash identified at the die sort stage. Consumer grade flash makes it into flash disks that are mass marketed through retail outlets. After identifying the industrial-grade flash memory, Adtron continues the characterization process by independently testing and validating the manufacturer's specifications in our facilities in Phoenix Arizona. At this stage the flash memory is put through extensive write endurance testing and the disturb characteristics are evaluated to further ensure the reliability of an Adtron flash disk."

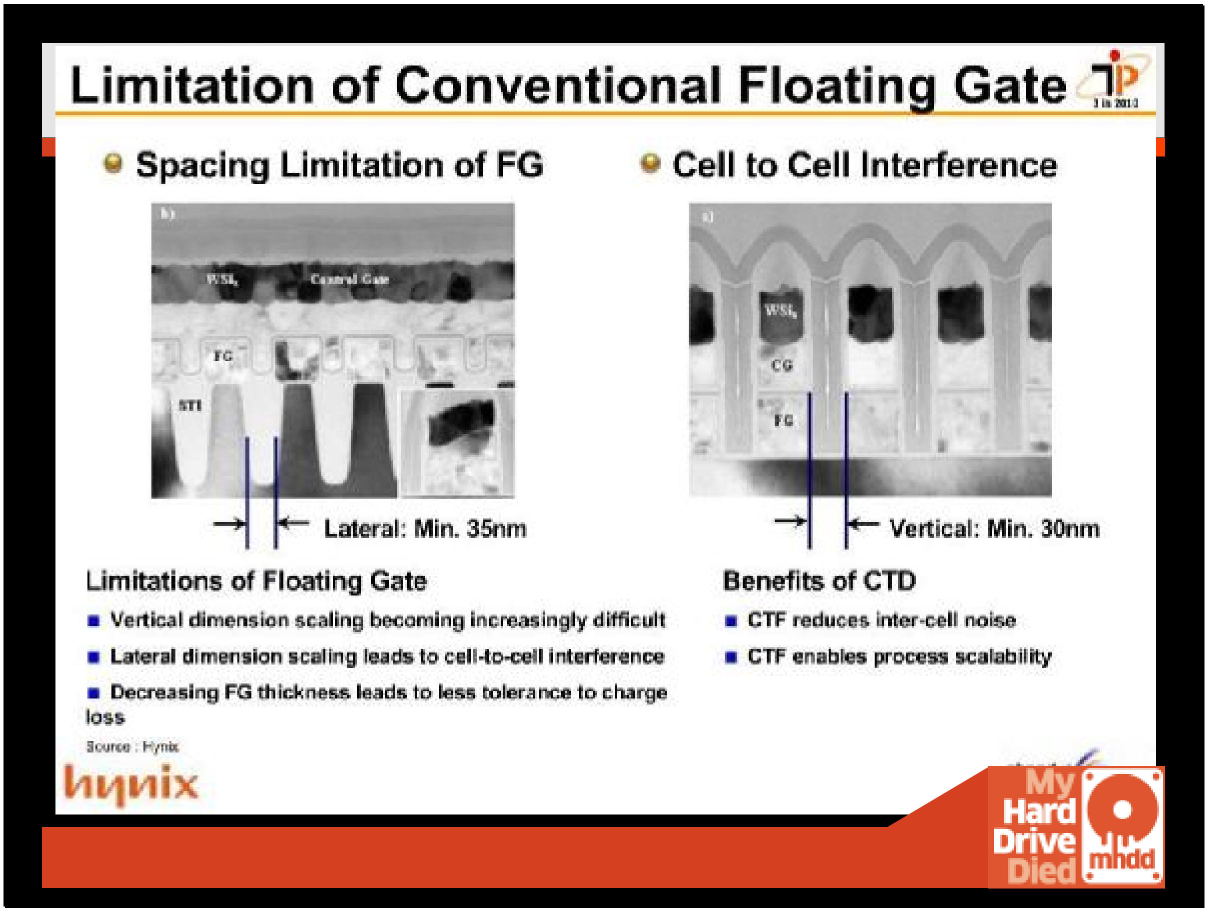

Presentation by Arun Kamat on Webcast 2008
"2012 is only 3 years away"
https://www.denali.com/ui/tv/Webcast?s=webcast2008_hynix2012
Used with Permission for Scott Moulton's Class

The spacing between two cells causes the bit flip problem, similar to Super-Paramagnetic problems with hard drives.

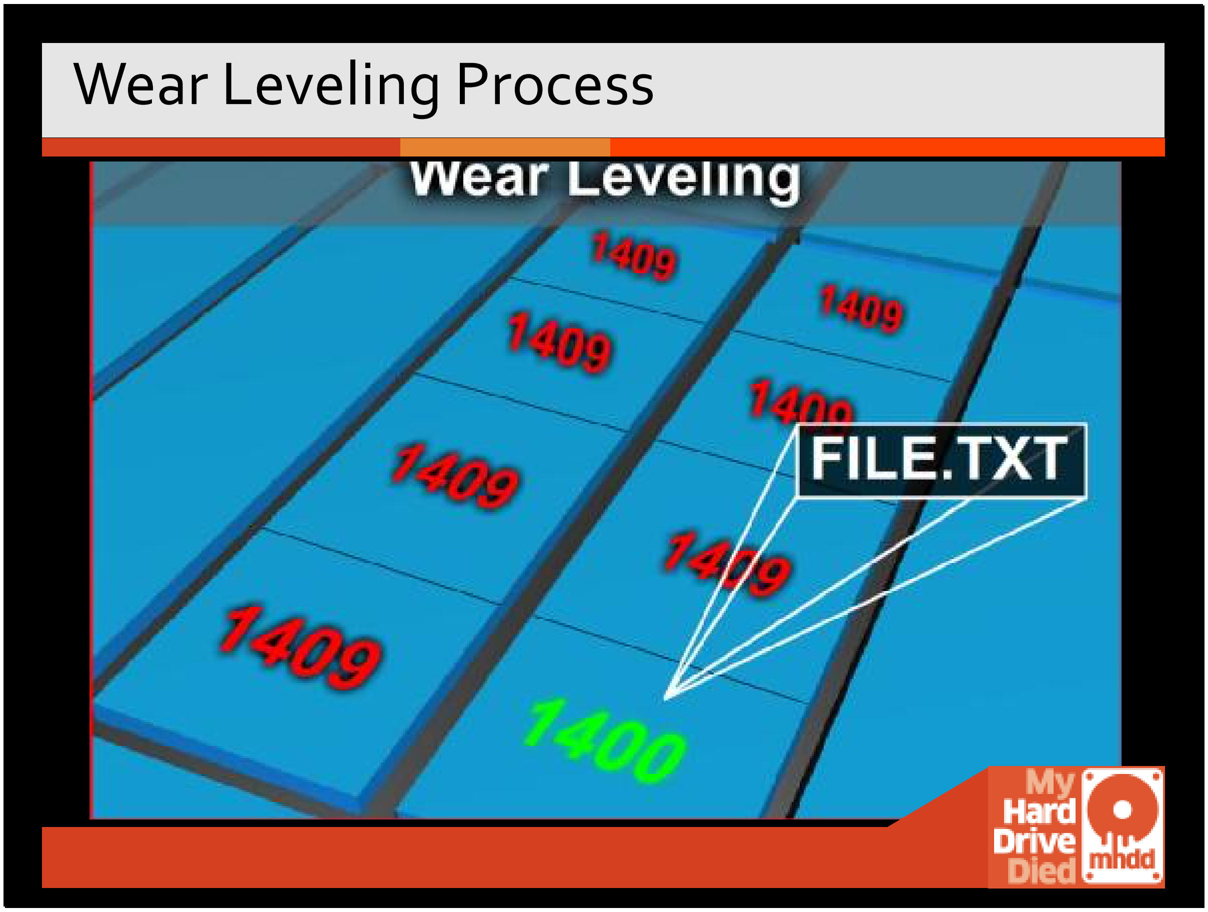

Wear leveling is a technique for prolonging the service life of some kinds of erasable computer storage media, such as flash memory.

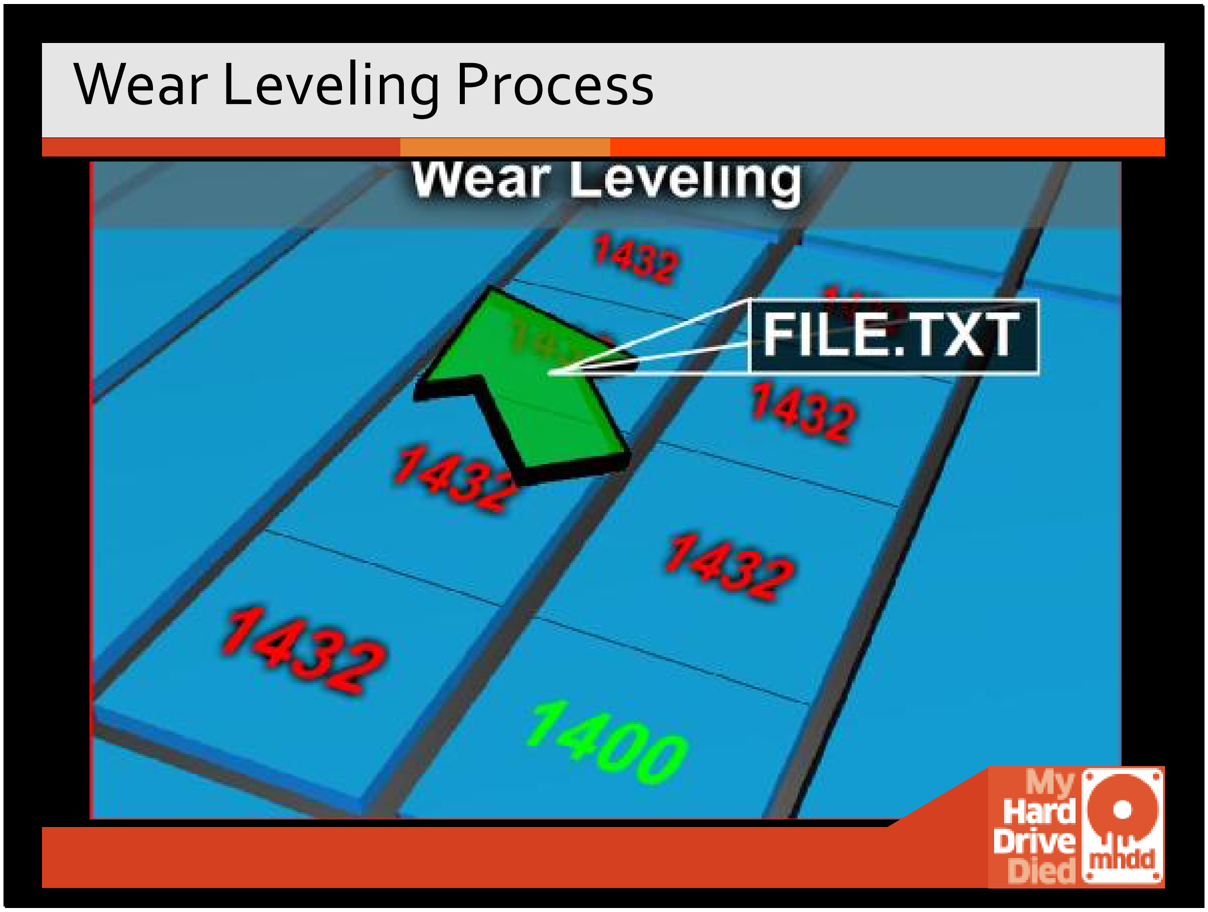

When content is changed it must be moved to a new location before the data can be saved, and then the data will be stacked in the garbage collection queue.

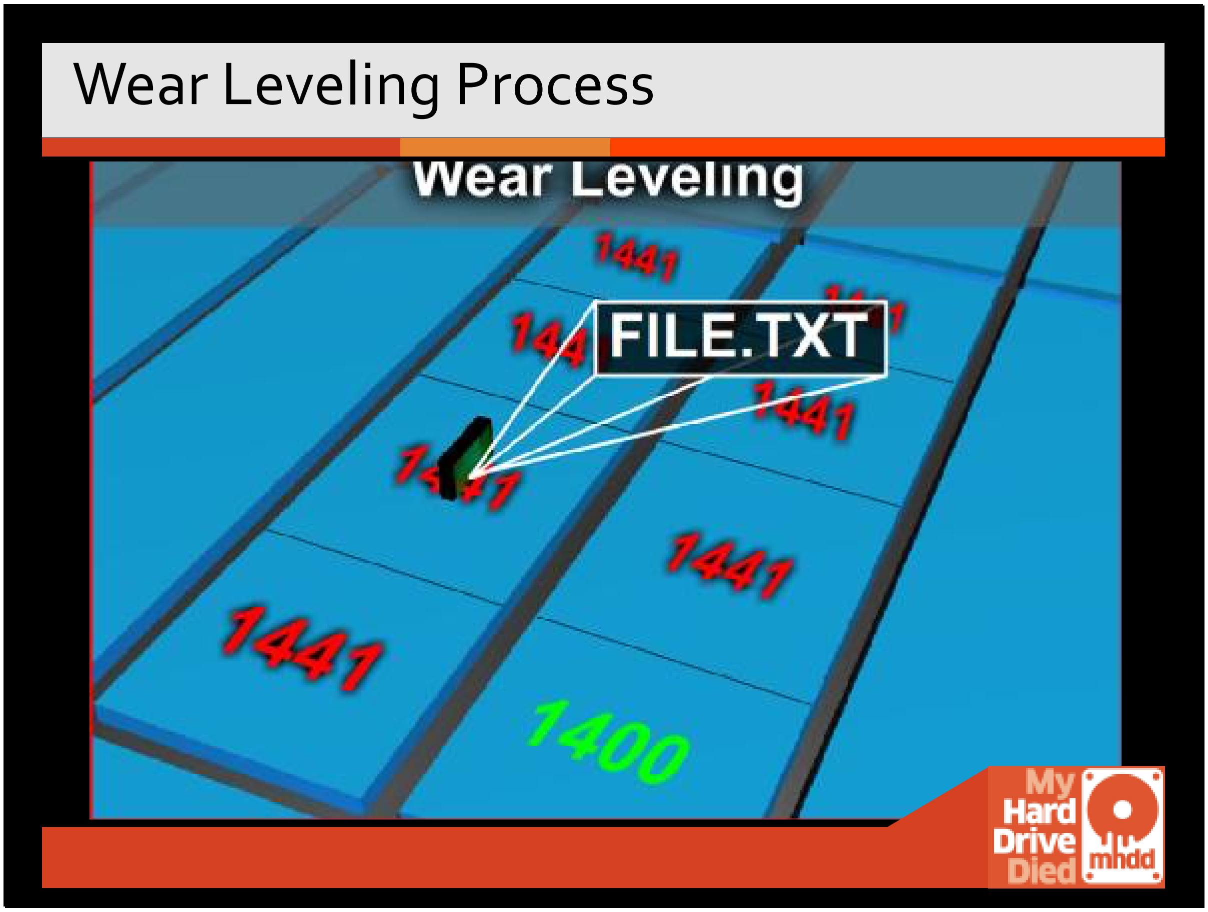

Once the data has been moved the original location is then free to be cleared.

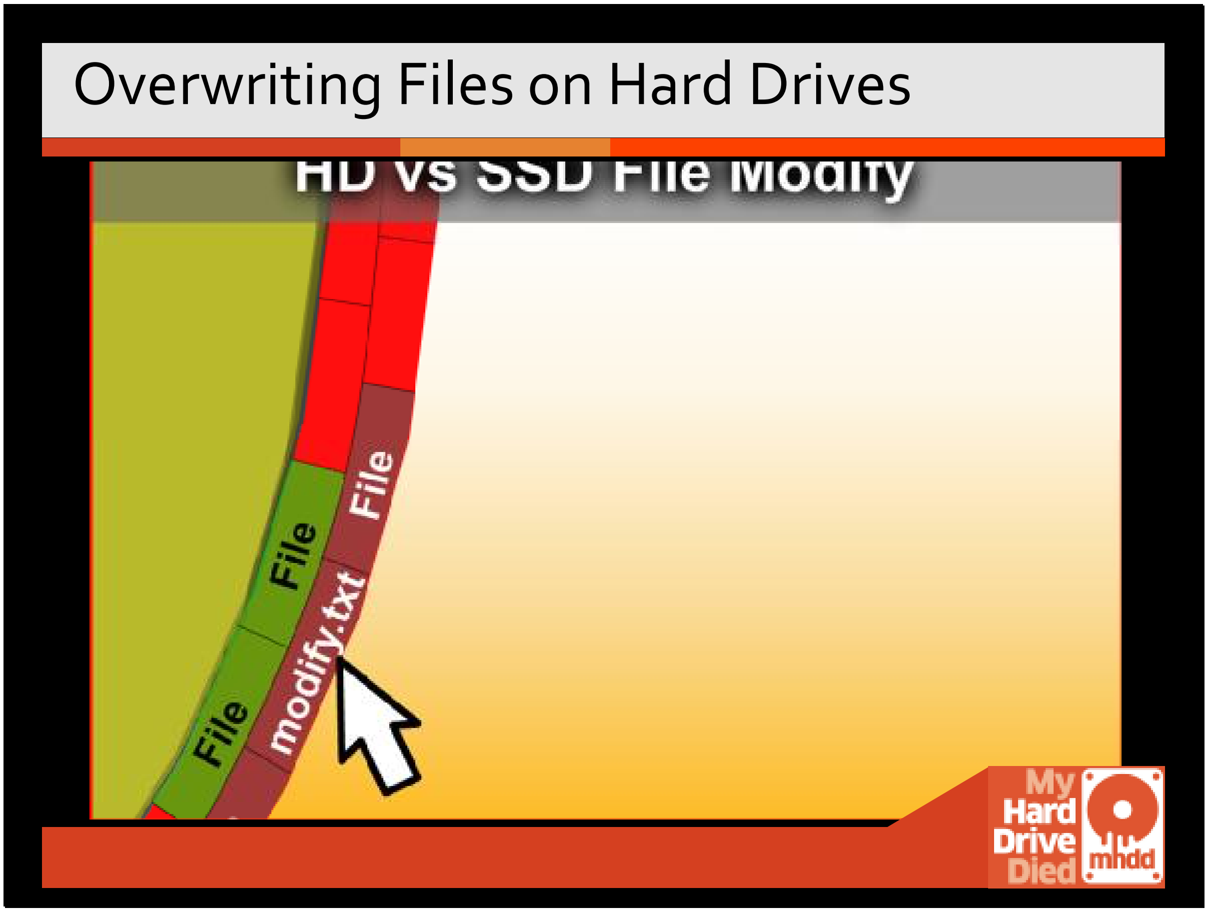

When you have documents or data on a hard drive it is possible to edit the content in place.

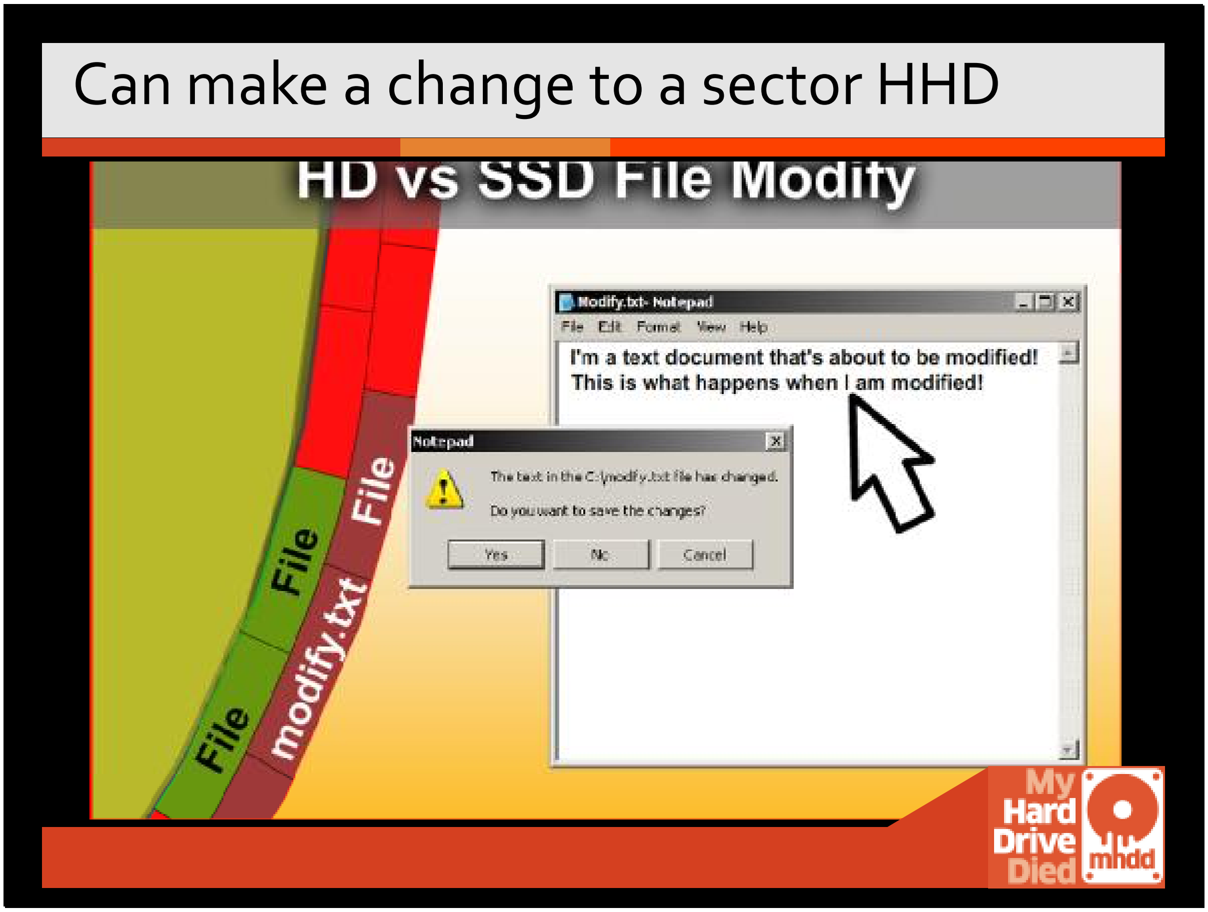

It is possible to make changes to a file (edit a sector in place) in a certain location on a hard drive. You can modify a file in place and not have to worry about it moving around.

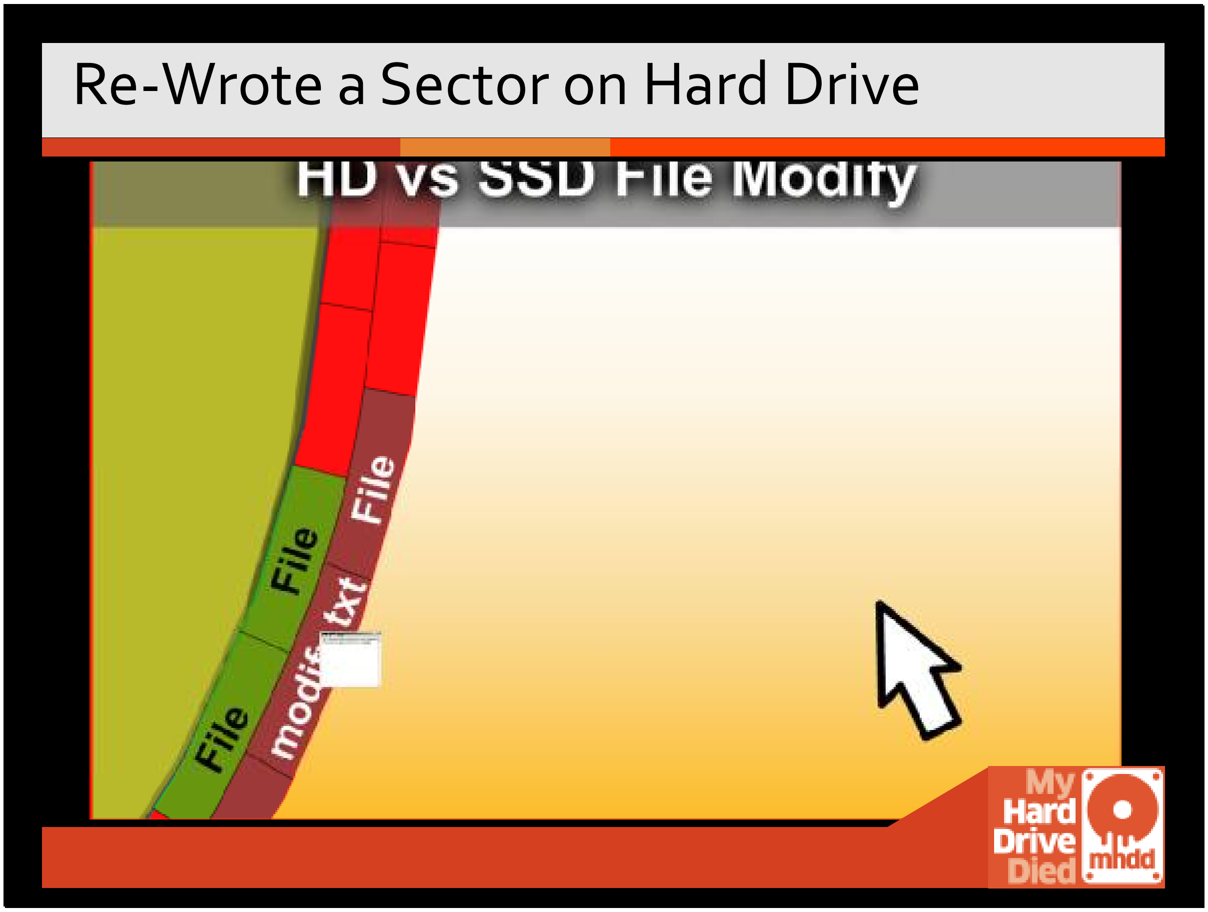

As data is saved it can be saved in the exact location (not always is it due to different circumstances with different software).

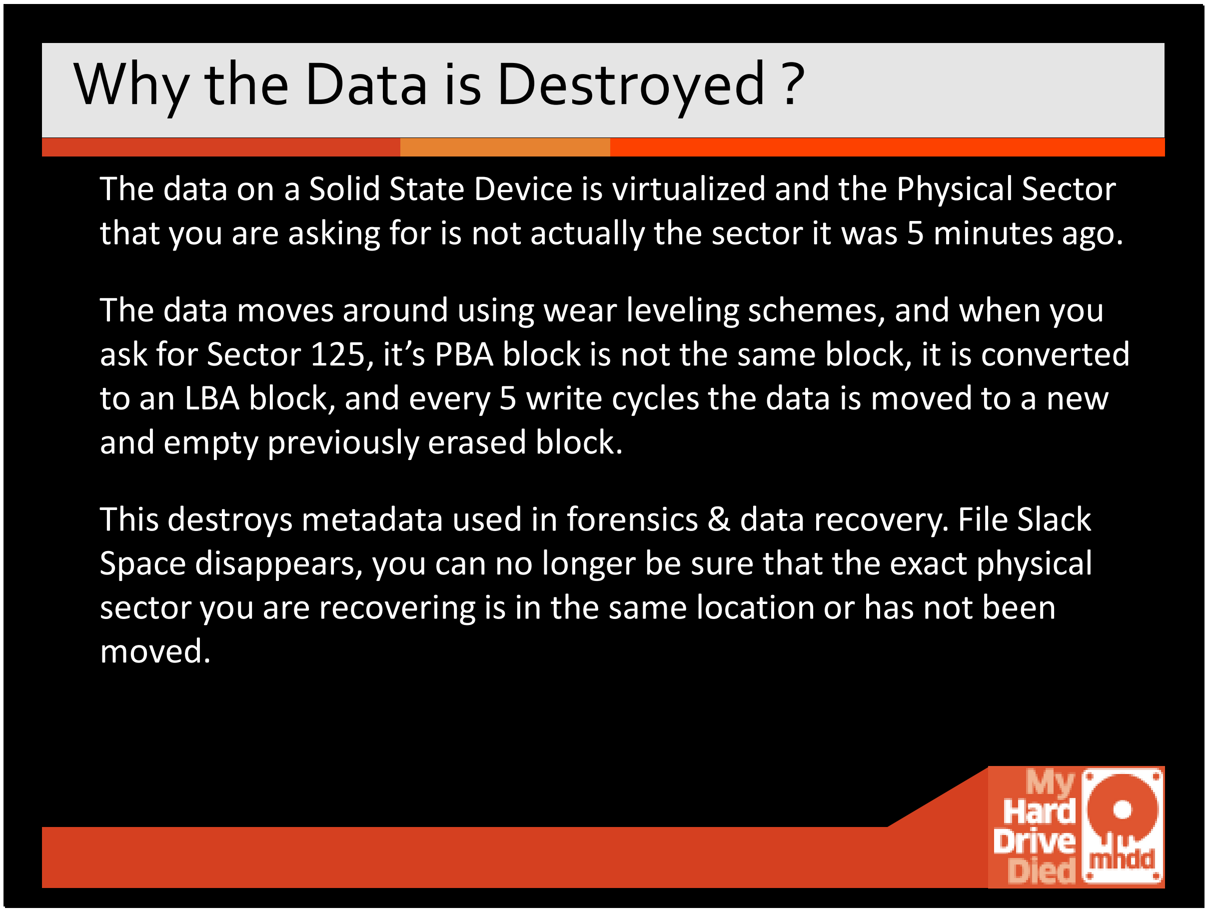

The data on a Solid State Device is virtualized, and the Physical Sector that you are asking for is not actually the sector it was 5 minutes ago. The data moves around using wear leveling schemes, and when you ask for Sector 125, it's PBA block is not the same block, it is converted to an LBA block and every 5 write cycles the data is moved to a new and empty previously erased block. This destroys metadata used in forensics & data recovery. File Slack Space disappears, and you can no longer be sure that the exact physical sector you are recovering was in the same location or has not been moved.

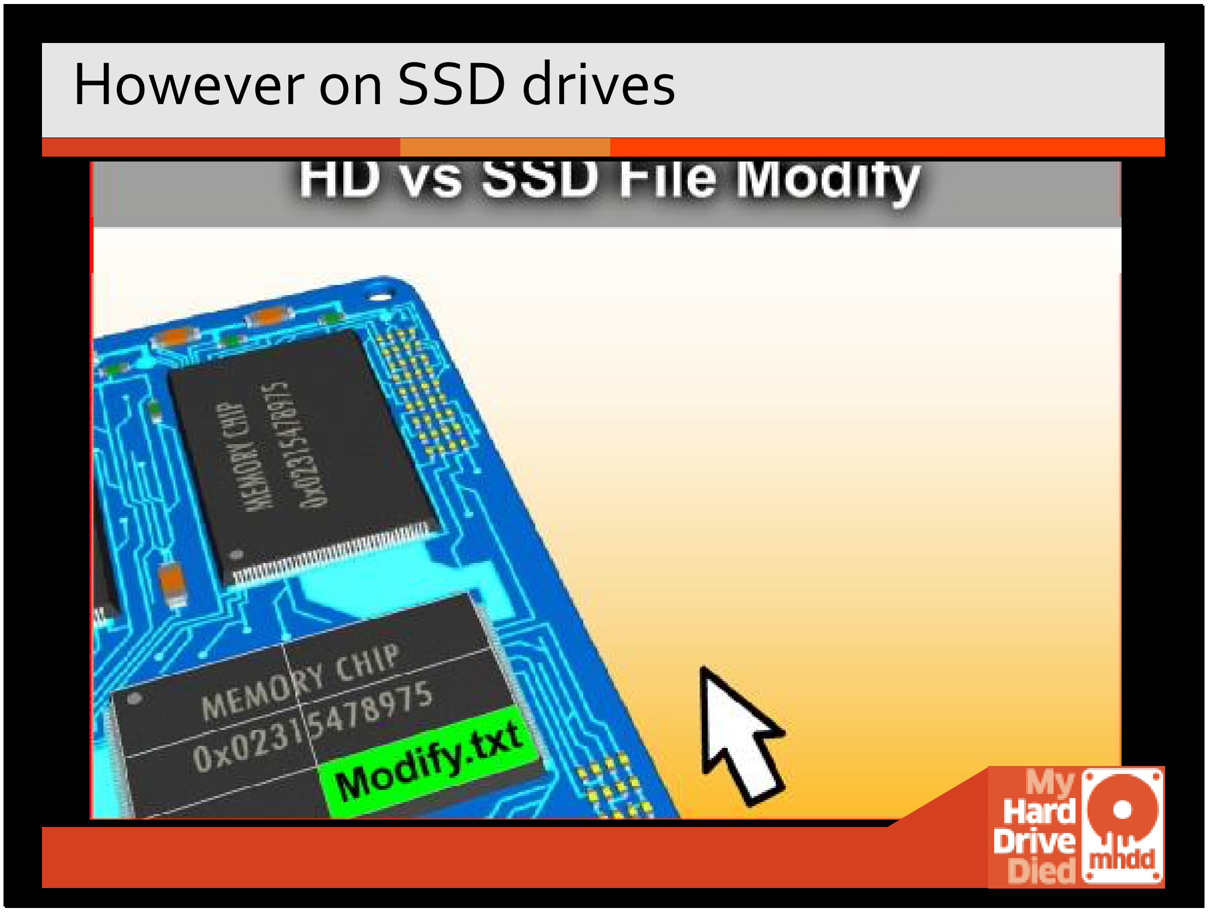

You cannot make a change (cannot edit a file in place in an SSD) within those cells without erasing the state of those cells. Legacy: PBA (Physical Block Address) conversion to LBA (Logical Block Addresses). At lower level, there is a translation table, which converts Physical Block Address to the LBA Block Address, which then supplies the data from the new location.

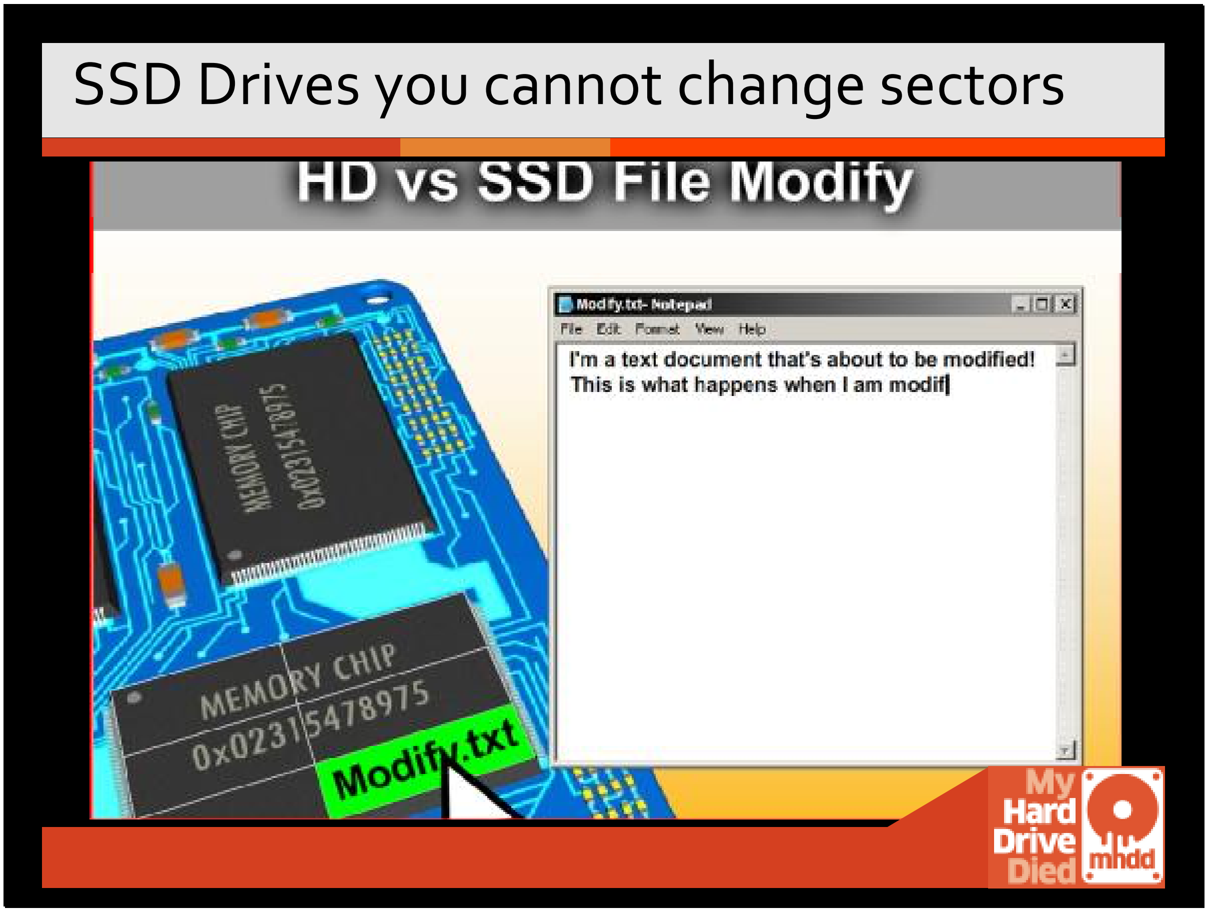

When the file is modified it will have to find a new home in which to be stored.

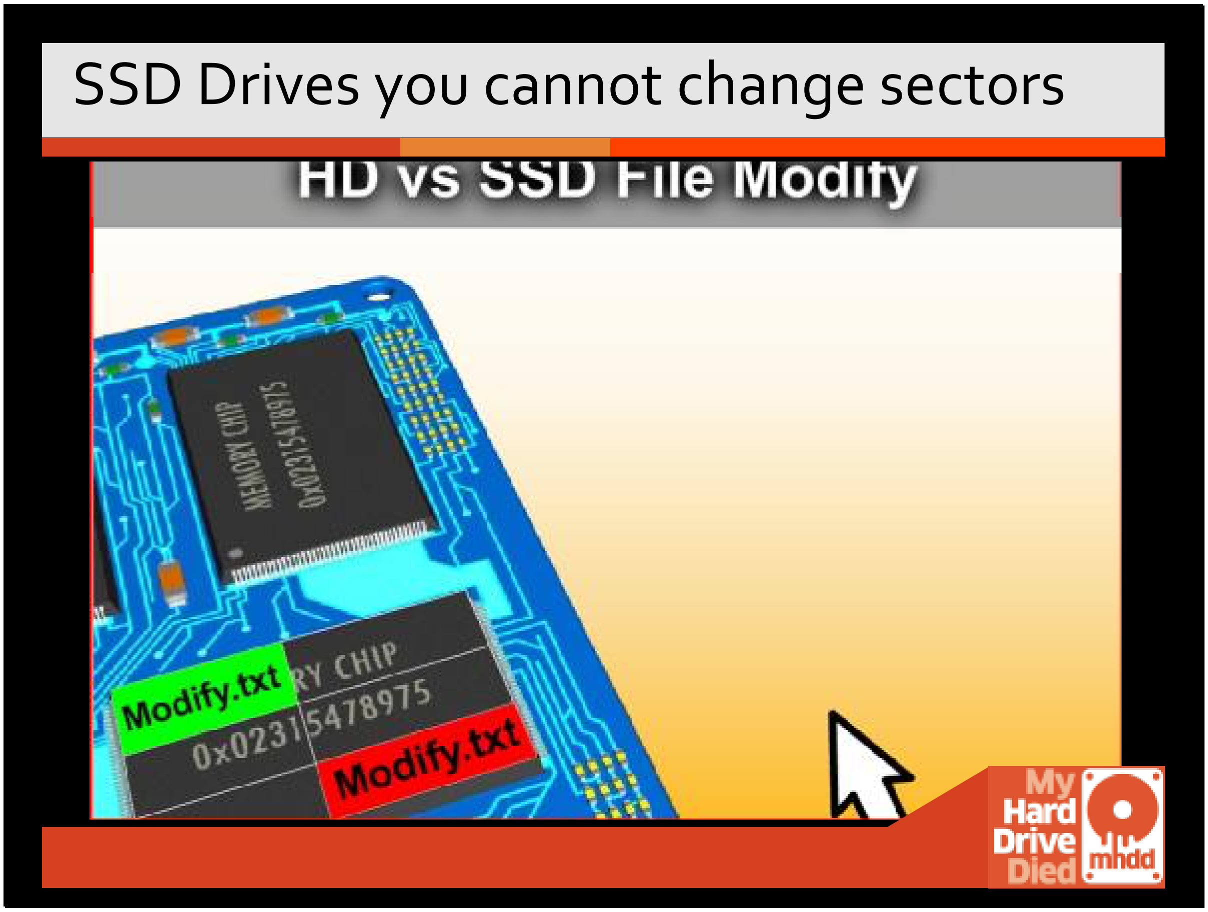

Now the modified file exists at a new location as it cannot be changed in its current location. Then the sector has to be added to the garbage collection routine due to the fact that it cannot be changed or erased in its location. It must be done on the new location.

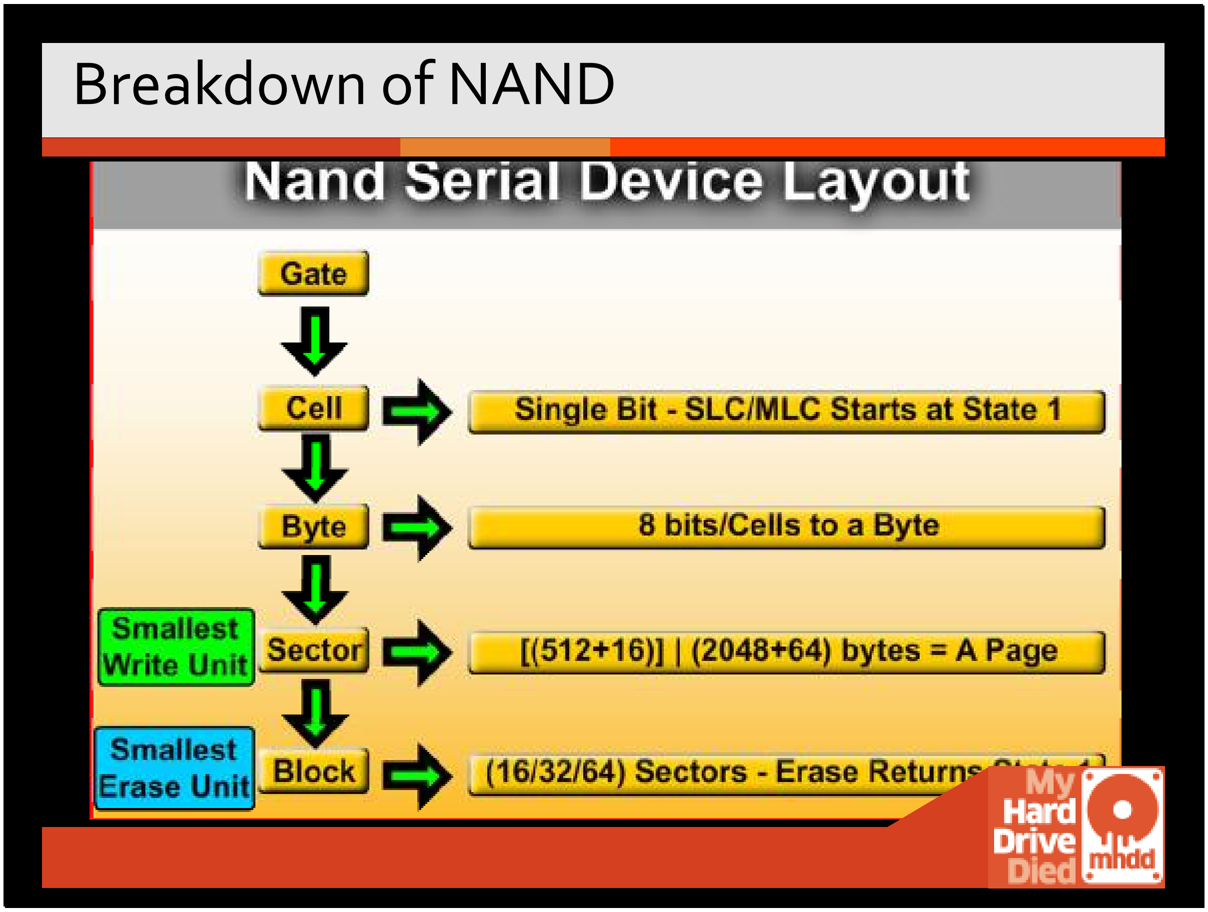

This is a basic chart of memory sizes of existing NAND. If a single bit fails in this block, the entire block is put in a bad block list (128K).

* Numbers in brackets are options chosen by the manufacture.

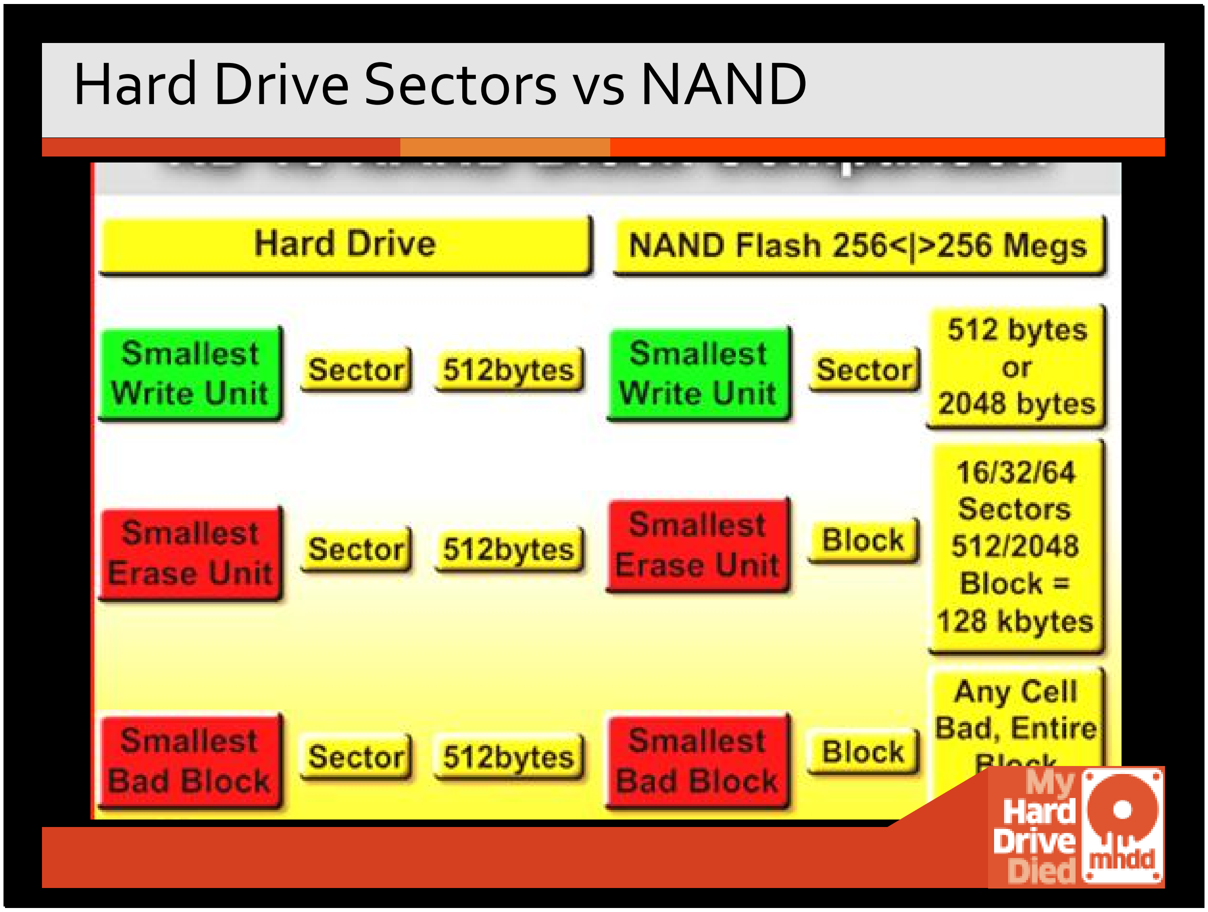

This is a comparison to hard drives and the data sizes, including the smallest bad block. Losing 128k every time one cell dies can cause the disk to shrink over time.

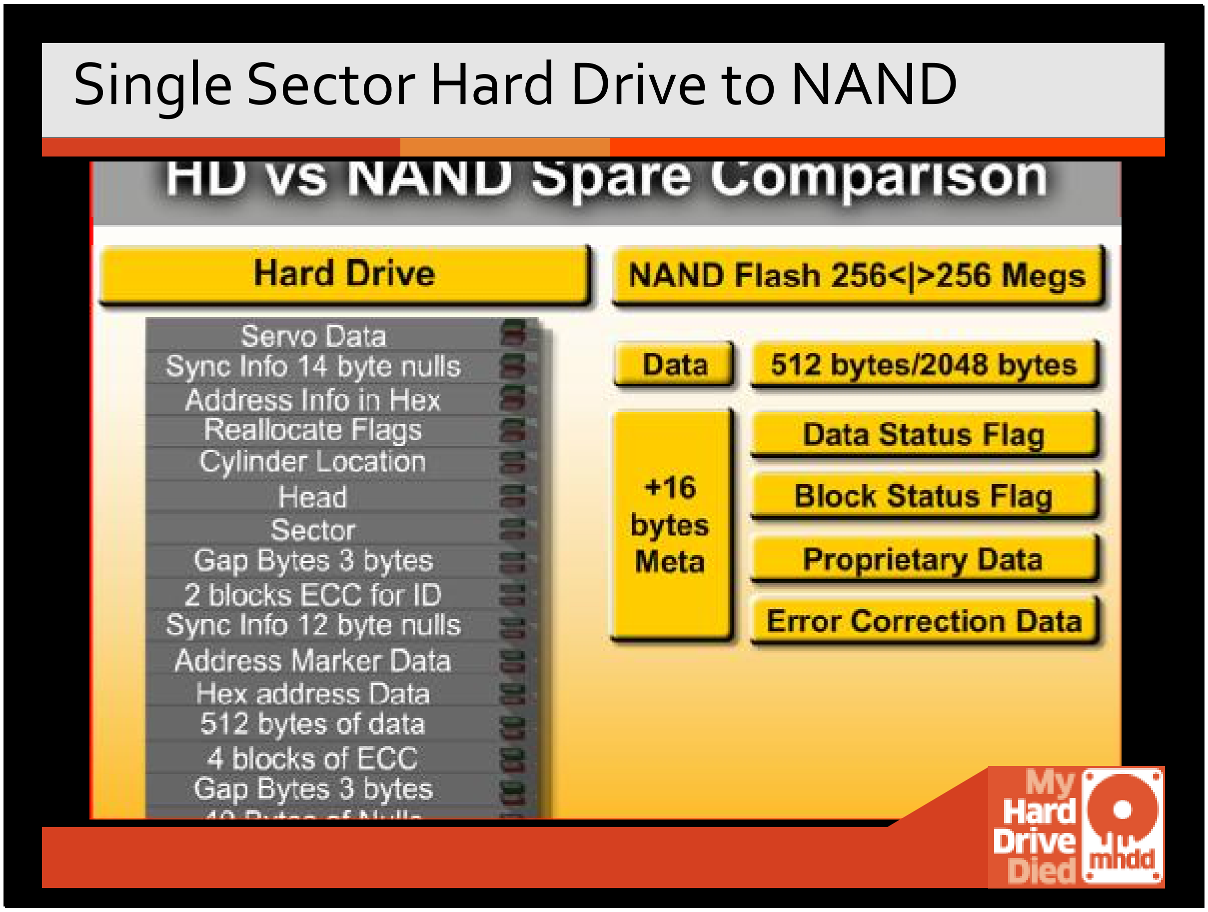

A sector on a disk, once all the content is added, is approximately 591 bytes. You can see on an SSD that it is much smaller.

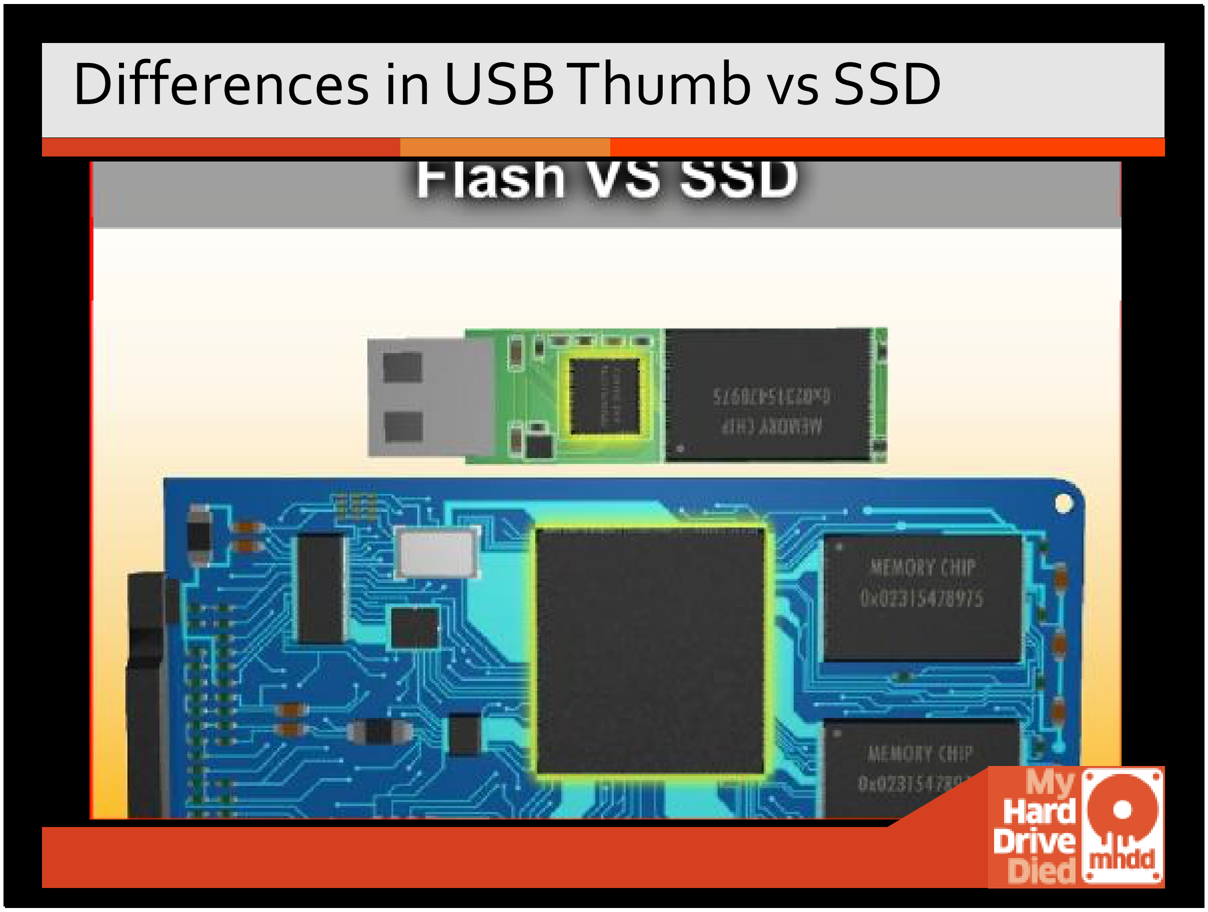

Similar to RAID by nature: Discreet controller and a Host Based Controller.

USB Memory stick is a host based controller. Host based controller uses your processor.

The SSD is more like a Discreet controller, as once power is supplied the device will continue to process data and keeps on going with its garbage collection routines, etc.

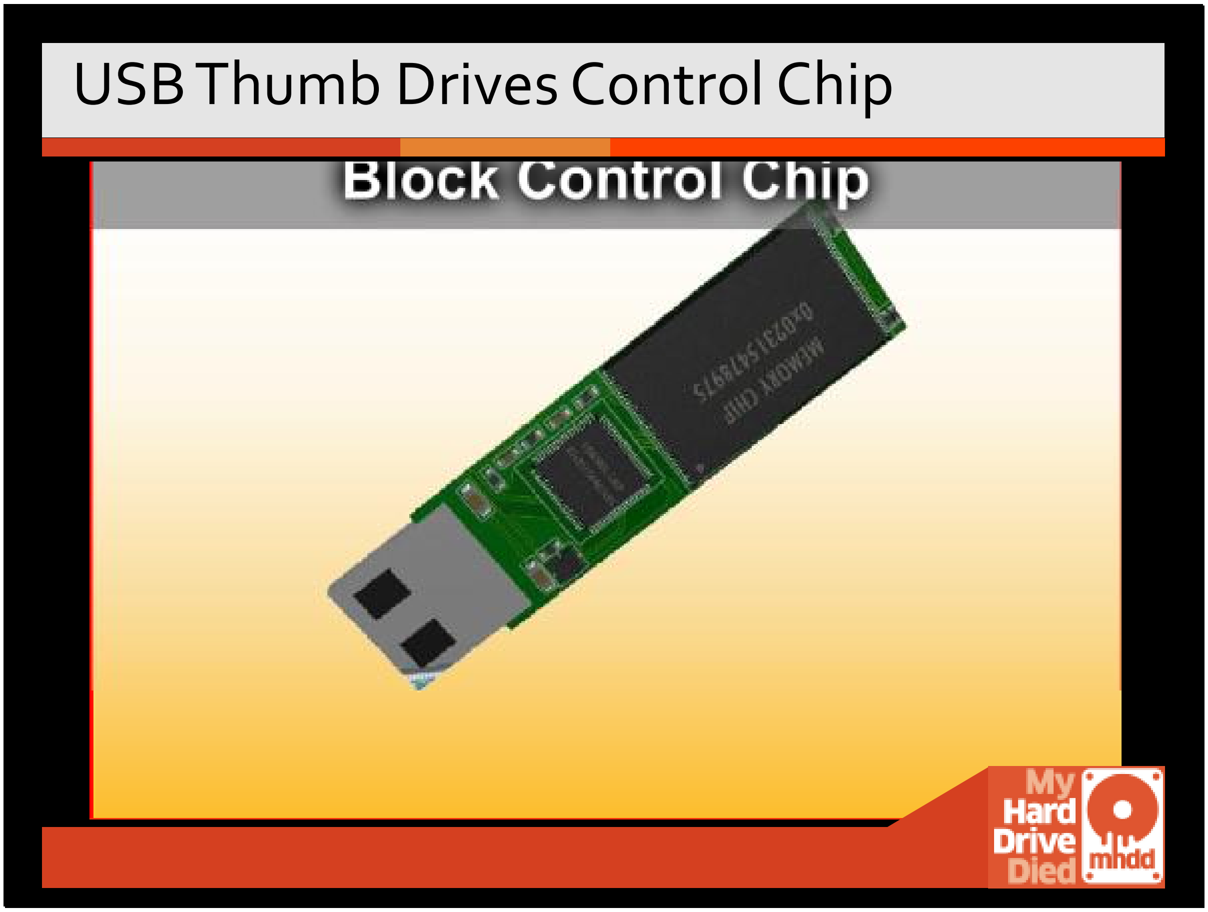

Firmware. Loads when it’s power on.

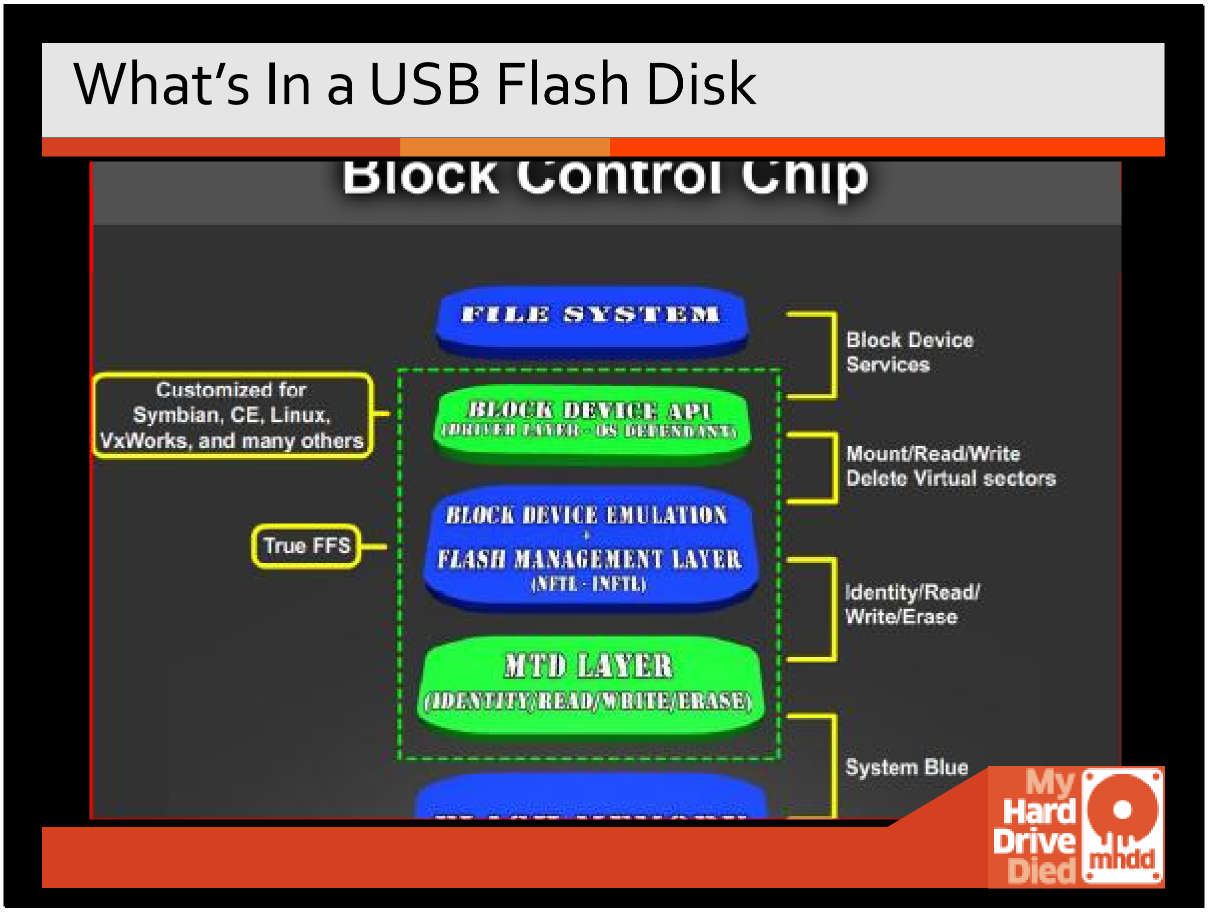

SanDisk bought M-Systems. This is a view of the basic functional implementation of a product by SANDISK to protect the contents of flash memory and is called True FFS.

Their IP property that SanDisk owns from SanDisk wrote all the stuff that Samsung uses for wear leveling. Every operating system has a mass storage driver of some sort and it is used to manage the data on the device.

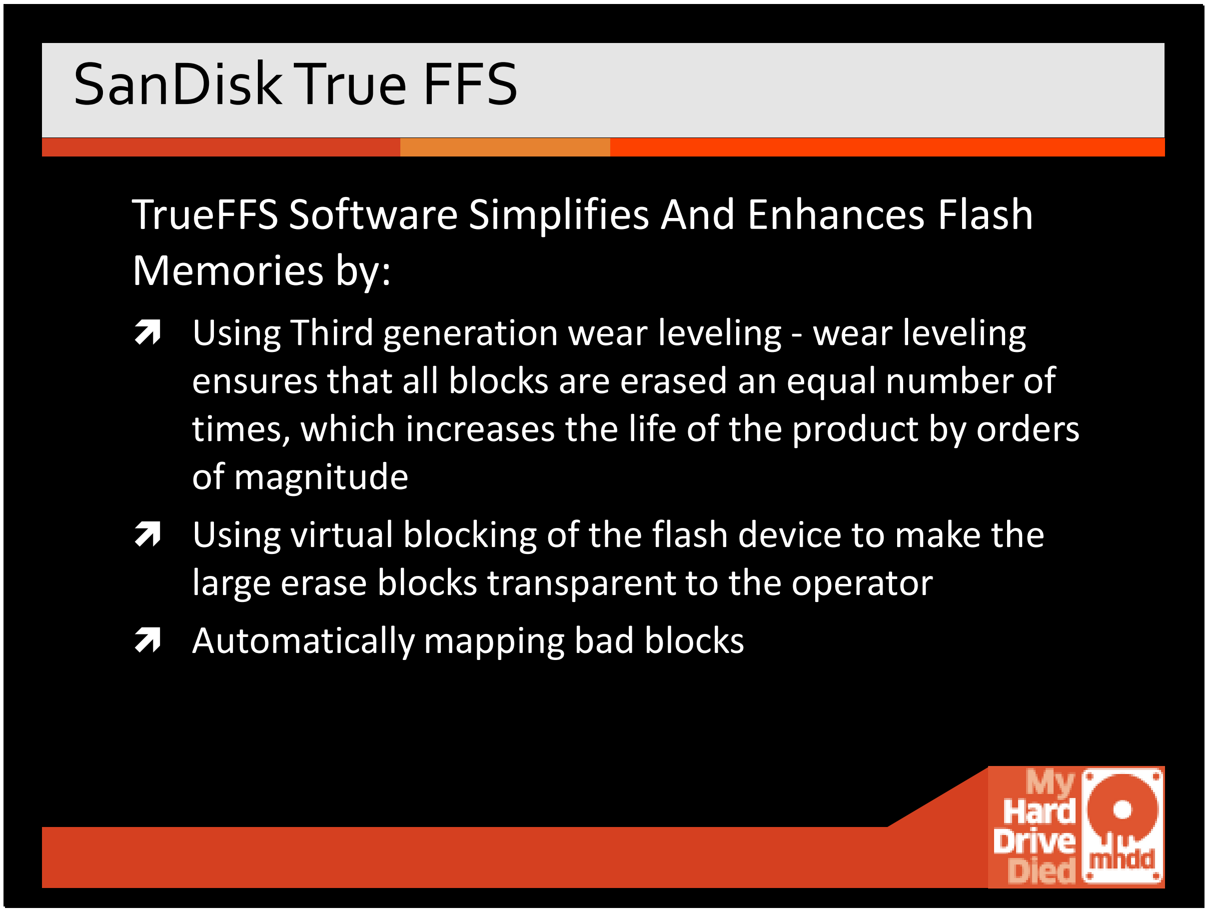

TrueFFS Software Simplifies And Enhances Flash Memories by:

- Using Third generation wear leveling. Wear leveling ensures that all blocks are erased an equal number of times, which increases the life of the product by orders of magnitude.
- Using virtual blocking of the flash device to make the large erase blocks transparent to the operator.
- Automatically mapping bad blocks

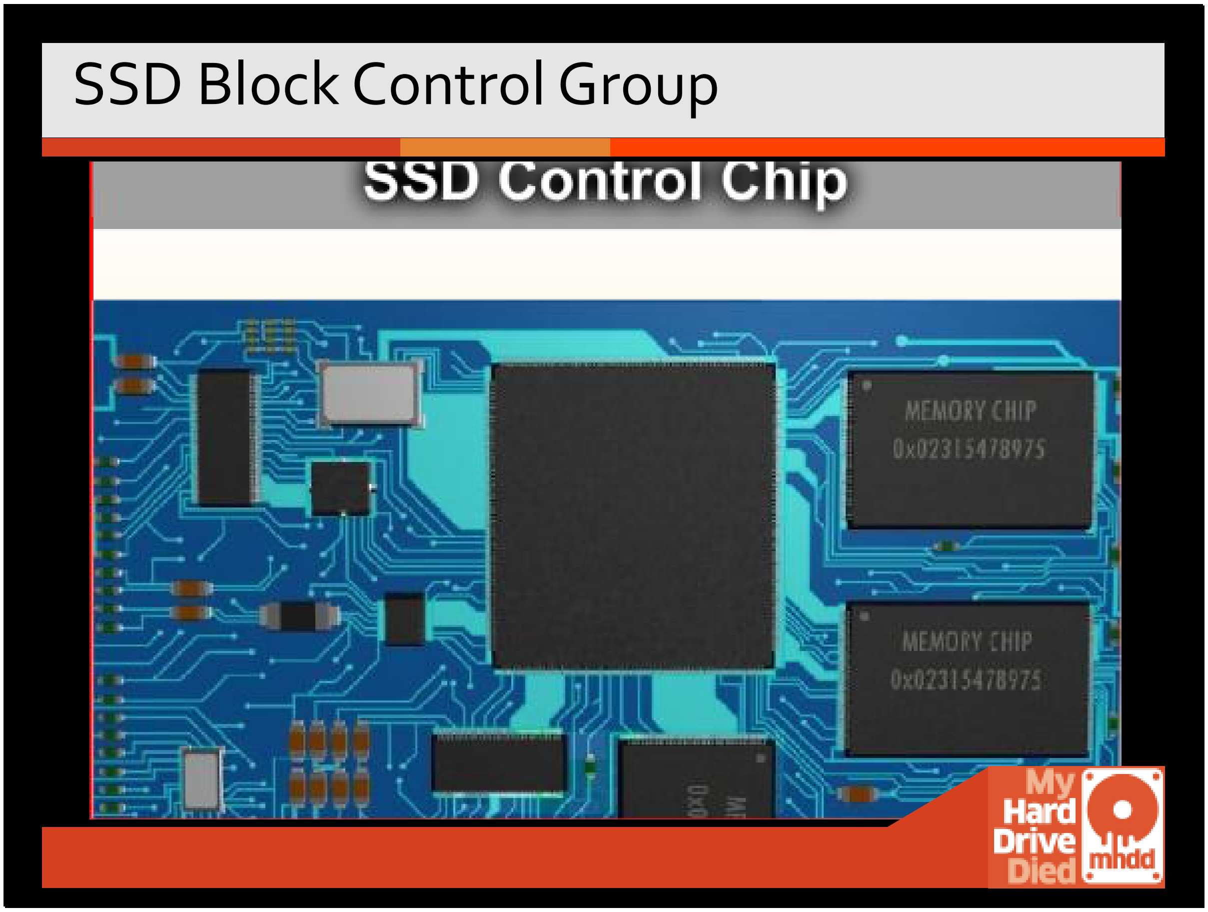

The control chips for SSD drives are much different that USB and have the ability to process data on their own.

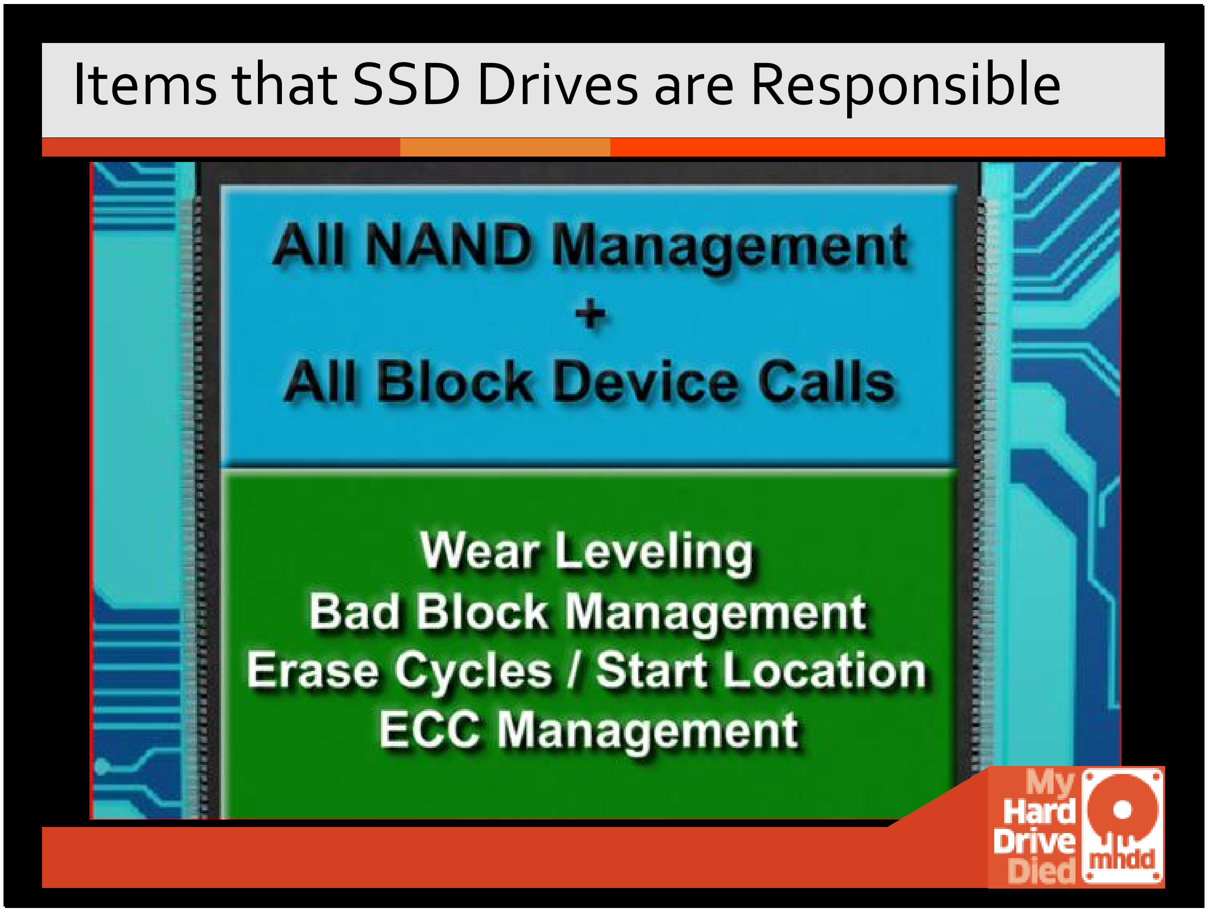

These are the expectations of SSD's from the vendors. Each SSD must be able to do all these on their own in order to be properly implemented.

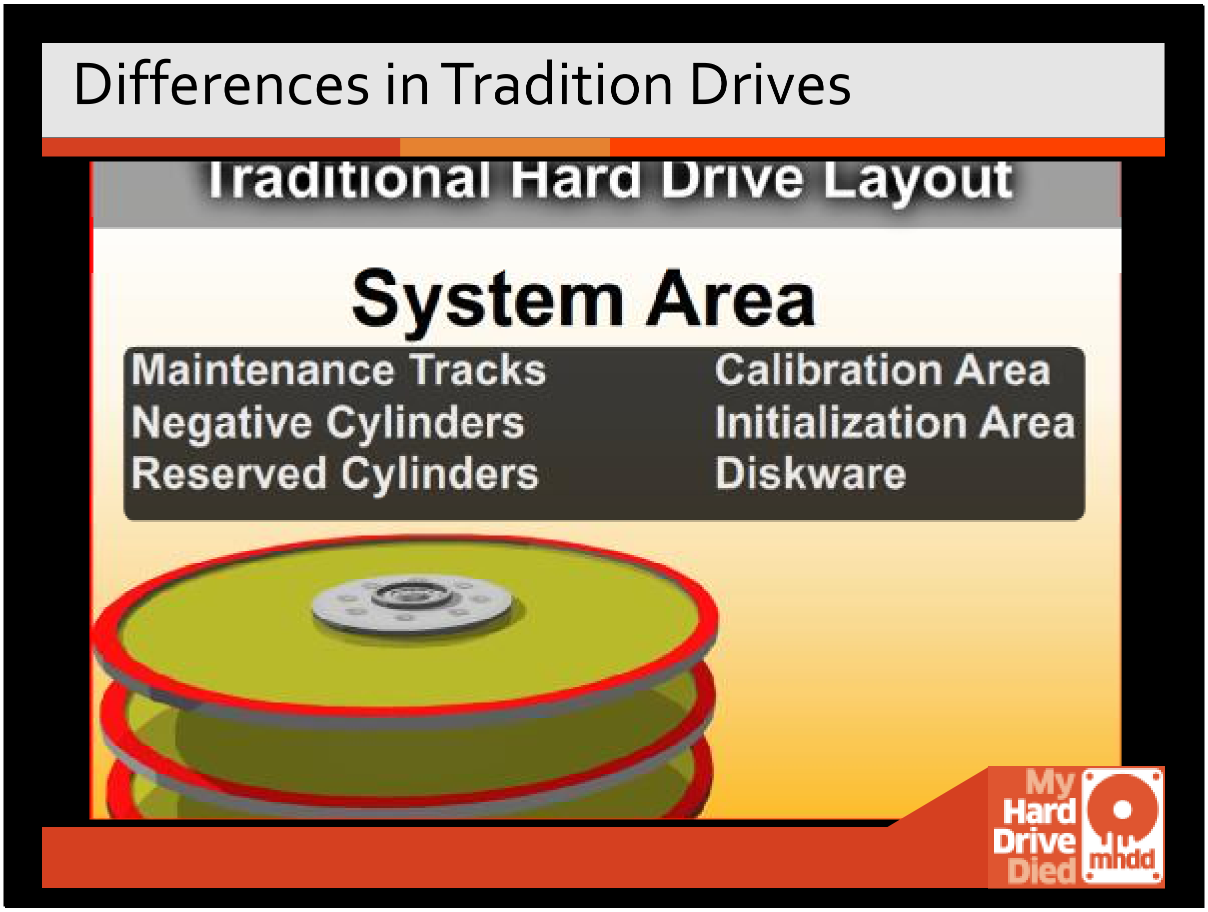

This is a comparison for what is changing on Flash that are legacy items, which seem to no longer be needed on hard drives. Maintenance Tracks is the oldest name. Negative Cylinders is the primary word used; it will read backwards.

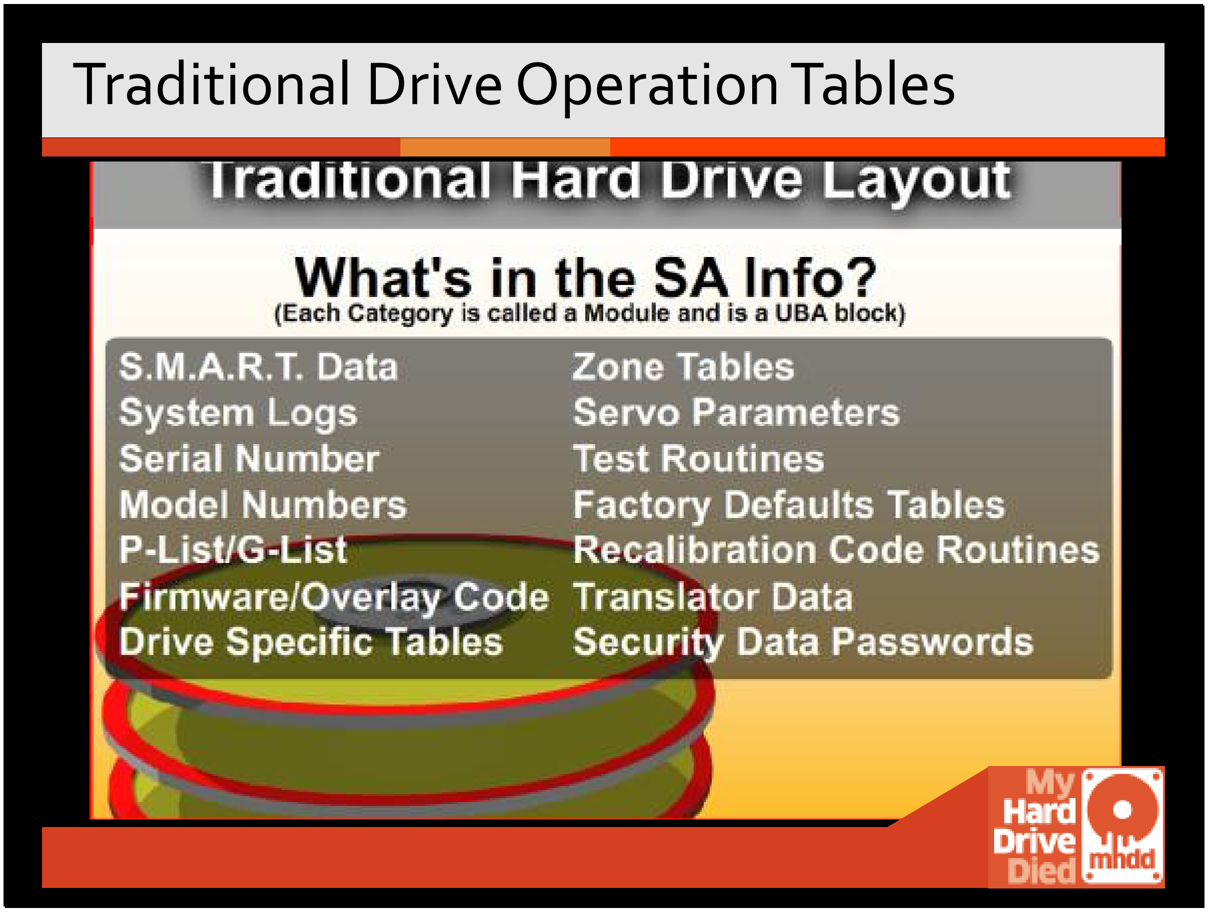

No point to many of these legacy items in SSD's. No need for actual P-List or Recalibration Routines because there's no head to move around.

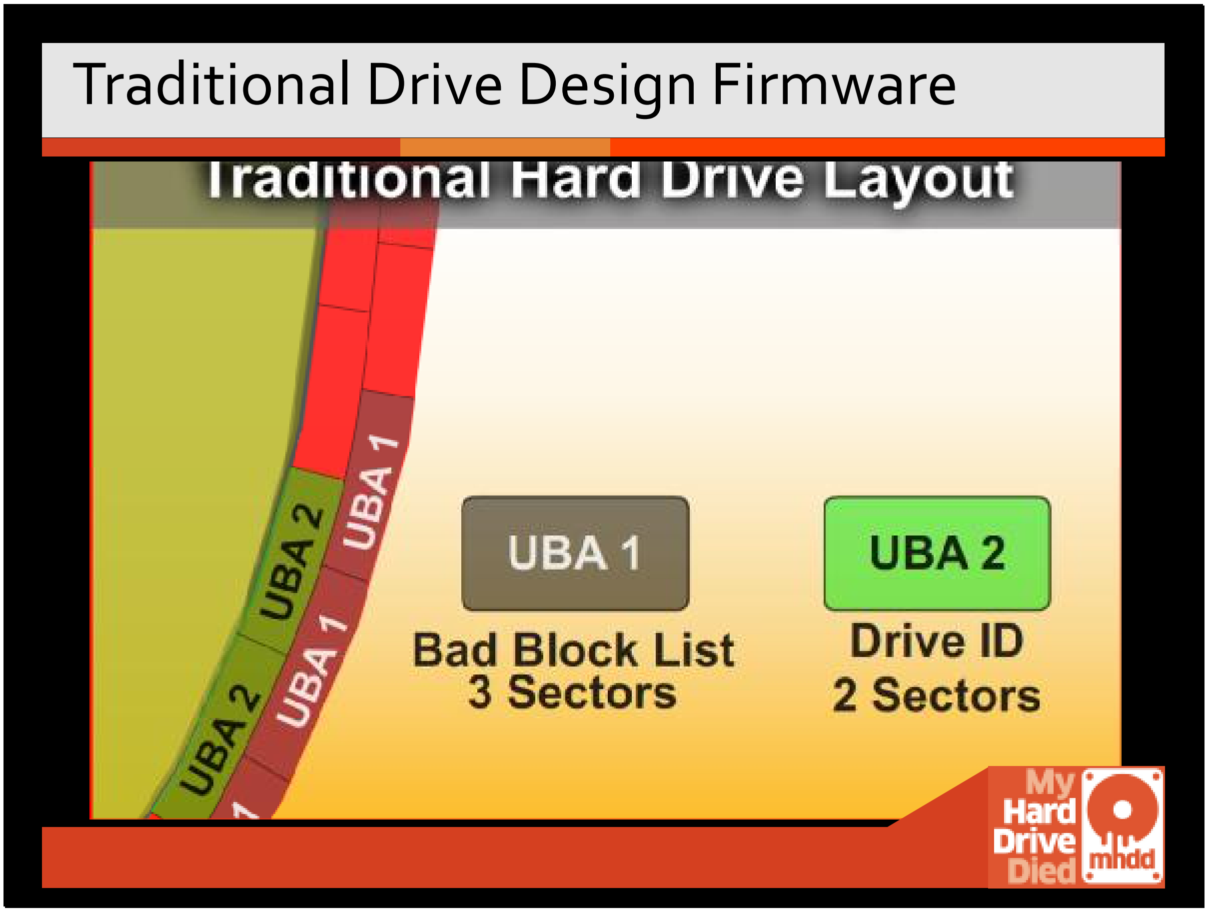

UBA blocks will probably be gone or replaced with something else for firmware or pointers to locate device space where a similar layout to the modules will be implemented.

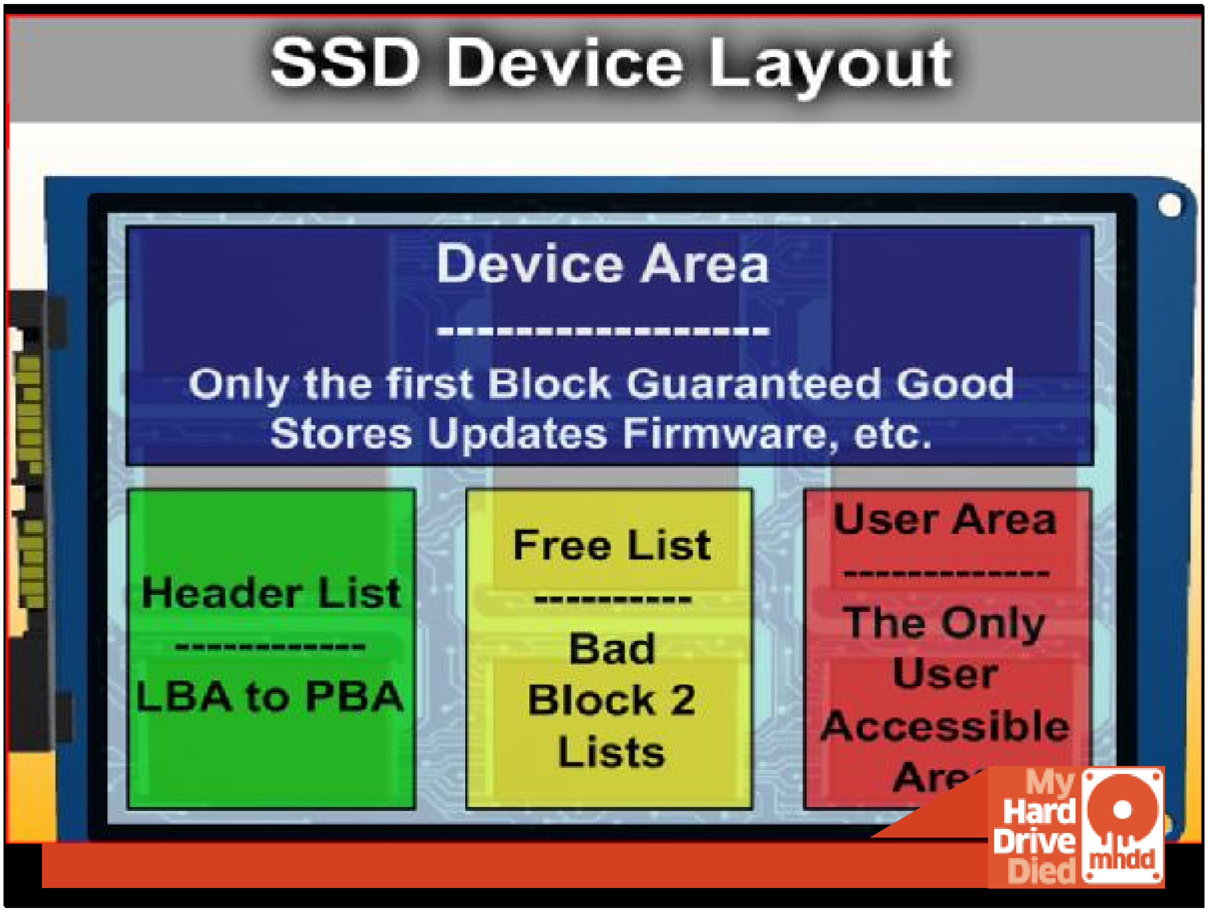

In solid state devices the blocks virtualize all the content over the device itself. The fundamentals of the basic generic table for most devices seem to be in this layout:

- 1st block is the only one guaranteed good and is where the device stores the important data
- This is then followed by the pointer for all the tables, which translates CHS to LBA to PBA
- The Free List block is for speed which already has a list of free locations
- The User Area block is the only area you have access to
- All content moves around with most wear-leveling except first block

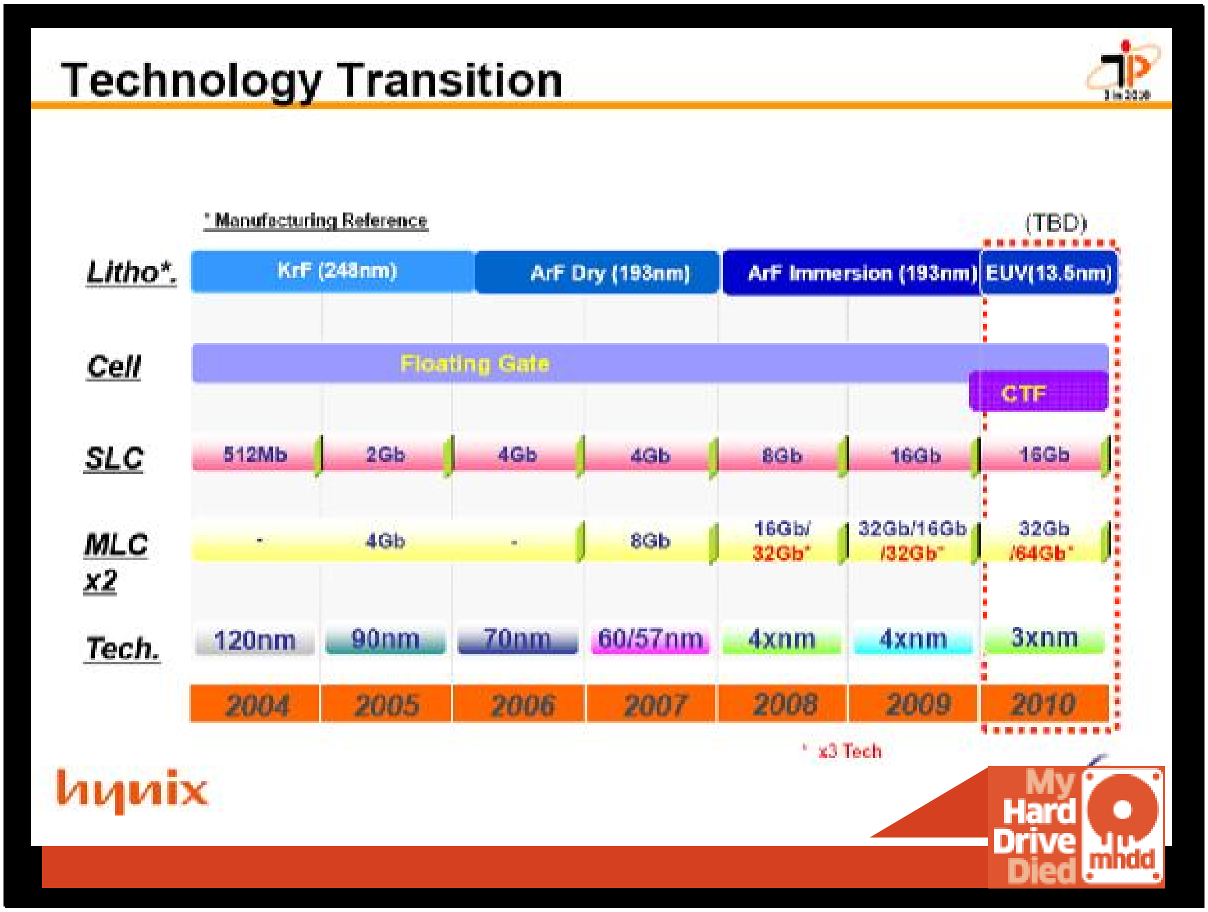

Presentation by Arun Kamat on Webcast 2008
“2012 is only 3 years away”
https://www.denali.com/ui/tv/Webcast?s=webcast2008_hynix2012
Used with Permission for Scott Moulton’s Class

MLCs are already at 32 Gigabit, SLCs are at 16 Gigabit.

Recovery of SSD drives compared to Hard Drives.

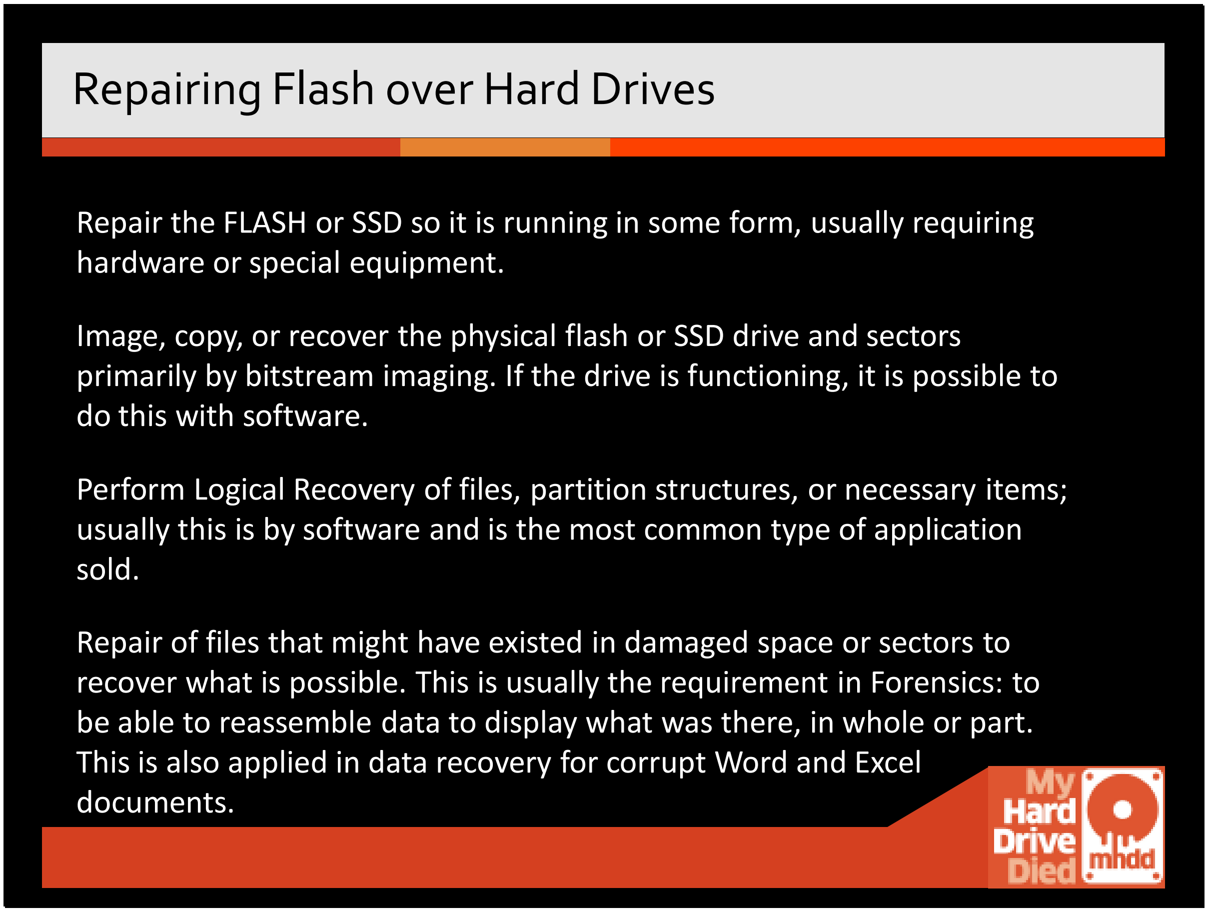

- Repair the FLASH or SSD so it is running in some form, usually requiring hardware or special equipment.
- Image, copy, or recover the physical flash or SSD and sectors primarily by bitstream imaging. If the drive is functioning, it is possible to do this with software.
- Perform Logical Recovery of files, partition structures, or necessary items; usually this is by software and is the most common type of application sold.
- Repair of files that might have existed in damaged space or sectors to recover what is possible. This is usually the requirement in Forensics: to be able to reassemble data to display what was there, in whole or part. This is also applied in data recovery for corrupt Word and Excel documents.

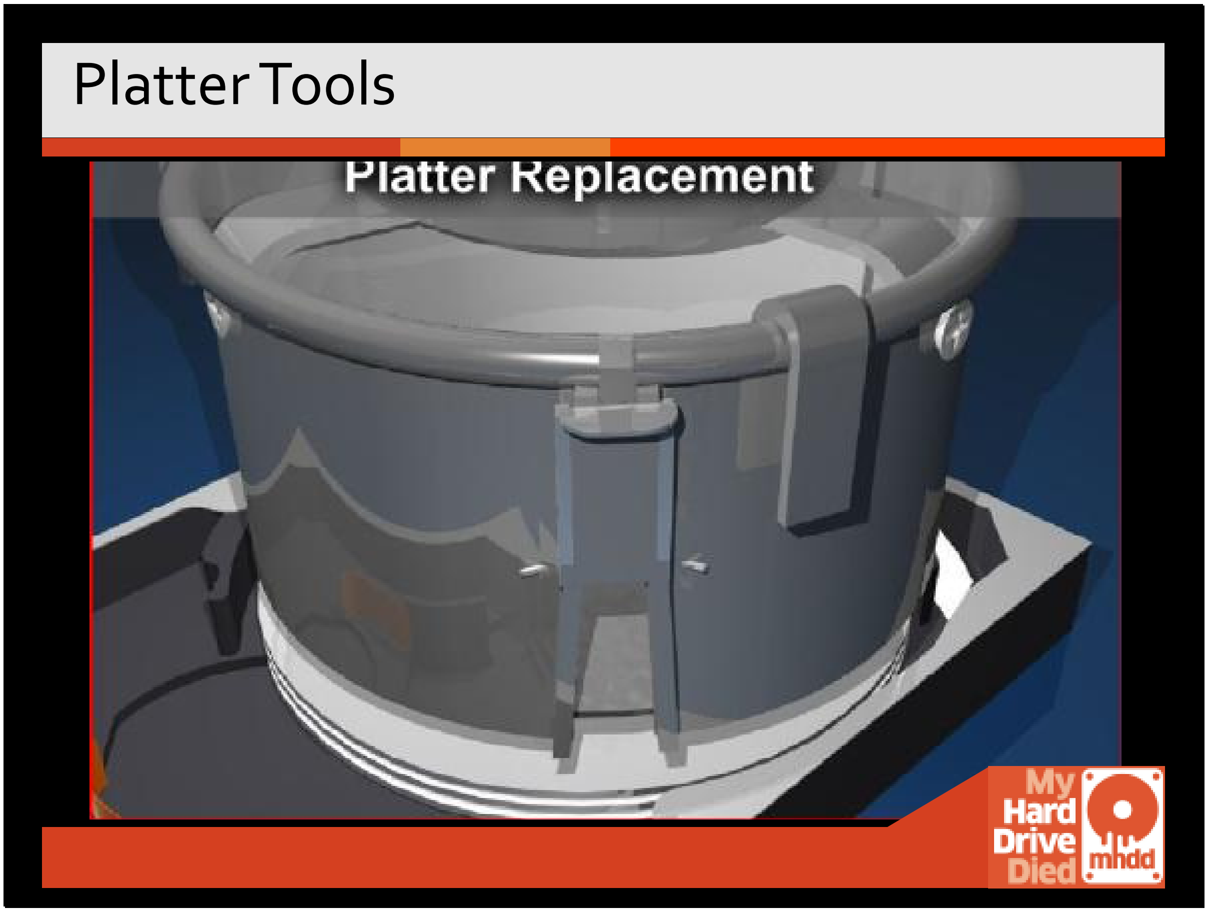

SSD's have a lot less to do. When working with SSD's, tasks like taking apart a hard drive and replacing platters go away.

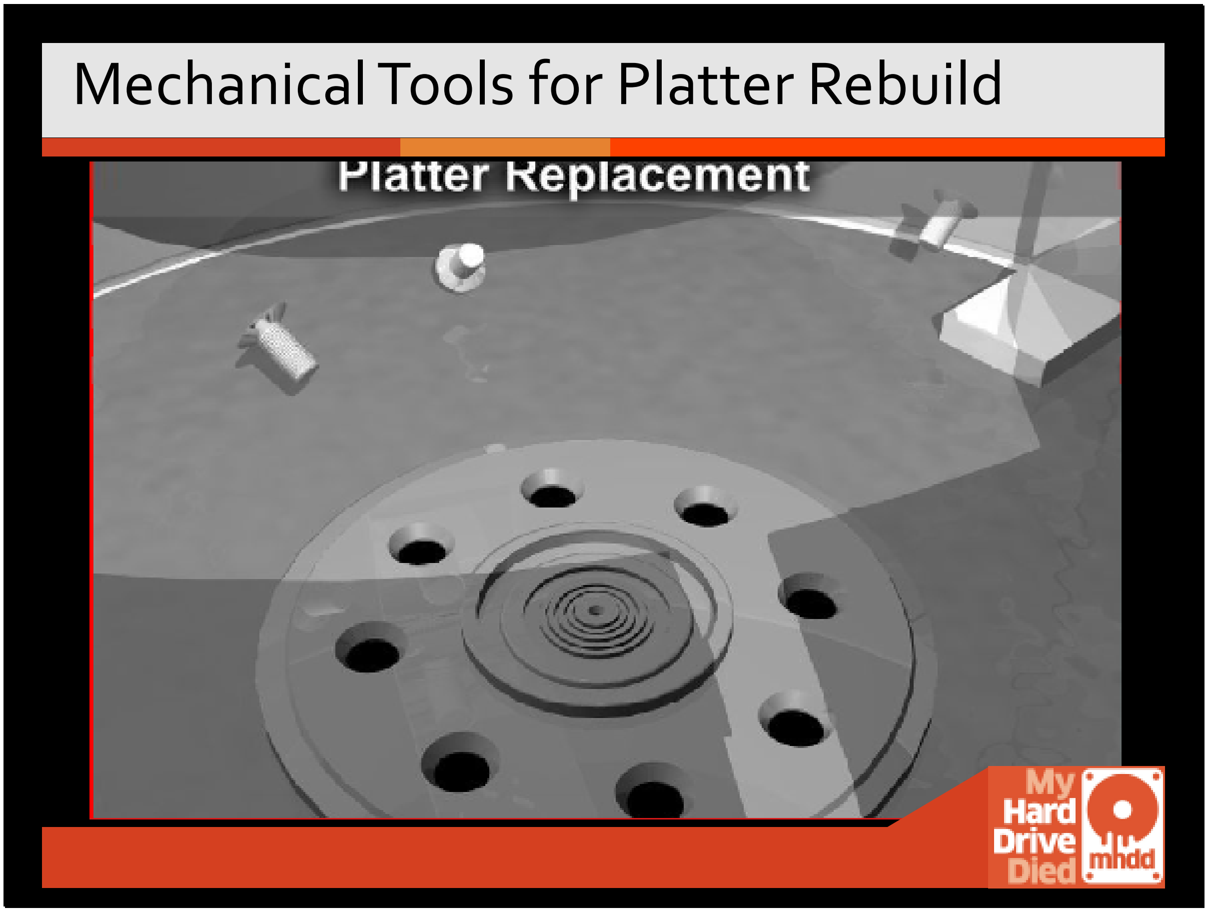

We are removing the mechanic and moving into a new realm with electronics.

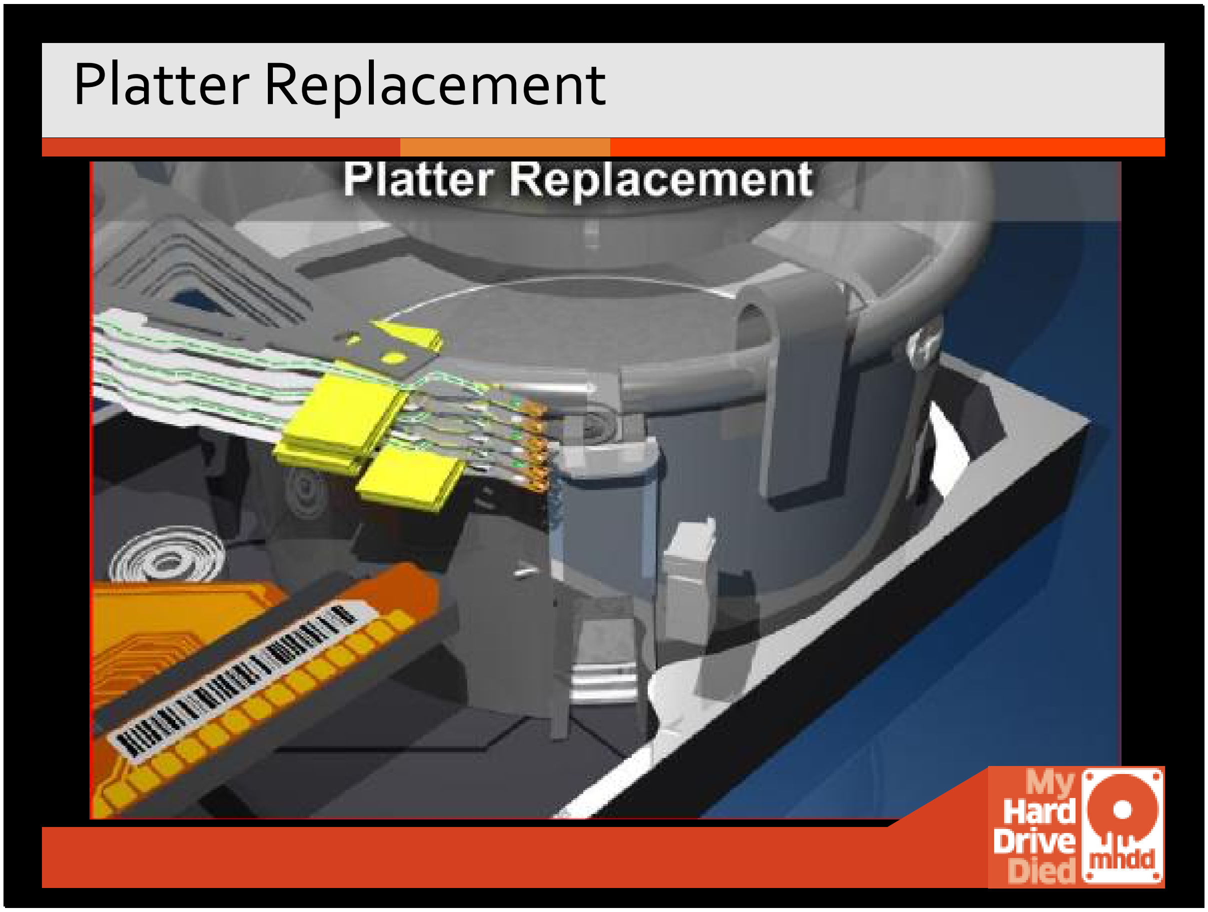

There will be no heads to realign or replace, and no simple donor drives to use as parts.

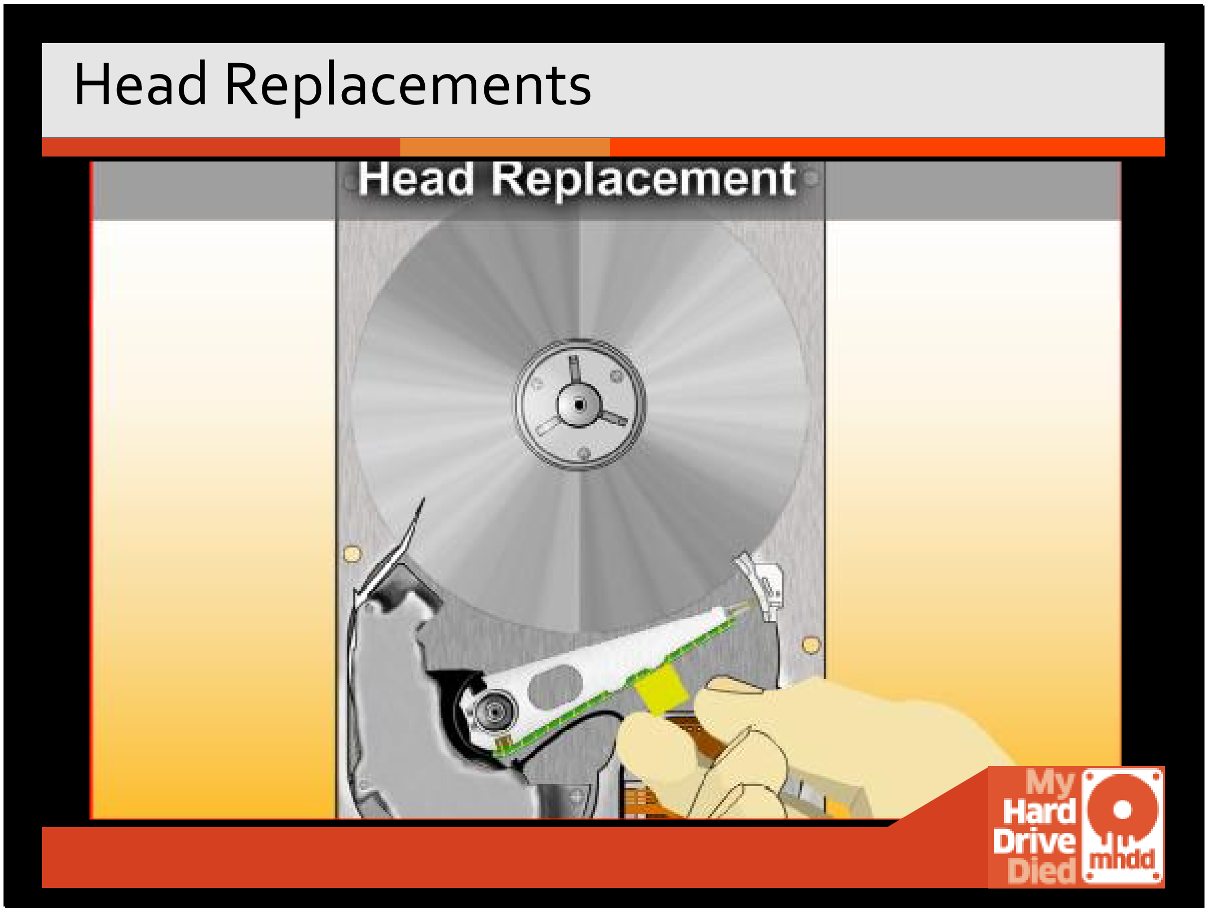

This will be the last remnant of the legacy items to go as we move from laptop drives to SSDs and eventually replace all spinning disks with other media.

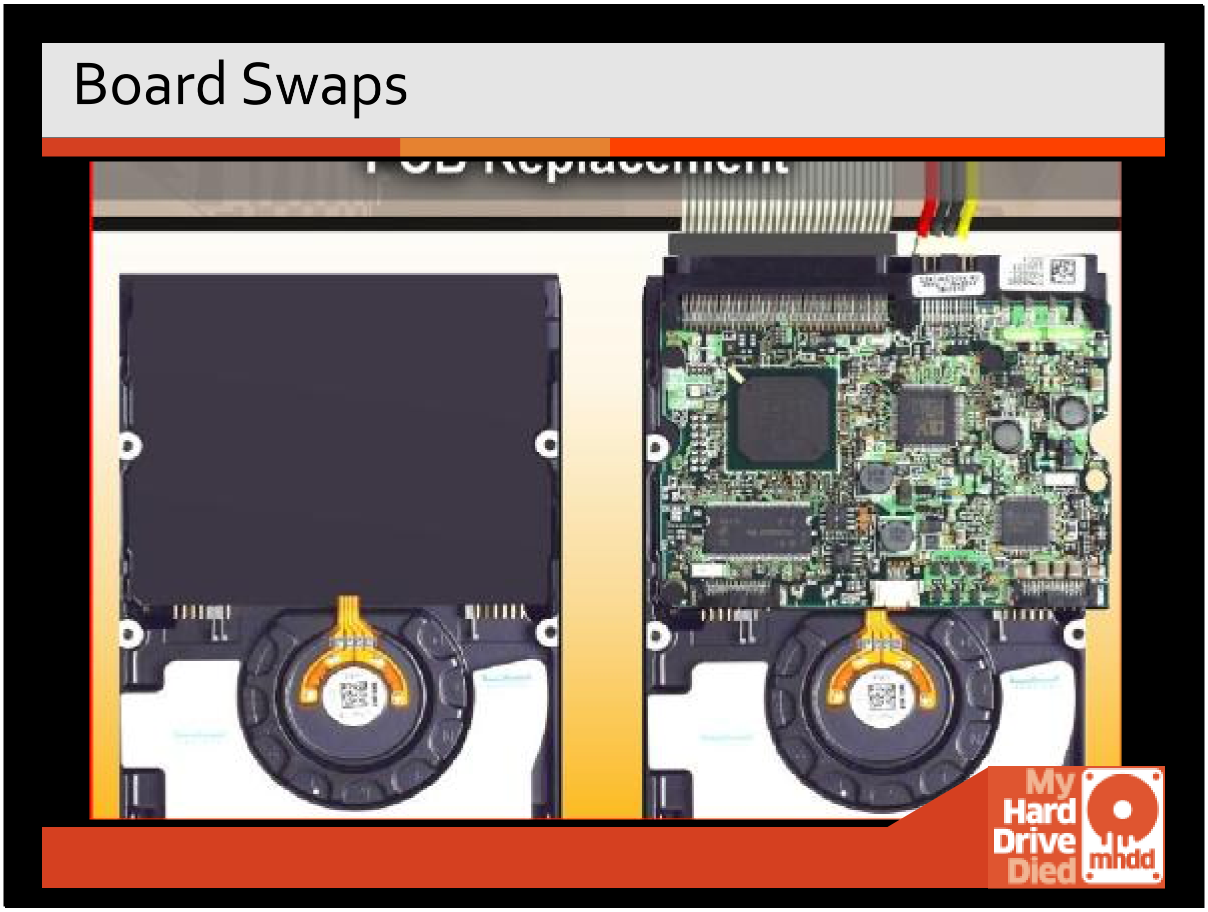

SSDs do not have printed circuit boards to swap around and cause a lot more difficulty than hard drives. There are no live board swaps.

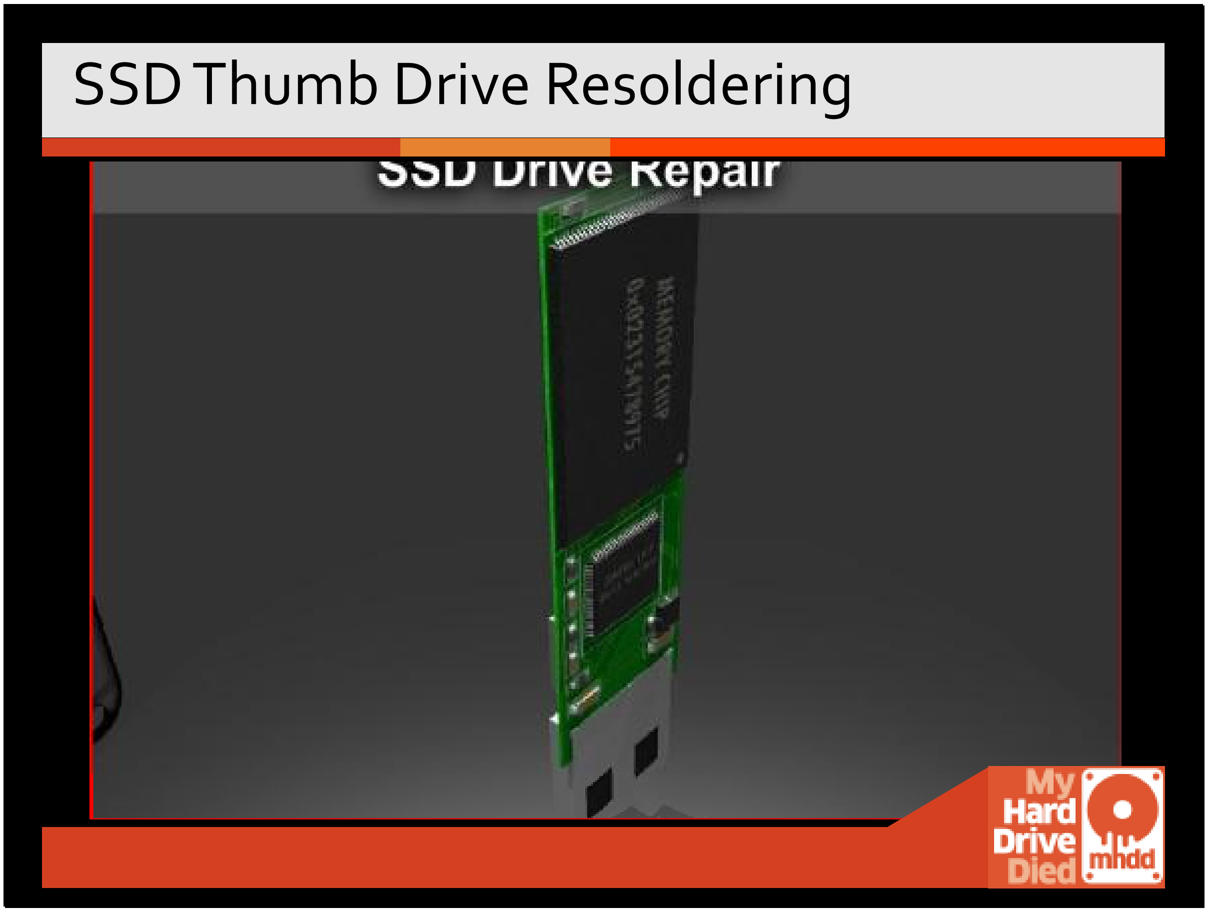

If you need to remove a chip and re-solder it, try ChipQuik. It is great for working on small jobs.

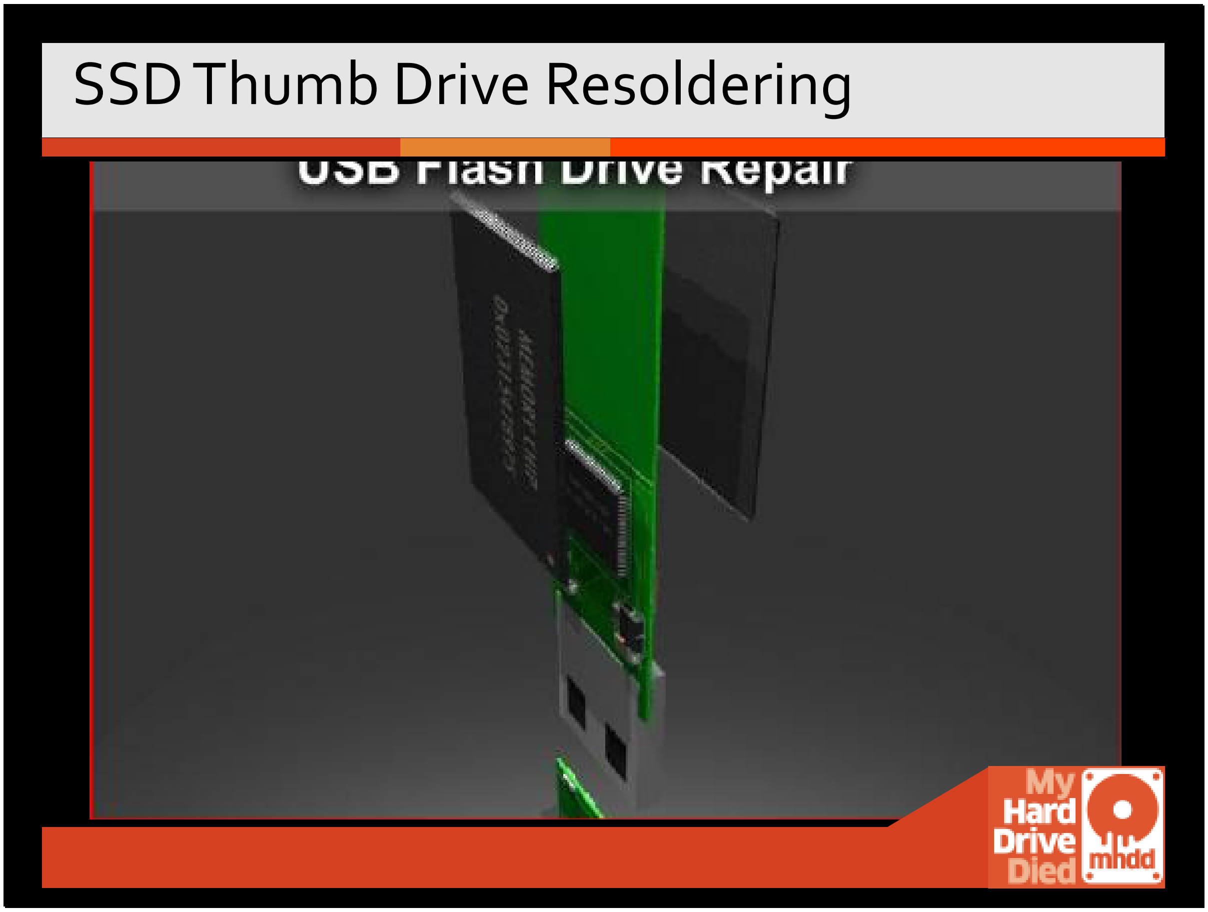

ChipQuik stays molten longer, allowing you to remove the chips without thousands of dollars in equipment.

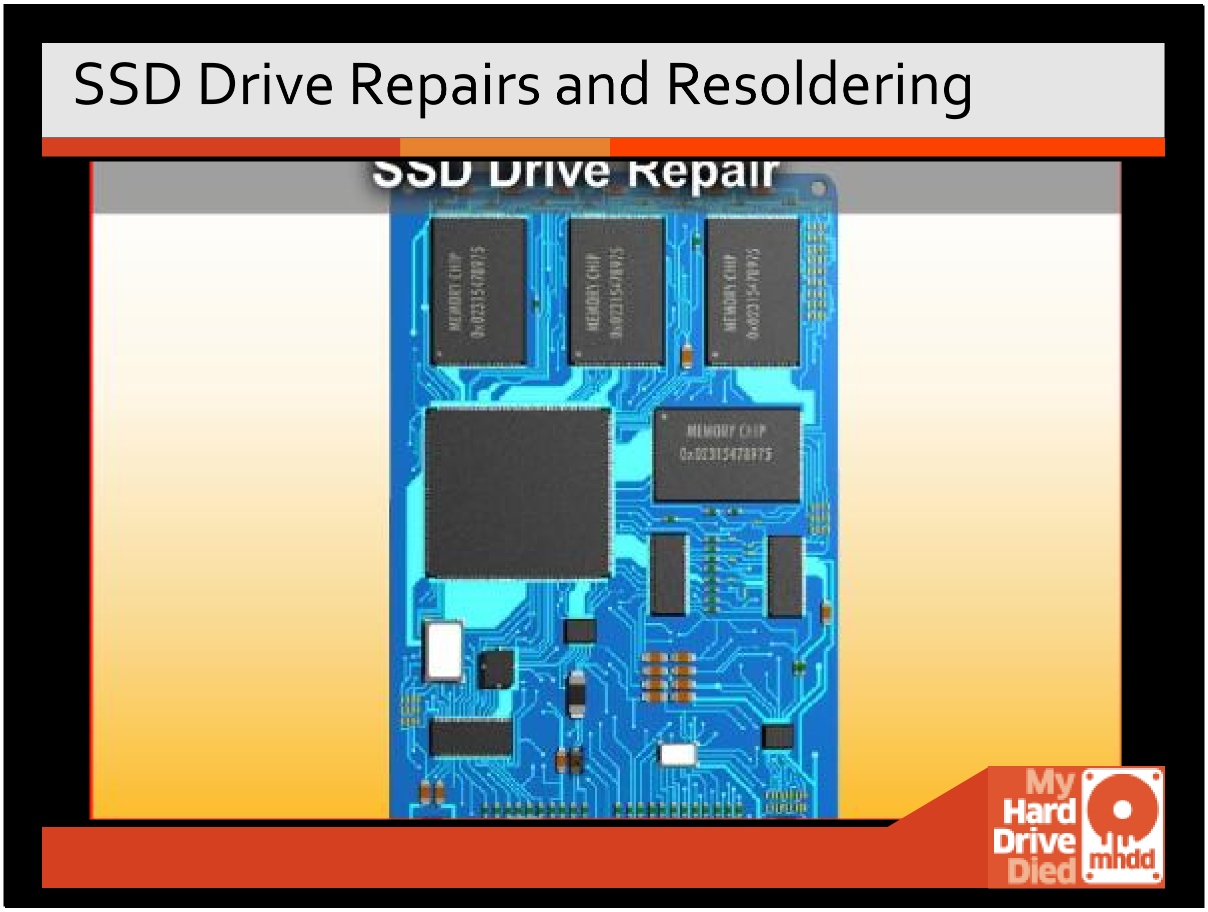

SSD drives are a lot more complex and costly to disassemble and reassemble. The hard part is that you have to carefully de-solder and re-solder all the chips. The expensive part is that you have to destroy another drive just like it to get the board or parts.

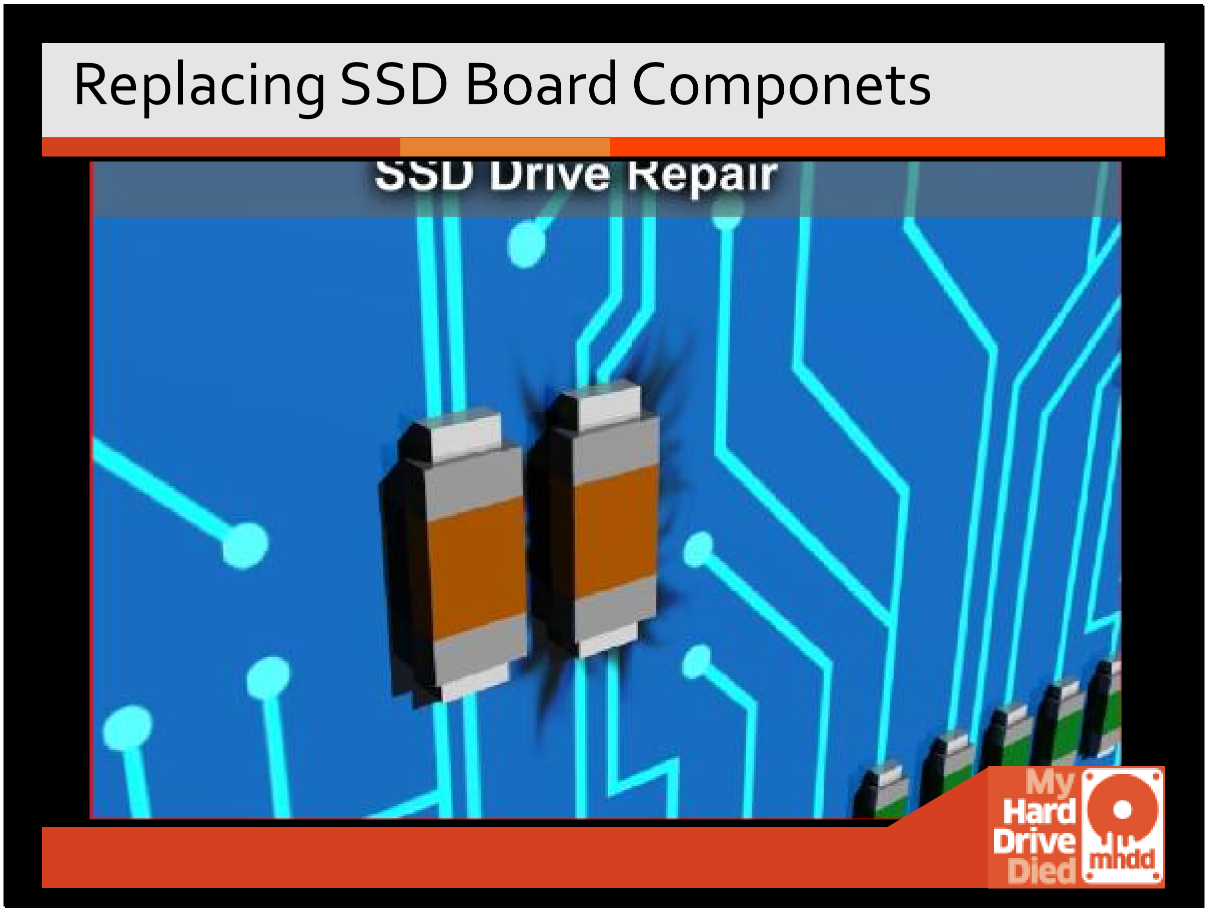

Some items are the same as Printed Circuit Board (PCB) repair. You still have the option of looking at something that may have gotten fried and replacing it, which is also the same when dealing with memory sticks. The nightmare comes when you have a cracked board.

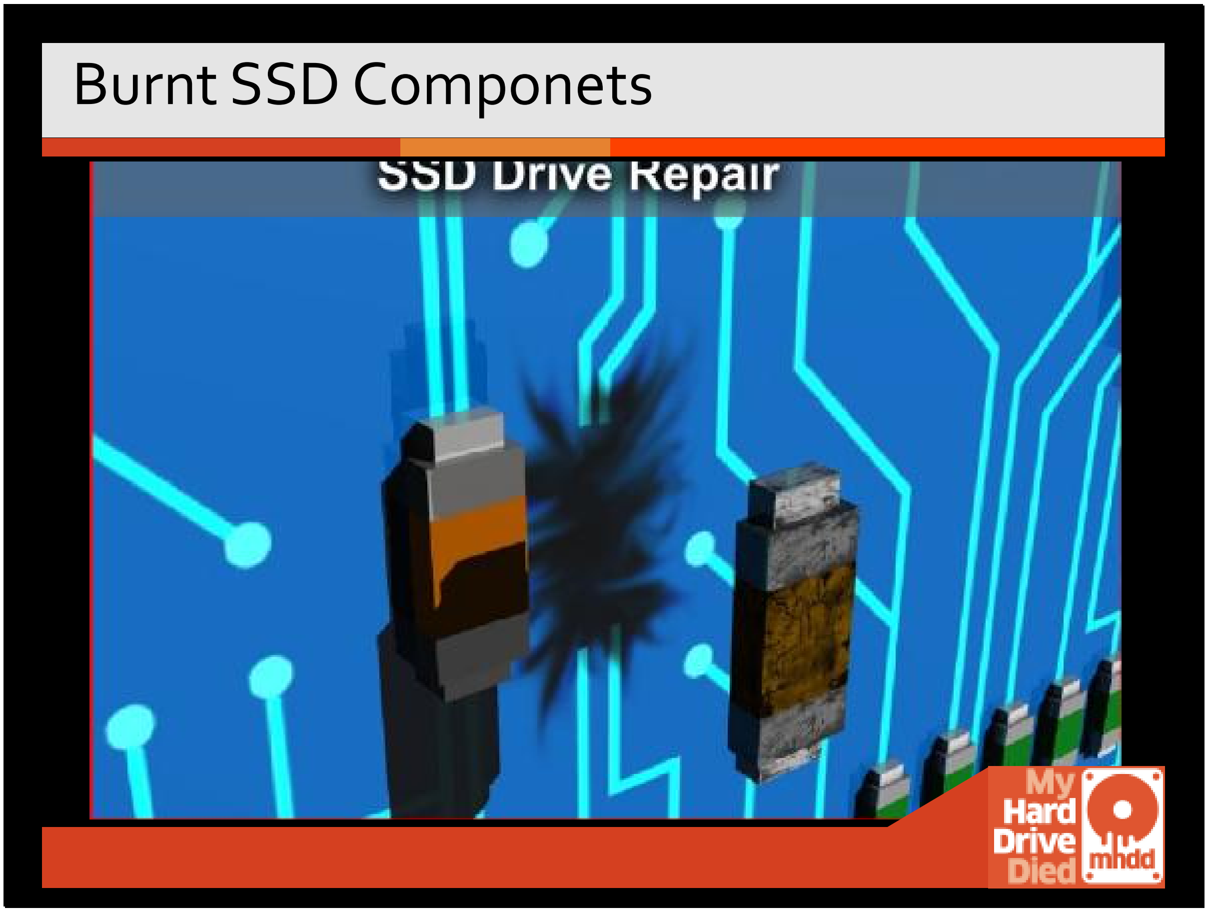

You may still be soldering to repair damaged components.

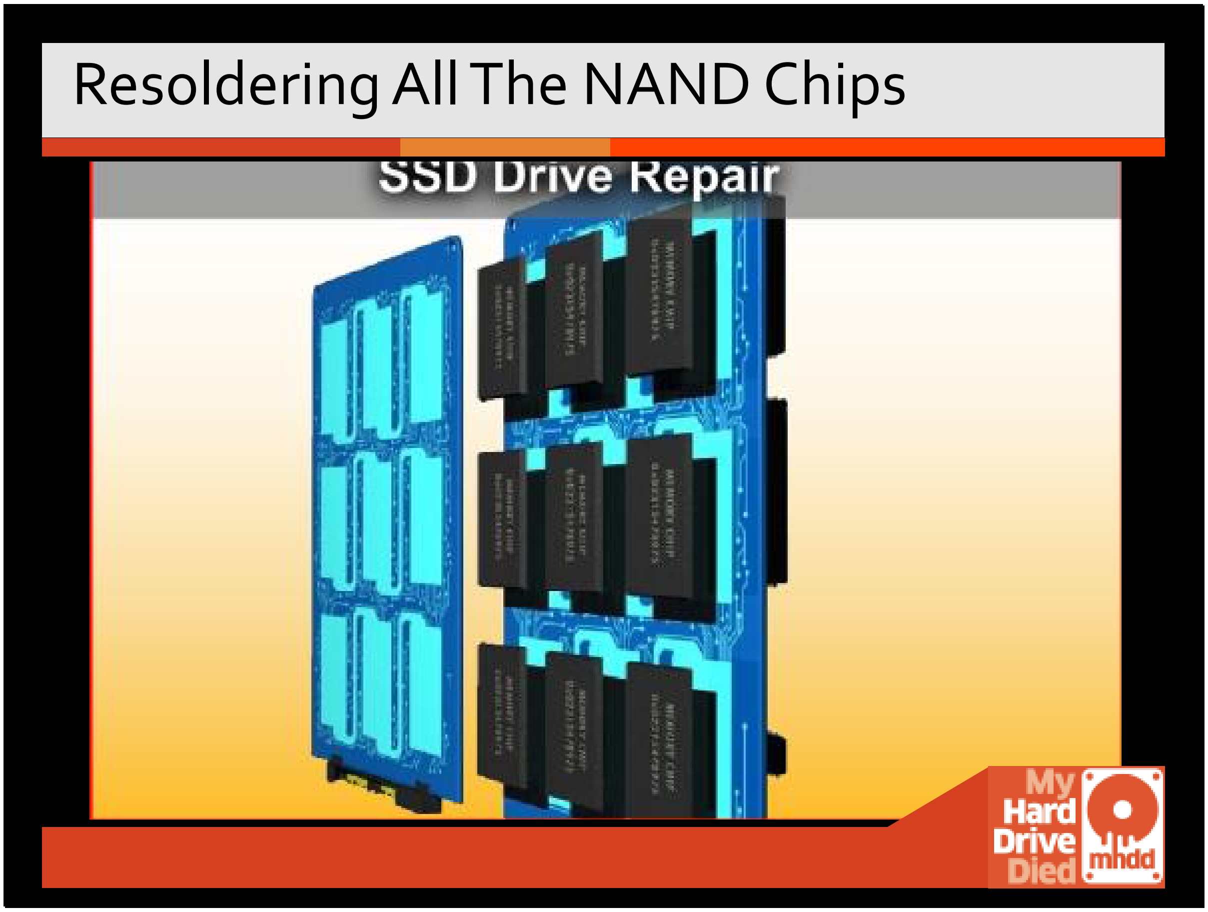

In some cases you will be rebuilding the entire board. This is obviously not very easy.

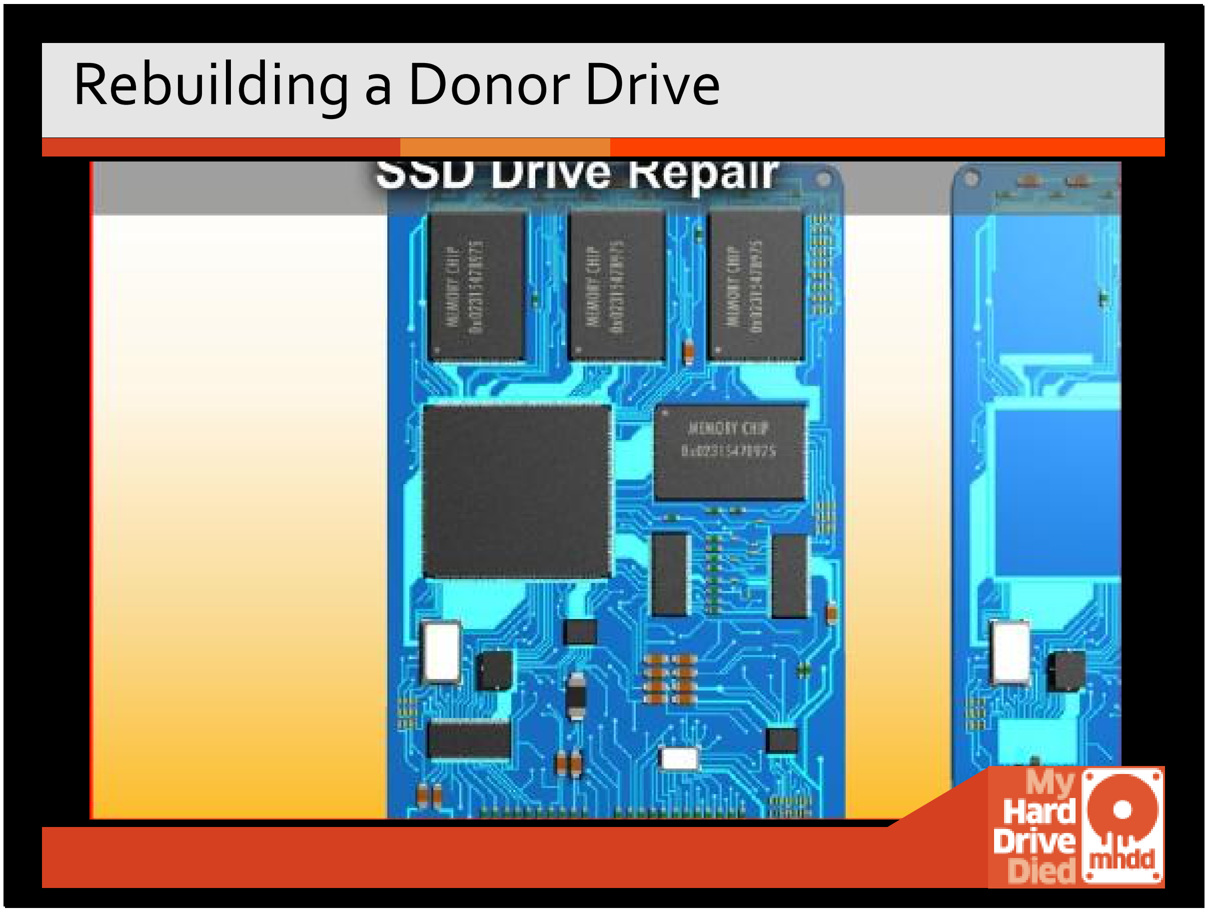

Expect that during your repair process of the SSD you will spend at least the price of a donor drive and destroy it attempting to repair the bad drive you want to recover from. There is always a great chance you will accidently destroy the bad drive in the process. It is hard to guarantee it will work, and you have to become great at soldering: 8 – 10 hours of work easily.

There are new techniques being developed every day, as this is new technology and we are only beginning to see advances in the area of dealing with Flash and its algorithms.

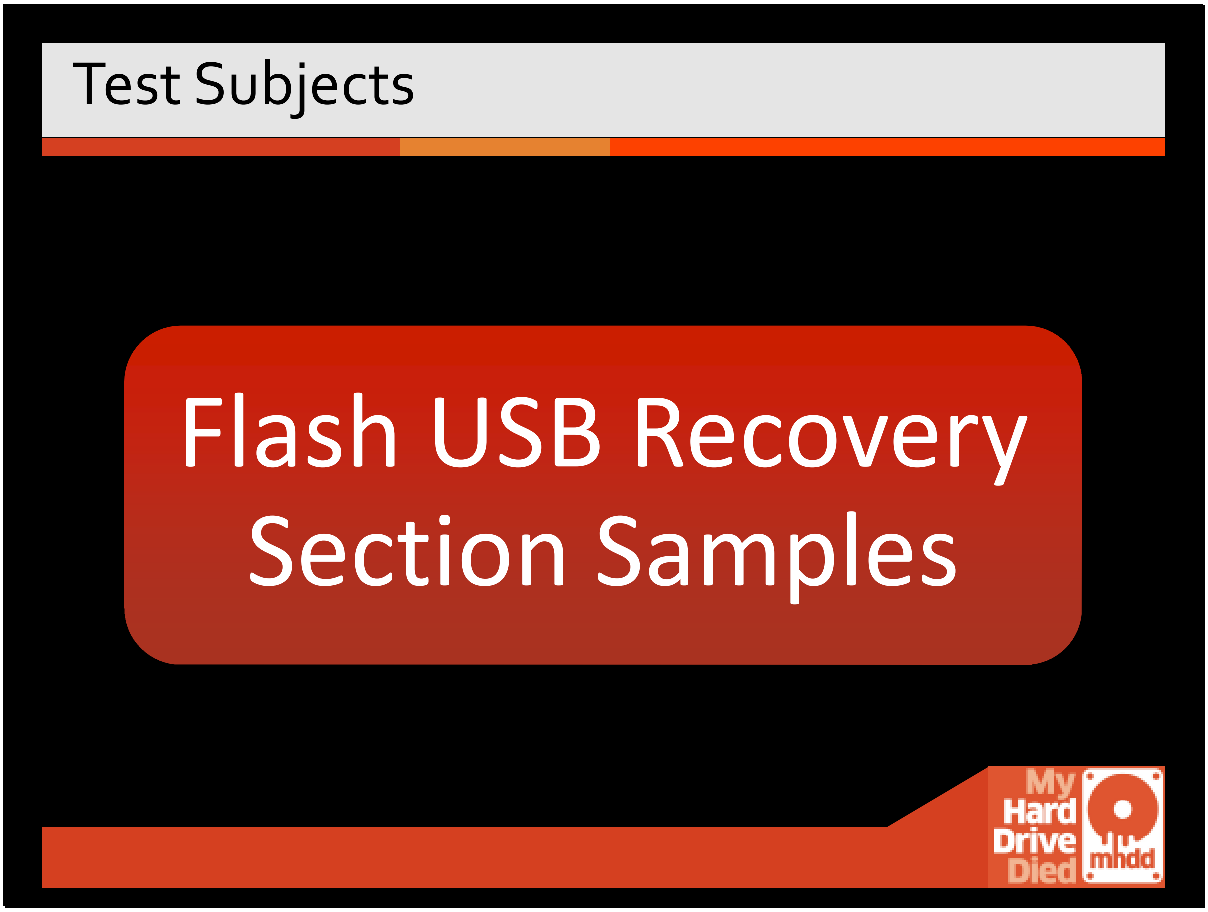

Samples of Flash USB Sticks we have damaged, repaired, and imaged.

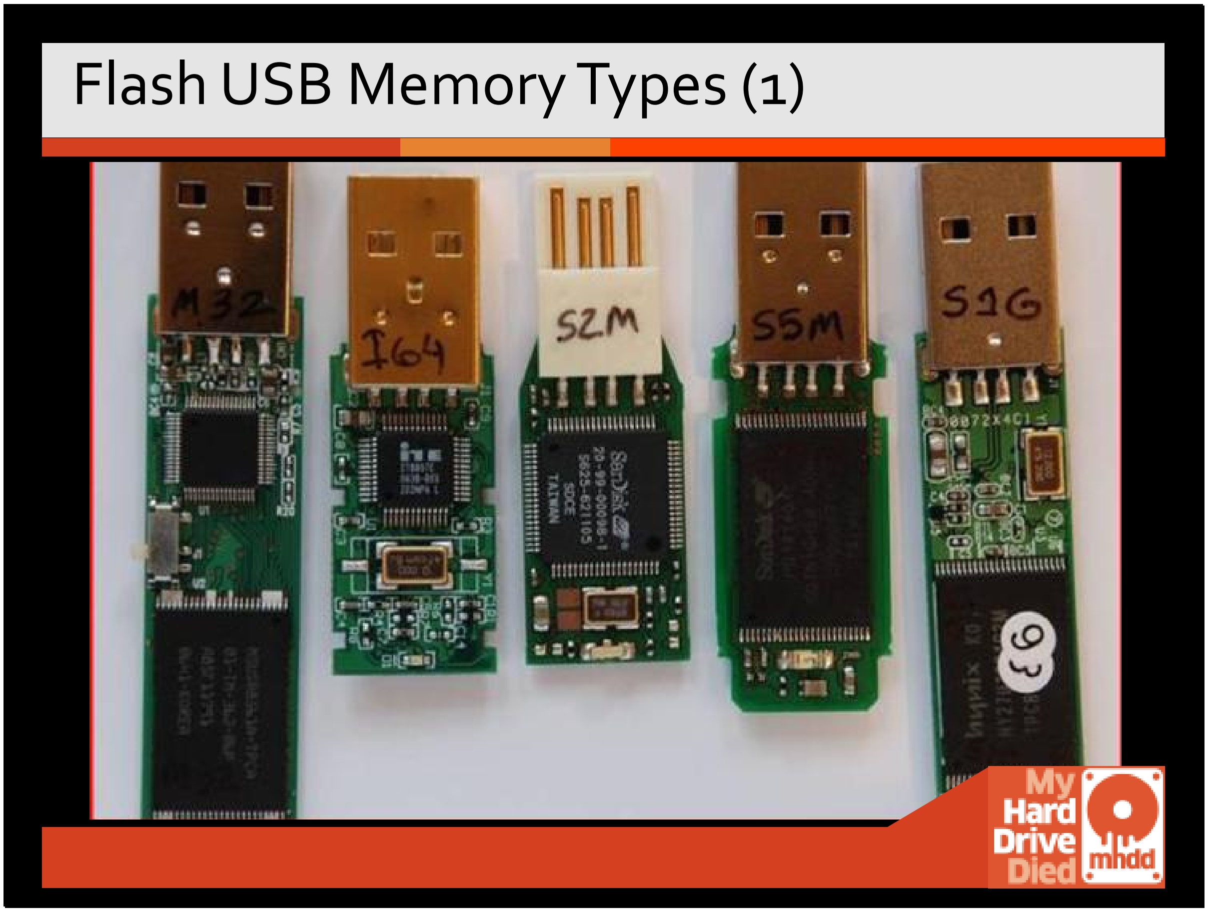

Samples of Flash USB Sticks we have damaged, repaired, and imaged.

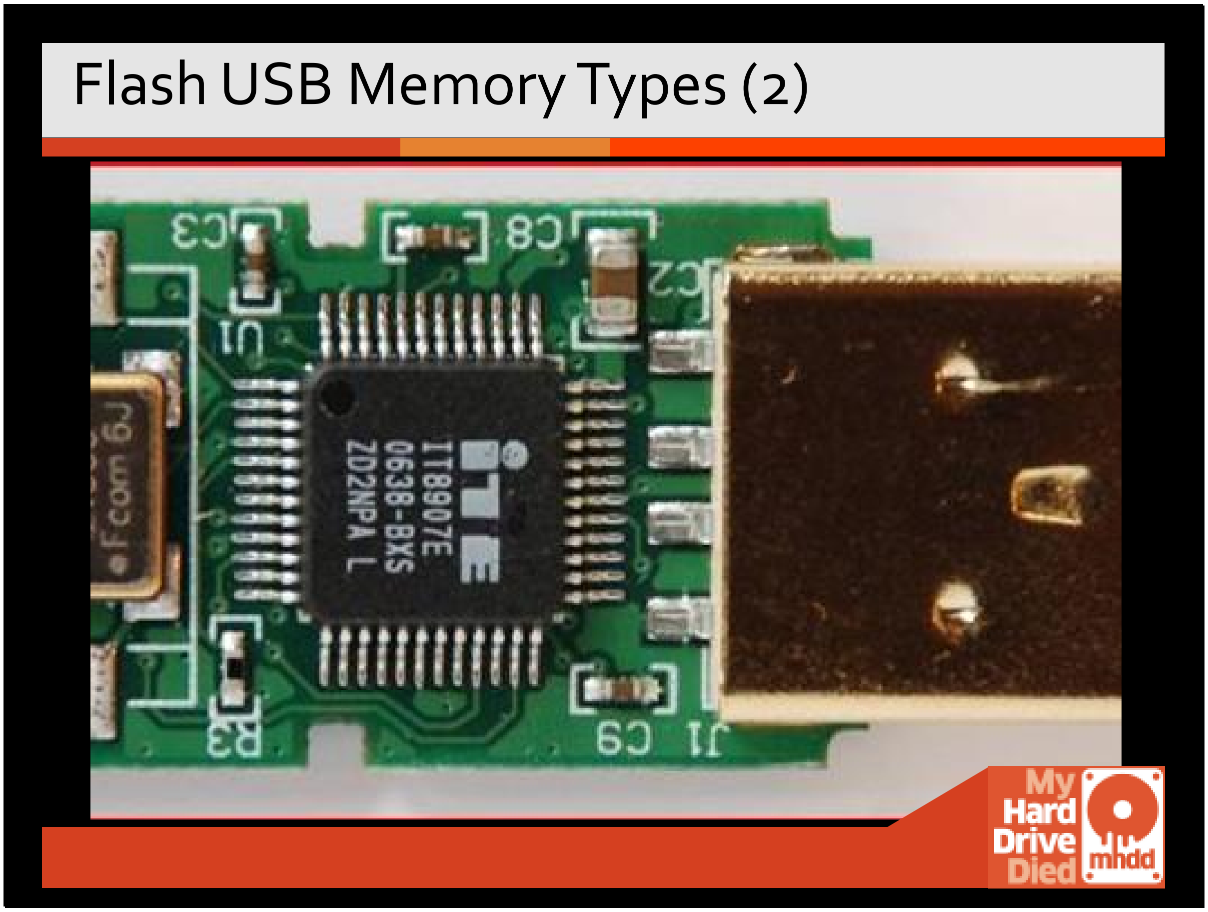

This is an example of the control chip for the USB device.

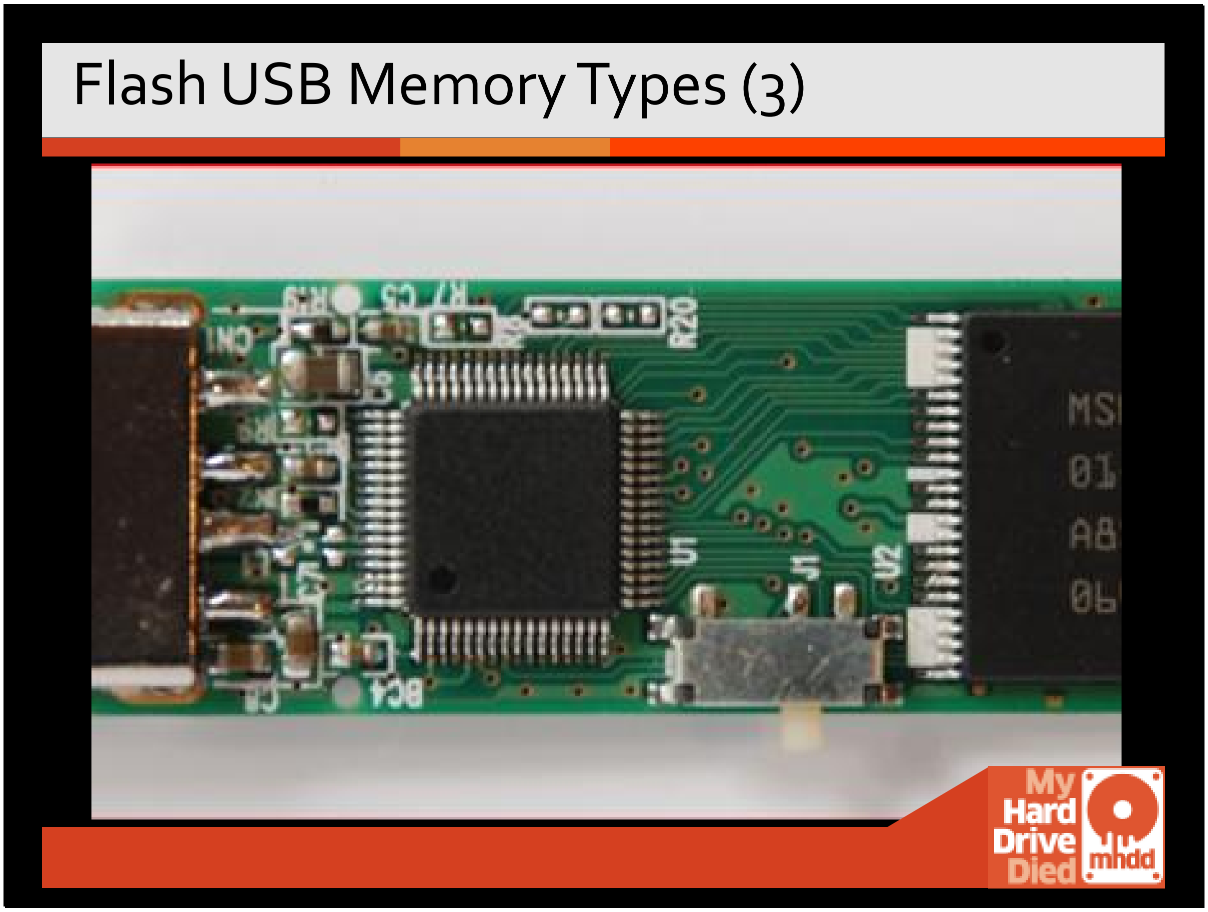

This is the chip that you would have to remove and image in order to acquire the content. There are new developments in this area every day. Different devices have different algorithms for the encoding and wear-leveling that makes it difficult to just read the content from the chips. So far I have heard of about a dozen different algorithms. It is plausible that over time a device will be developed to incorporate these different sequences, but as of a year ago not much research had been published on this topic. As I write this there are currently two new developments in process just released; A PC3000 Flash and Salvation Data Flash Device.

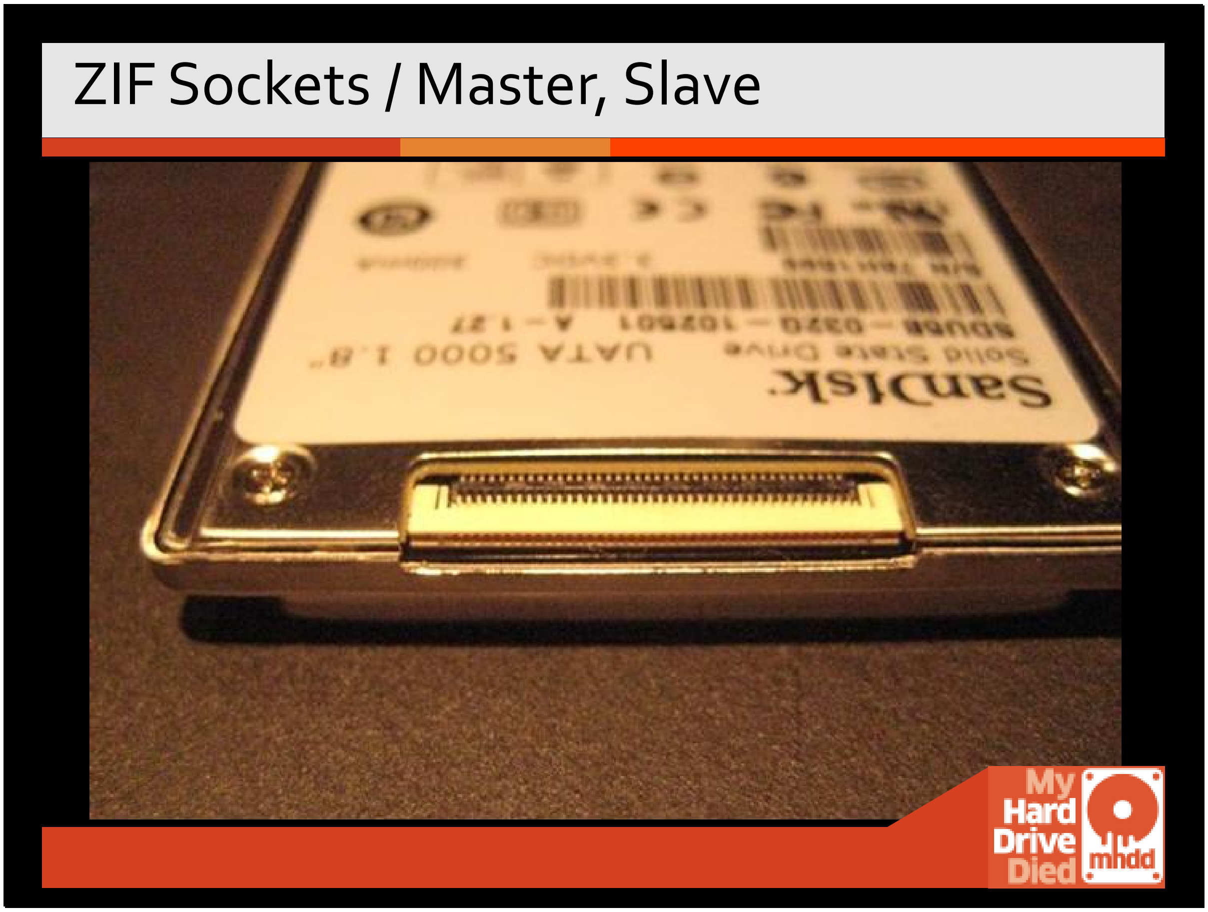

Code:N
This is a photo of a ZIF socket. They are still a parallel interface (PATA) however they generally have a permanent master/slave setting.

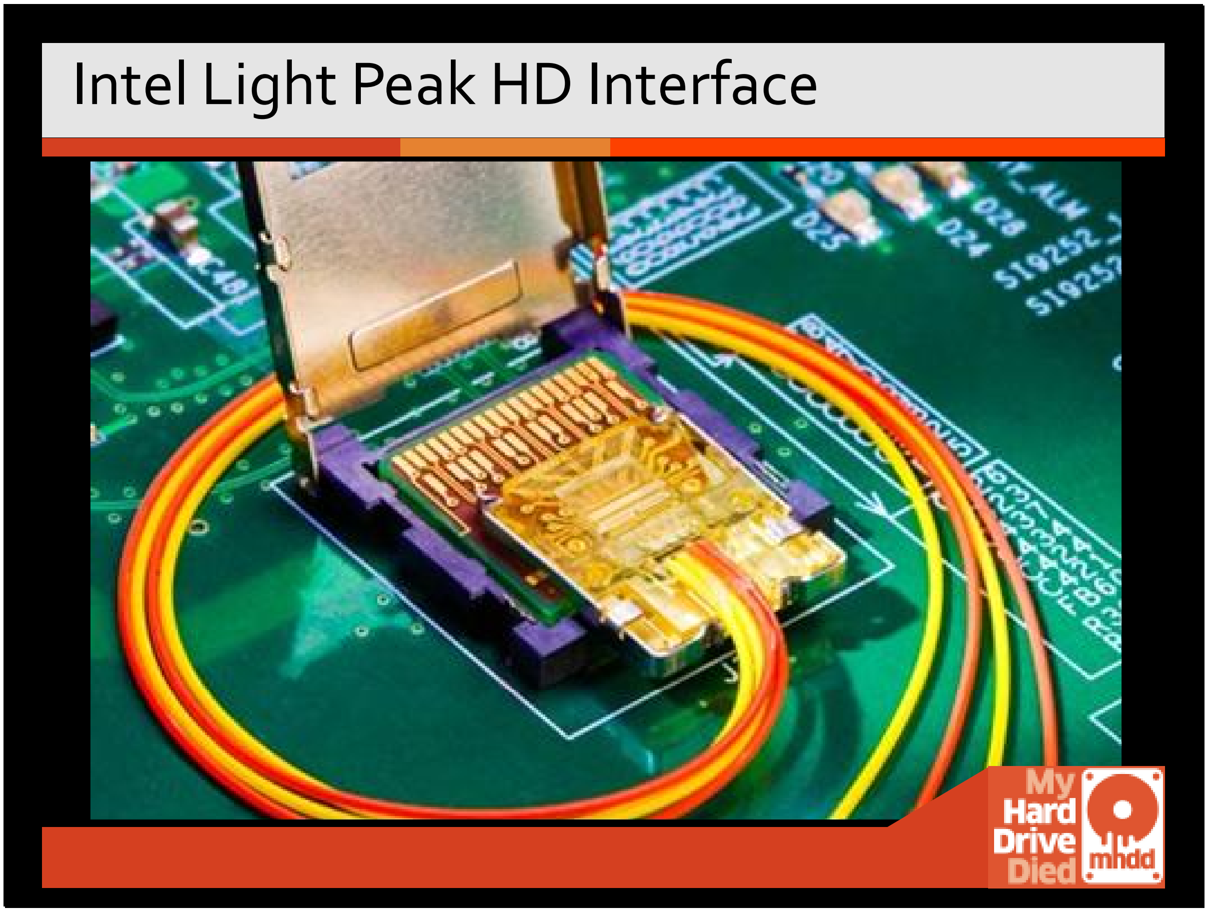

Code:N
Intel has designed a new hard drive interface they expect will be released in 2010/2011 as a primary controller interface. They have found a way to make an Optical Transceiver cheap enough to put on the motherboard. While you will still need a power source for the drive, this should make communications much faster with the drive and improve IO speed. Bandwidth will start at 10 Gb/s.

REF: http://techresearch.intel.com/articles/None/1813.htm

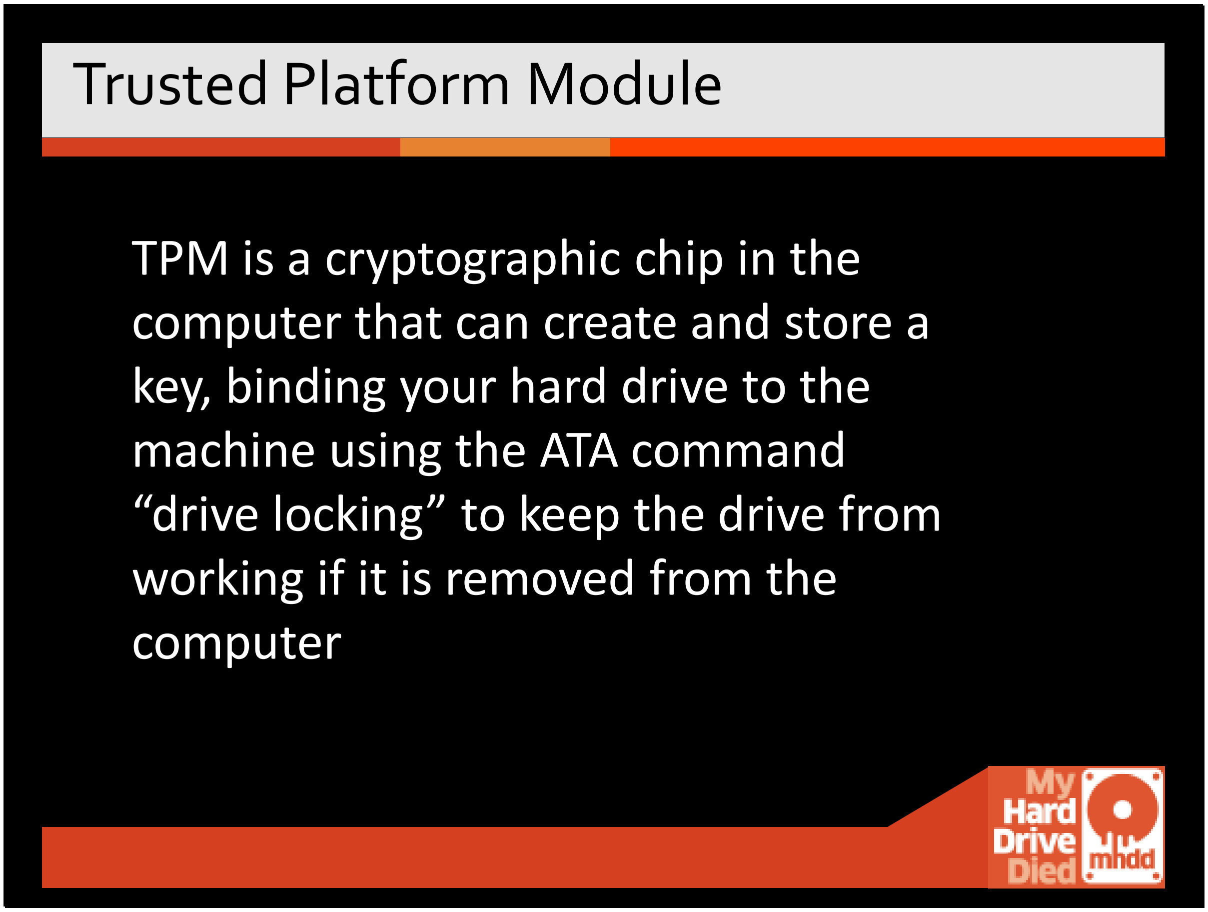

Encryption exists on some SSD and USB Thumb Drives. TPM is a cryptographic chip in the computer that can create and store a key, binding your hard drive to the machine using the ATA command "drive locking" to keep the drive from working if it is removed from the computer.

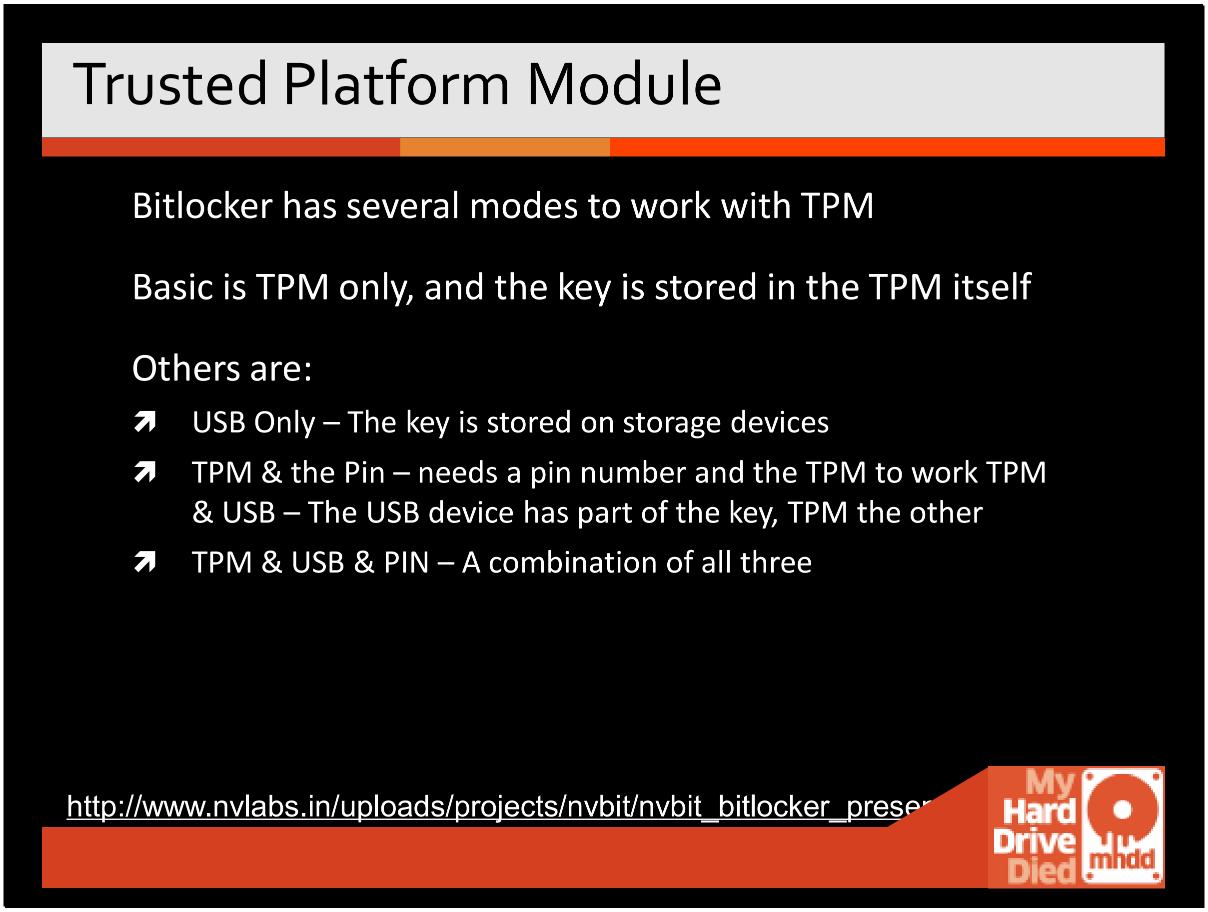

Code:R
These guys are doing a lot of work on the TPM and Bitlocker side.

Get more info here:
http://www.nvlabs.in/uploads/projects/nvbit/nvbit_bitlocker_presentation.pdf
http://www.nvlabs.in/uploads/projects/vbootkit/nitin_vipin_vista_vbootkit_poc_RC1_edited_video.avi

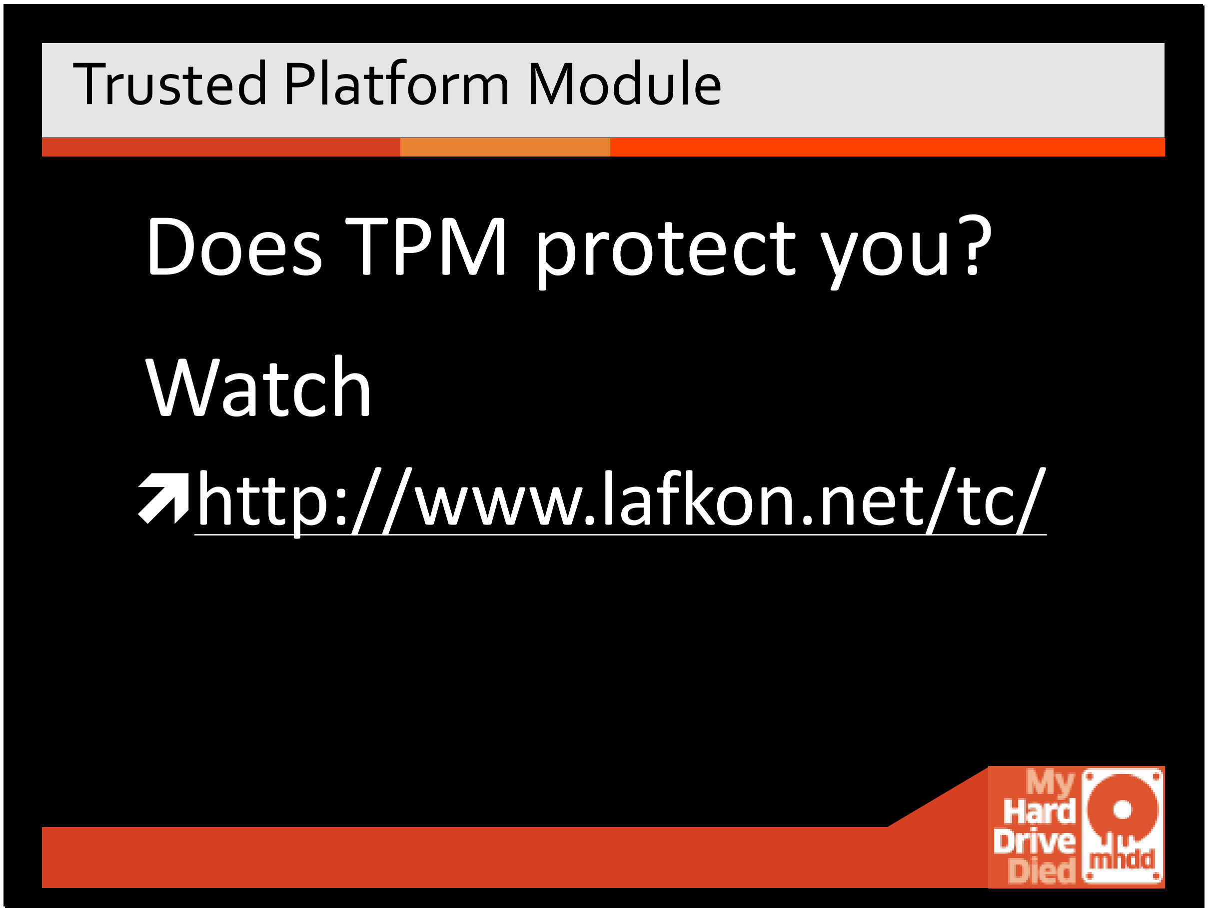

Protection by TPM. TPM can lock the drive causing you problems accessing it, but in addition the data is encrypted. It is possible to get past the password issue, but then you might not be able to read the data because of the encryption.

http://www.bit-tech.net/news/2008/06/06/gigabyte-tpm-explained/1
Think Government Key Escrow. http://www.cdt.org/crypto/risks98/

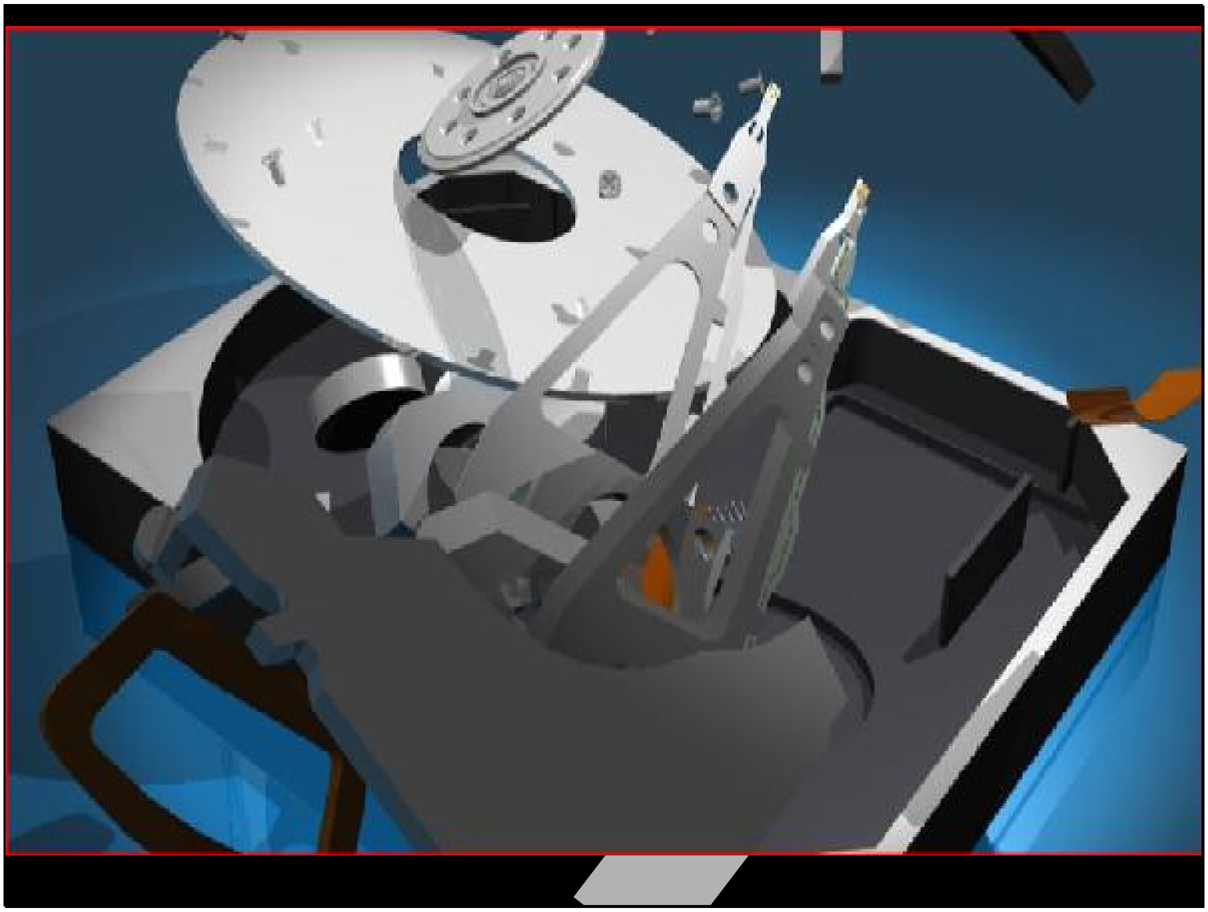

This www.MyHardDriveDied.com & Forensic Strategy Services, LLC. @ Scott A. Moulton Section intentionally left blank.

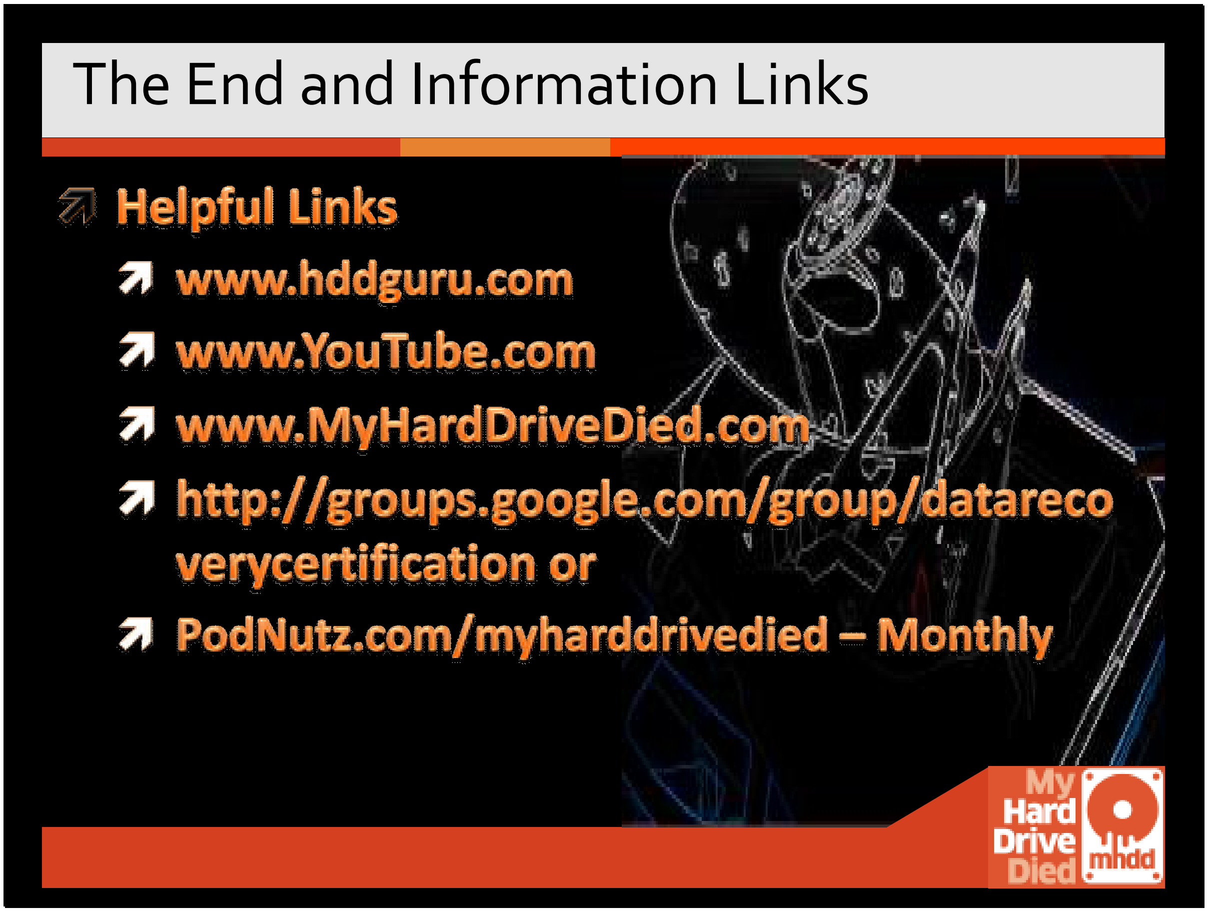

There are other places where you can get some helpful information. Here are some of the websites you can use for reference.

Helpful Links for Research on Data Recovery

- www.hddguru.com
- www.YouTube.com
- www.MyHardDriveDied.com
- http://groups.google.com/group/datarecoverycertification or
- http://bit.ly/3A5qex

ComputerAmerica.com – Two Hours Monthly
PodNutz.com/myharddrivedied – Monthly