System Impacts of Storage Trends

Hard Errors and Testability

STEVEN R. HETZLER

Steven R. Hetzler is an IBM Fellow at IBM's Almaden Research Center, where he manages the Storage Architecture Research group. He has spent 25 years in data storage research and development. Steven is currently focusing on novel architectures for storage systems and on applications for non-volatile storage. His most notable patents include split-data field recording and the No-ID headerless sector format, which have been used by nearly all magnetic hard-disk-drive manufacturers for a number of years. Steven was educated at the California Institute of Technology, where he received his PhD and Master's degrees in Applied Physics in 1986 and 1982, respectively. He joined IBM Research in November 1985 and was named an IBM Fellow in 1998.
Hetzler@almaden.ibm.com

The continuing increase in disk storage capacity has been extremely beneficial to the IT industry. However, the growth in capacity has not been accompanied by increases in the hard error specifications. We will examine how this affects system design. SSDs aren't immune to the effects of hard errors either. In fact, in many ways, designing systems to properly accommodate SSD hard errors is more complex than for hard disk-based systems. In this paper we will explore hard errors, develop some improved methods for specifying them, examine testing requirements, and propose some new specifications for SSDs.

Background

Hard disk (HDD) capacity has been increasing exponentially since they were first introduced. In the past few years, however, the demand for enterprise storage appears to have exceeded the ability of enterprise disks to deliver storage that meets customer budgets. This can be seen in Figure 1: the proportion of capacity-optimized disks in enterprise applications is increasing. Commonly, the performance-optimized disk will be what the industry calls enterprise class (10,000 and 15,000 rpm) disk. Capacity-optimized disks will be 7,200 rpm SATA-class. Note that such disks may not use the SATA interface, but are distinguished by their slower rotation speed and higher capacity compared with enterprise class disks. SATA-class disks come predominantly in two grades: enterprise and consumer.

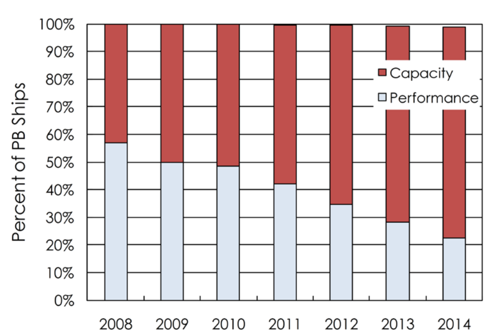

Figure 1: Percentage of PB shipped in enterprise storage systems by HDD class [1]

Unfortunately, there are side effects to the capacity increase in hard disks. The problem arises because not everything scales with capacity. This requires that system designs compensate for these effects. Let's start by examining the impacts that SATA-class disks will have on storage systems.

Hard Errors

Ideally, a storage device should always be able to return the data stored on it. When a storage device is unable to deliver a piece of data, it is referred to as a hard error. In such a case, the device has exhausted all means to recover the data. Thus, while the device is still operational, data has been lost.

Hard error specifications are one of the attributes that haven't scaled with capacity growth. However, the situation is obscured by the archaic units used to specify hard errors. Typically, storage devices are specified as less than one event per some number of bits transferred: e.g., <1 error in 10^{14} bits read [2]. We will refer to such a specification as a hard error interval, or HEI. While this may have been appropriate when disk interfaces were more primitive, it's not very informative today.

First, such a specification is not even statistically meaningful. The specification represents a minimum, which we all know isn't obtainable. The distribution is unspecified as well, so we can't tell if the mean is one error in twice the specified interval, or in 1000 times the interval. Thus, we can't model with the specification as is. The only reasonable approach is to operate as if it's a mean, and go from there. If the manufacturers are bothered by treating the specification as a mean, they could publish the distributions.

Second, a hard error results in an entire sector being lost, not one bit. Further, a block device can only transfer data in increments of a sector. Thus a per-bit specification is not as useful as a sector failure probability. Assuming a typical 512 byte sector size, we would have:

```
Sector HE probability = 1/(HEI * 4096)
```

Third, the large exponents give an artificial impression that the devices are quite reliable. For example, 10^{14} bits seems incredibly large. However, there are 0.08×10^{14} bits in a terabyte! Thus, the reliability isn't quite what it seems.

A Better Hard Error Spec

I feel a specification should be sufficiently clear to allow one to determine its impact at a glance. Therefore, I propose that hard errors be specified as probability per TB transferred.

	SATA	Ent. SATA	Ent. Hi Perf.
Typical hard error interval (bits)[2], [3], [4]	10^{14}	10^{15}	10^{16}
Sector hard error rate	4×10^{-11}	4×10^{-12}	4×10^{-13}
Hard error rate (prob/TB)	8×10^{-2}	8×10^{-3}	8×10^{-4}

Table 1: Typical HDD hard error specifications. The SATA column refers to a consumer-grade SATA disk, the Ent. SATA column to an enterprise-grade SATA disk, and the Ent. Hi Perf. column to a high-performance enterprise disk.

This seems to be quite reasonable as a specification, and one might ask why the industry hasn't adopted it. I think the reason is clear—8% per TB doesn't sound very impressive. The 0.08% for high performance enterprise disk isn't terribly comforting, for that matter.

FURTHER ISSUES WITH CURRENT SPECS

Another serious issue for system designers is that the specifications as stated don't give any detail on the root causes of hard errors. As specified, it is related to disk read and write activity. Thus, we might assume that read and write operations create hard errors. This could be from head-disk contact, lube depletion, scratches, etc.

Alternatively, it might be just related to the drive operating, not tied directly to reading and writing. Such effects might include having a high bit error rate, defects in synchronization, defects in servo, process errors (off track positioning), etc.

In reality, all these likely contribute to differing extents, although manufacturers don't publicly release this information. From a system impact point of view, however, we typically need to use public information as a starting point.

Capacity-Related Specification

In storage systems, there are other metrics we can use to understand the impact of hard errors. I find it instructive to recast the hard error specification in terms of the drive capacity. Let's examine the probability of being able to read the entire drive capacity successfully—that is, without a hard error.

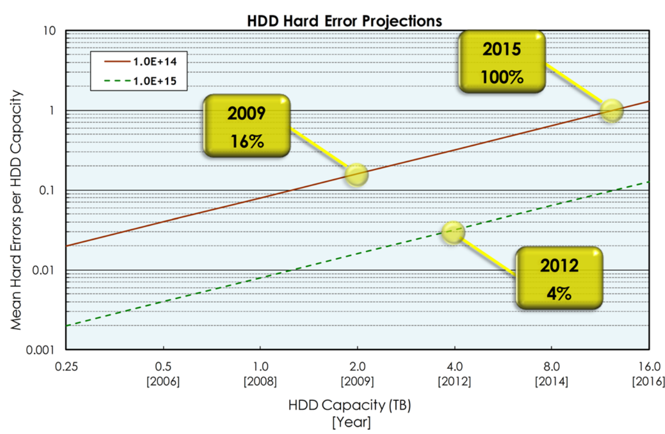

Figure 2: Plot of mean hard errors per HDD capacity vs. HDD capacity. The vertical axis is the log of hard error probability. The horizontal axis shows hard disk capacity in TB by the estimated year of first shipment. The solid line is for consumer-grade SATA disks, and the dashed line for enterprise-grade SATA disks.

In 2009, the largest SATA HDD capacity was 2 TB. At the specified hard error interval of $1x10^{14}$, the probability of encountering a hard error in transferring the entire capacity of the HDD would be 16%. At a capacity growth rate of 40% per year, one would therefore expect to encounter one hard error when transferring the capacity of the drive by 2015. It is interesting to ponder what this means. For enterprise-grade SATA disks, in 2012 the probability of a hard error when transferring the entire capacity of the drive would be 4%. As we shall see, this isn't dramatically different from the situation with consumer-grade SATA disks.

The impacts of such behavior depend on the application. Consider a typical consumer HDD used in a digital video recorder application. First, the result of a hard error might be the loss of a few frames of video. Further, the error rate of the input stream, such as from a satellite or cable, is higher—about 440/TB [5], and thus the contribution from the HDD hard error rate is negligible.

In a desktop PC application, it depends on the average data rate. If we assume that an average desktop HDD is used for 2,000 hours per year and that the time-average data rate for that period is 100 KB/s, this works out to 0.7 TB/year, or a 5% probability of one hard error per year. For reference, this example is the equivalent of 25 IOPS at a 20% duty cycle. In such a case, the impact of such hard error rates are likely to be seen soon, especially as HDD capacity and disk usage grow.

In storage systems, the situation is more acute, as the reliability depends on the ability to read the entire capacity of the disks.

A BRIEF REVIEW OF RAID-5 AND FAILURES

Because of its low cost and low performance overheads, RAID-5 remains a popular choice for protecting storage systems against disk failures. (We can safely ignore the effects of striping in our analysis here, as it doesn't affect the outcome.) RAID-5 uses a single parity disk to protect an array of disks. Since we have a full disk's worth of parity, a single failure will not cause data loss. There are three possibilities of dual failures we must consider in determining the probability of data loss. First, the array can lose a first disk, then lose a second disk before the missing data has been rebuilt onto a spare disk. This is commonly referred to as an *array kill event*. Second, the array can lose one disk, then encounter a hard error while rebuilding the missing data. This is commonly referred to as a *strip kill event*, where some portion of a strip's worth of data is lost. Third, the array can have two hard errors in the same parity strip. Having sector failures line up like this can be ignored to first order, since in hard disks sector failures do not exhibit a tendency to correlate between drives. The probability of two hard errors in a strip is order sector failure squared and the number of strips is also order sector failure (recall Figure 2); thus the resulting probability will be about 10^{-10}, which is too small to impact the failure rate.

QUICK RAID-5 DATA LOSS ESTIMATOR BY DISK CAPACITY (TB)

	1	2	4
Prob. disk loss/year	6.8%	6.8%	6.8%
Rebuild TB	7	14	28
Expected hard errors in rebuild	0.56	1.12	2.24
Prob. strip kill/year	3%	4.7%	6.1%
Prob. 2nd disk in rebuild	$1.6x10^{-4}$	$2.4x10^{-4}$	$3.9x10^{-4}$
Prob. array kill/year	$1.1x10^{-5}$	$1.7x10^{-5}$	$2.6x10^{-5}$
Prob.	3%	4.7%	6.1%

Table 2: Estimated failure probability per year for a 7+P RAID-5, assuming disks have a $1x10^{14}$ hard error specification and 1 MH MTBF. Each column shows the results for arrays built with disk drives of the indicated capacities.

We can easily create a simple RAID-5 loss estimator to see the impact of hard errors. Table 2 shows the results of the estimator for a system with seven data disks and one parity disk, where the disks have a stated reliability of 1 MH MTBF (one million hours mean time between failures). The latter equates to an annual failure rate (AFR) of 0.9% (suddenly a 1 MH MTBF doesn't seem that impressive). Using the AFR value, a binomial can be used to compute the probability of a first disk loss during a year (shown on line 2). The amount of data needing to be read to complete a rebuild after a data loss is seven here (the number of data disks). Given what we have previously shown regarding the hard error rate per TB, it should be clear what the result is going to be. Line 3 shows the expected number of hard errors encountered during the rebuild. Note that in the case of drive capacity > 2 TB, the expected number of hard errors is greater than one. Thus, we expect rebuilds to fail more often than they succeed! Line 4 shows the probability of strip kill per year, which is greater than 1% for all cases shown here. Line 5 shows the probability of a second disk failing during the rebuild process, which here is assumed to take about eight hours. The array kill probability is 100x smaller than the strip kill probability, and thus we need only consider strip kill here for the final probability of data loss. Switching to enterprise-grade SATA disks only reduces the strip kill rate by a factor of 10. So a system built in a year or so with 4 TB enterprise-grade SATA disks would still have >0.5% chance of failure per year, which is not very comforting.

It is interesting to note that we have created an array with a greater failure rate than the disks we built it out of (3% vs. 0.9%)! This is not quite a fair comparison, but it does say that such a design is of dubious reliability.

So it appears that RAID-5 is unsuitable for use with consumer SATA disks. However, as HDD capacities increase, if the hard error specifications remain constant (as they have for 10 years) [6], then enterprise-grade SATA disks and even high-performance enterprise disks will eventually produce the same result. Thus, the movement to stronger protection, such as dual-parity RAID-6, is inevitable.

SCRUBBING

Scrubbing is a technique where the data on the array is checked for errors (scrubbed) while all disks are functional. Any hard errors encountered can thus be corrected. While scrubbing should be able to limit the accumulation of hard errors, one has to wonder how valuable it is at large capacities. Consider a RAID-6 as eight data plus two parities, using 2 TB consumer-grade disks. The scrub reads 20 TB with an 8% per TB chance of a hard error, so we expect the scrub to leave us with about one new hard error in the array! An interesting question is how scrubbing should be implemented, when it will always leave hard errors behind. I liken it to washing a car during rain. The accumulated dirt can be removed, but the car will never be clean.

Importance

The enterprise storage market is increasing its use of lower-cost disks (such as enterprise-grade SATA). However, the data reliability requirements need to remain the same. Thus, action must be taken at the system level to understand and plan for the impacts. Clearly, stronger protection than RAID-5 is warranted for SATA disks of either grade. However, moving to RAID-6 will increase system costs and reduce performance. Scrubbing policies might also require adjustment.

Given the trend lines shown in Figure 2, it would seem beneficial for the HDD hard error rate behavior to be improved.

Hard Errors with SSD

It is tempting to believe that hard errors are a problem associated with magnetic recording and that solid state storage technologies will thus suffer to a far lesser extent. However, this isn't necessarily the case, and the situation warrants examination. For the following discussion, we will consider 2-bit MLC NAND flash as the underlying solid state storage.

NAND flash behaves quite differently from an HDD in response to storage operations (e.g., reading and writing). In a disk drive, the hard error rate has not been shown to significantly depend on how many times a data location is written, or on how old the data is. Experience has shown that the hard error rate may be treated as a constant for a device. The situation with NAND flash is quite different. NAND flash has both a finite endurance (that is, a limit on the number of times a location can be written), and a finite retention. Further, these parameters are coupled—the more a location is written (also called cycled), the shorter the data retention span. Therefore, the bit error rate may be expressed as a surface in three dimensions, as illustrated in Figure 3. The shape must be similar to what's shown. Since it is possible to cycle the device to complete failure even at short cycle times, the error rate multiplier must increase accordingly. Similarly, the data degrades over time; the error rate multiplier must increase with age. Since we expect the surface to be smooth, it will have a shape like that in Figure 3.

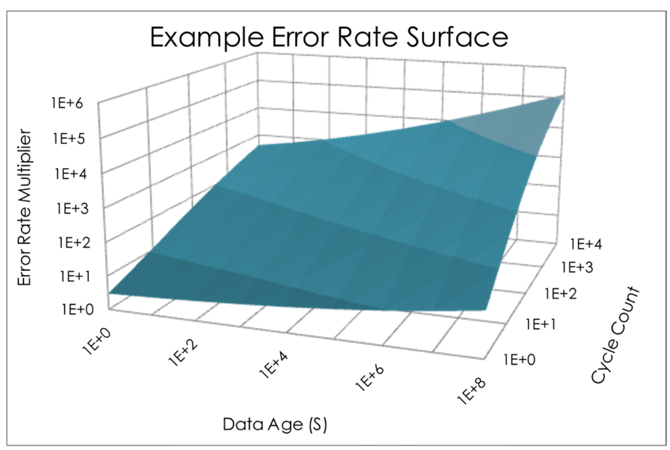

Figure 3: Example NAND flash error rate surface. The horizontal axis is data age in S, from 1 S to 3 years, the depth axis is the cycle count, and the vertical axis is the bit error rate multiplier, relative to the bit error rate at 1 cycle and 1 ms data age [7].

The behavior of flash is more complicated than described above. The bit error rate also depends on temperature and on the number of times a block has been read (read, disturbed). Thus, we have a five-dimensional surface.

System Reliability with SSDs

Reliability Targets

It is important to understand how to create a reliability target when designing a storage system. The targets should be expressed in a manner that reflects usage. Therefore, I propose that the storage system targets be specified as a target probability of a failure event per unit time. I choose this approach since it reflects how the user experiences reliability. Other methods of expressing reliability, such as per byte or per I/O, don't reflect user experience and are difficult for the user to measure (as we have seen for hard errors). Therefore the target should be expressed in a manner that clearly shows whether the system meets the target.

The design of targets is often based on the desired single user experience, such as single customer install. I prefer to use program-based targets, which cover the entire field population for the life of the program. This is how the business team of a storage system manufacturer will determine whether the program will meet its financial targets.

Inputs to a program-based target will include the install base and an estimate of the field failure the program can tolerate. The latter will depend on the type of failure. For example, does the failure result in a warranty event, a loss of access, and data loss, or will it significantly impact the customer's business?

A Simple Program-Based Reliability Target Estimator

Program-level targets are coarse enough not to require high precision. We only need precision to an order of magnitude here. Thus, for the simple estimator we will ignore effects such as duty cycle, read/write ratio, etc.

Table 3 shows an example for an enterprise storage system program using SSDs. Assume we are using enterprise-grade SSDs capable of 30,000 IOPS and that the desired field lifetime for the program is five years. Assume we plan to ship 50,000 SSD units each year. This gives a field population of 250,000 SSDs for the full program. Hard errors are sector-based events for most SSDs, and a typical 4 KB I/O is eight sectors. We can therefore compute the total program sector operations, arriving at $8x10^{18}$ sector operations for the program. Note how large this value is!

SSD unit IOPS	30,000	Enterprise SSD
Field lifetime	5	Years per product
Field SSD units	250,000	For full program
I/O size	4	KB
Total program sector ops	$8x10^{18}$	Program life usage
Field events/program	(0.1–50)	Depends on the event

Table 3: SSD-based program reliability estimator.

As mentioned earlier, the target number of field events depends on the event. If the event causes data loss, a value as high as 50 would equate to about one field data loss event per month. A value of 500 would be one a week, which is clearly too high for an enterprise program. Thus, 50 is an upper bound on acceptability. If the event causes a significant customer disruption, such as a data corruption event, then the target might even be less than 1. A business team might claim that the proper target is 0, but this wouldn't be practical. So a 10% chance of one such event in a program might be reasonable here.

Another important metric can be derived from this information—the number of unit-years of SSD operations. This program represents 1.5 M unit-years of SDD operation.

TEST CAPABILITY

Now that we have a field reliability target, we need to determine how to develop a test program to confirm that the system meets the target. HDD-based systems provide us with decades of experience to draw upon. A typical qualification test for an HDD program is to use 1,000 drives for 1,000 hours. Let's examine an SSD test based on these parameters. Assume that the device can operate at 80% duty cycle during the test, as we need to perform tests on the data. This leaves us with 24,000 IOPS per SSD. Thus, for the full 1,000 SSDs and 1,000 hours, we can test $7x10^{14}$ sector operations. However, this is $9x10^{-5}$ of the program sector operations. Thus, it will be very difficult to extrapolate the program behavior based on the results of such a test. (We haven't examined the confidence levels given 0 or even a few events seen during the test.) The bottom line is that this is only a 100 unit-year test, and we don't have much historical experience to draw upon.

SSD Specifications (Speed Kills)

One approach used to increase the confidence in such a test is to utilize an acceleration factor, such as temperature. There isn't room to examine the efficacy of such an approach in this paper (I have strong reason to suspect that such acceleration is questionable [8]), but we should examine SSD hard error specifications first.

SSDs are capable of significantly higher random I/O rates than hard disks. Given that SSDs are significantly more expensive than hard disks of a given capacity, it is reasonable to assume that they will be used only where the higher performance is required. Thus, we should consider how they behave in high random I/O workload.

The hard error specifications for SSDs are typically just copied from hard disk specification in the same market segment [9]. However, an SSD will be operated at substantially higher I/O rates than an HDD. Thus the hard error specification needs to be scaled accordingly. Let's examine this premise.

	Ent. SATA HDD	**Ent. Perf. HDD**	**SSD consumer**	**SSD enterprise**
IOPS	120	250	10,000	30,000
Equivalent MB/s	0.48	1	40	120
Hard error bit interval spec.	$1x10^{15}$	$1x10^{16}$	$1x10^{15}$	$1x10^{16}$
Mean years/hard errors	66	320	0.8	2.6
Proposed scaled bit interval spec.			$\mathbf{1x10^{17}}$	$\mathbf{1x10^{18}}$

Table 4. SSD and HDD hard error specifications and their impacts.

Table 4 is a comparison of enterprise-grade SATA disk, high performance enterprise disk, consumer-grade SSD, and enterprise-grade SSD. The IOPS row lists the sustained random I/O per second each device might attain. Clearly, the SSDs exhibit superior performance characteristics. Row 2 is a conversion of the IOPS rate to MB/s, assuming a typical 4KB I/O size. The hard error specifications are shown in row 3. The consumer SSD specification is the same as the enterprise SATA HDD, and enterprise SSD and high performance HDD also share the same specification. Given the values in rows 2 and 3, we can easily compute the mean years/hard error, as shown in row 4. It should be obvious given the IOPS ratios between devices, but the SSDs operating at these specifications will be expected to exhibit hard errors substantially more often than HDDs. The consumer SSD would not be expected to go a full year between hard errors.

It would appear that the SSD specifications weren't derived using the above analysis. Since the SSDs are about 100 times faster than the HDDs in the same market segment, they will need a hard error specification 100 times better just to maintain the same reliability as HDDs, as shown in the last row. Claiming SSDs are more reliable than HDDs would require an even higher specification than I have proposed.

Going back to the program example of Table 3, there are roughly 10^{19} sector operations in the program. Table 1 shows that the sector hard error probability for the enterprise SSD specification is $4x10^{-13}$. Thus we might expect something on the order of 10^6 sector errors during the program. Such behavior is normally of little concern, as RAID can be used to correct operational hard errors. However, the situation with SSDs is more complex than with HDDs. RAID reliability computations like those illustrated in Table 2 make two key assumptions regarding failures: failures are assumed to be both independent of each other and independent of time. Unfortunately, neither is true with SSDs. Since flash has a wear-out mechanism where the bit error rate (thus the hard error rate) is a function of the number of times a location has been written, failures are time-dependent. Further, since RAID systems spread the data to evenly distribute the write accesses (identical in the case of a mirror), we expect the failures to correlate across SSDs in an array. Thus both assumptions are false, and we can't assume that RAID will be sufficient. We need to create a more sophisticated RAID model to compute reliability and a more sophisticated RAID design to accommodate SSDs.

If SSDs met my proposed scaled hard error specifications, then the strip kill exposure problem would likely be 100 times smaller. However, proving that SSDs meet such a specification will be problematic. We determined above that a 1,000 SSD 1,000-hour test could perform about $1x10^{15}$ sector operations. My proposed enterprise SSD hard error specification is a sector failure probability of $3x10^{-15}$; thus, the counting statistics in a 10^{15} operation test will be limited. (The nature of the error rate surface will necessitate even further testing to get endurance, retention, temperature, and read-disturb effects.)

SSD Testing

Since a flash cell has a finite (and small) write endurance, SSDs use wear leveling to spread the write load more evenly across the device. This is widely assumed to increase the reliability, as it substantially increases the amount of data that can be written to the device prior to wear out. However, there are side effects to wear leveling that negatively impact the reliability. One of the most significant is that it restricts the ability of the system integrator to test the device.

Wear leveling is essentially a virtual address layer that decouples the external logical block address (LBA) from the internal physical block address. Such devices necessarily perform background operations, such as reclamation, and so an LBA may be relocated without providing external information. This means that an external system can't know the physical location of a given LBA, its write count, or the data age. Therefore it isn't possible to construct a deterministic test using a wear-leveled device. The best one can do is to write the entire device sequentially until wear-out. This is time-prohibitive and still leaves many unknowns.

It is interesting to note that system integrators, who are very reluctant to employ hard disks that they can't test, don't seem to have similar reservations about SSDs. Essentially, they are forced to trust the vendors.

IN FLASH, ERRORS NEVER SLEEP . . .

Another side effect of wear leveling is that it enhances the cluster failure of sectors, since it has a tight distribution of write cycle counts. Coupled with the growth of the bit error rate as the data ages, wear leveling also makes the transition from a working device to one which is significantly out of specification much more rapid. This can be seen in Figure 4, which is from actual measurements of a commercial MLC SSD using 5,000 cycle rated flash. We can see that data of age 1000 hours (solid curve) meets the $4x10^{12}$ sector loss specification at about 3770 cycles. Now, if this data is allowed to age just 48 more hours, the hard error rate doubles! Looked at another way, the cycle specification at 1048 hours is 105 cycles shorter than it is at 1000 hours. Thus, we can see that having a large fraction of the sectors at similar cycle counts in this range can put a substantial fraction at risk in a very short time.

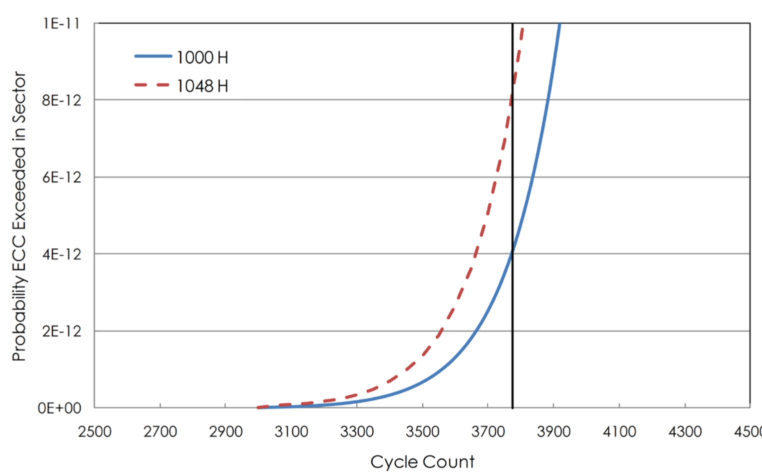

Figure 4: Measured probability of sector hard error vs. cycle count at different data ages. The vertical axis is the probability the ECC is exceeded for a sector (a hard error). The x axis is the write cycle for the sectors. The solid curve is for a data age of 1000 hours, and the dashed curve is for a data age of 1048 hours.

Conclusion

Disk capacities have been growing exponentially for decades, continuing to feed the market's insatiable demand for storage. However, the hard error specifications haven't kept pace. This has led to a situation where it would no longer be advisable to use RAID-5 with enterprise-grade SATA disks. If the hard error rates aren't improved, it will be only a few more years before high performance enterprise disks hit the same limit. While stronger RAID levels, such as dual parity RAID-6 can help, much of the additional protection will go towards hard error protection, as opposed to dual disk failures. It is also likely that scrubbing policies will require modification, as we will reach the point where scrubbing is unlikely to leave disks without hard errors.

It is widely assumed that SSDs are more reliable than hard disks, since they have no moving parts. As we have seen, this assumption isn't necessarily true and can't be verified with the data available. In fact, using SSDs at the JEDEC hard error specifications [9] could actually increase the time rate of failure in enterprise systems. This is because SSDs are deployed in higher-performance applications, but the specifications have been copied from hard disks which don't perform at these levels. I propose that SSDs require more stringent hard error specifications appropriate to their workloads.

There is a significant dearth of published information on the reliability of SSDs. One reason may be that the wear leveling used to increase device lifetime in SSDs has the side effect of preventing users from testing the devices. Properly designing a system using SSDs requires a deep understanding of the error rates. Thus, it would be beneficial for the industry to develop testable SSDs.

References

[1] N. Yezhkova and R. Villars, "Worldwide Enterprise Storage Systems 2010-2014 Forecast Update," IDC #226223, December 2010, p. 27.

[2] Caviar Black specification sheet, Western Digital 2879-701276-A11, March 2011, p. 2: http://www.wdc.com/wdproducts/library/SpecSheet/ENG/2879-701276.pdf.

[3] RE4 specification sheet, Western Digital 2879-701338-A05, September 2010, p. 2: http://www.wdc.com/wdproducts/library/SpecSheet/ENG/2879-701338.pdf.

[4] Ultrastar 15K600 data sheet, Hitachi Global Storage Technologies DSU-S156009EN-03, October 2009, p. 2: http://www.hitachigst.com/tech/techlib.nsf/techdocs/DAC6606EE8DDF5D7862576490028557B/$file/US15K600_DS_final.pdf.

[5] For example, the target error rate for satellite DVB-S2 is 1x10-7 packet error rate, or about 1 in 18 Gb; see Alberto Morello and Vittoria Mignone, "DVB-S2: The Second Generation Standard for Satellite Broad-band Services," *Proceedings of the IEEE*, vol. 94, no. 1, January 2006, p. 214.

[6] For example, DiamondMax D540X data sheet, Maxtor Corp. 9/01 6158v2 update, April 2002. This was a desktop drive with capacities up to 160 GB.

[7] Steven R. Hetzler, "System Impacts of HDD and Flash Reliability," IDEMA Hard Disk and Solid State Drive Reliability Symposium, May 2008.

[8] Steven R. Hetzler, "System Design Impact of Storage Technology Trends," FAST '11: 9th USENIX Conference on File and Storage Technologies, T4, February 2011.

[9] JEDEC Standard 218, "Solid-State Drive (SSD) Requirements and Endurance Test Method," September 2010, p. 7.