As storage systems grow larger and larger, protecting their data for long-term storage is becoming ever more challenging.

BY DAVID S.H. ROSENTHAL

Keeping Bits Safe: How Hard Can It Be?

THESE DAYS, WE are all data pack rats. Storage is cheap, so if there is a chance the data could possibly be useful, we keep it. We know that storage isn't completely reliable, so we keep backup copies as well. But the more data we keep, and the longer we keep it, the greater the chance that some of it will be unrecoverable when we need it.

There is an obvious question we should be asking: how many copies in storage systems with what reliability do we need to get a given probability that the data will be recovered when we need it? This may be an obvious question to ask, but it is a surprisingly difficult question to answer. Let's look at the reasons why.

To be specific, let's suppose we need to keep a petabyte for a century and have a 50% chance that every bit will survive undamaged. This may sound like a lot of data and a long time, but there are already data collections bigger than a petabyte that are important to keep forever. The Internet Archive is

already multiple petabytes.

The state of our knowledge about keeping bits safe can be summarized as:

► *The more copies, the safer.* As the size of the data increases, the per-copy cost increases, reducing the number of backup copies that can be afforded.

► *The more independent the copies, the safer.* As the size of the data increases, there are fewer affordable storage technologies. Thus, the number of copies in the same storage technology increases, decreasing the average level of independence.

► *The more frequently the copies are audited, the safer.* As the size of the data

increases, the time and cost needed for each audit to detect and repair damage increases, reducing their frequency.

At first glance, keeping a petabyte for a century is not difficult. Storage system manufacturers make claims for their products that far exceed the reliability we need. For example, Sun claimed that its ST5800 Honeycomb product had an MTDDL (mean time to data loss) of 2.4×10^6 years.^{a,41} Off-the-shelf solutions appear so reliable that backups are unnecessary. Should we believe these claims? Where do they come from?

All that Sun was saying was if you watched a large number of ST5800 systems for a long time, recorded the time at which each of them first suffered a data loss, and then averaged these times, the result would be 2.4×10^6 years. Suppose Sun watched 10 ST5800s and noticed that three of them lost data during the first year, four of them lost data after 2.4×10^6 years, and the remaining three lost data after 4.8×10^6 years; Sun would be correct that the MTDDL was 2.4×10^6 years. But we would not consider a system with a 30% chance of data loss

after the start of the experiment. As Sirius did not start watching a batch of SC5800s 2.8 million years ago, how would they know?

Sirius says it will sell 2×10^4 SC5800s per year at $\$5 \times 10^4$ each (a \$1 billion-a-year business), and it expects the product to be in the market for 10 years. The SC5800 has a service life of 10 years. So if Sirius watched the entire production of SC5800s ($\$10^{10}$ worth of storage systems) over their entire service life, the experiment would end 20 years from now after accumulating about 2×10^6 system-years of data. If its claim were correct, Sirius would have about a 17% chance of seeing a single data-loss event.

In other words, Sirius claims the probability that *no SC5800 will ever lose any data* is more than 80%. Or, since each SC5800 stores 5×10^{13} bytes, there is an 80% probability that 10^{19} bytes of data will survive 10 years undamaged.

If one could believe Sirius' claim, the petabyte would look pretty safe for a century. But even if Sirius were to base its claim on an actual experiment, it would not provide results for 20 years and even when it did, would not validate the number in question. In fact, claims such as those of Sun and Sirius are not the result of experimentation at all. No feasible experiment could validate them. They are *projections* based on models of how components of the system such as disks and software behave.

Models

The state of the art in this kind of modeling is exemplified by the Pergamum project at UC Santa Cruz.³⁹ Its model includes disk failures at rates derived from measurements^{30,35} and sector failures at rates derived from disk vendor specifications. This system attempts to conserve power by spinning the disks down whenever possible; it makes an allowance for the effect of doing so on disk lifetime, but it is not clear upon what this allowance is based. The Pergamum team reports that the simulations were difficult:

"This lack of data is due to the extremely high reliability of these configurations—the simulator modeled many failures, but so few caused data loss that the simulation ran very slowly.

Before using Sun's claim for the ST5800 as an example, I should stipulate that the ST5800 was an excellent product. It represented the state of the art in storage technology, and Sun's marketing claims represented the state of the art in storage marketing. Nevertheless, Sun did not guarantee that data in the ST5800 would last 2.4×10^6 years. Sun's terms and conditions explicitly disclaimed any liability whatsoever for loss of, or damage to, the data the ST5800 stores⁴⁰ whenever it occurs.

in the first year was adequate to keep a petabyte safe for a century. A single MTDDL number is not a useful characterization of a solution.

Let's look at the slightly more scientific claim made at the recent launch of the SC5800 by the marketing department of Sirius Cybernetics:^b "SC5800 has an MTDDL of $(2.4 \pm 0.4) \times 10^6$ years." Sirius implicitly assumes the failures are normally distributed and thus claims that about two-thirds of the failures would occur between 2.0×10^6 and 2.8×10^6 years

a Numbers are expressed in powers-of-10 notation to help readers focus on the scale of the problems and the extraordinary level of reliability required.

b Purveyors of chatty doors, existential elevators, and paranoid androids to the nobility and gentry of this galaxy.¹

This behavior is precisely what we want from an archival storage system: it can gracefully handle many failure events without losing data. Even though we captured fewer data points for the triple inter-parity configuration, we believe the reported MTTDL is a reasonable approximation.”³⁹

Although the Pergamum team’s effort to obtain “a reasonable approximation” to the MTTDL of its system is praiseworthy, there are a number of reasons to believe it overestimates the reliability of the system in practice:

- The model draws its failures from exponential distributions. The team thus assumes that both disk and sector failures are uncorrelated, although all observations of actual failures^{5,42} report significant correlations. Correlated failures greatly increase the probability of data loss.^{6,13}

- Other than a small reduction in disk lifetime from each power-on event, the Pergamum team assumes that failure rates observed in always-on disk usage translate to the mostly off environment. A study⁴³ published after the Pergamum paper reports a quantitative accelerated life test of data retention in almost-always-off disks. It shows that some of the 3.5-inch disks anticipated by the Pergamum team have data life dramatically worse in this usage mode than 2.5-inch disks using the same head and platter technology.

- The team assumes that disk and sector failures are the only failures contributing to the system failures, although a study¹⁷ shows that other hardware components contribute significantly.

- It assumes that its software is bug-free, despite several studies of file and storage implementations^{14,20,31} that uniformly report finding bugs capable of causing data loss in all systems studied.

- It also ignores all other threats to stored data³⁴ as possible causes of data loss. Among these are operator error, insider abuse, and external attack. Each of these has been the subject of anecdotal reports of actual data loss.

What can such models tell us? Their results depend on both of the following:

- The details of the simulation of the system being studied, which, one

The more data we keep, and the longer we keep it, the greater the chance that some of it will be unrecoverable when we need it.

hopes, accurately reflect its behavior.

- The data used to drive the simulation, which, one hopes, accurately reflects the behavior of the system’s components.

Under certain conditions, it is reasonable to use these models to compare different storage-system technologies. The most important condition is that the models of the two systems use the same data. A claim that modeling showed system A to be more reliable than system B when the data used to model system A had much lower failure rates for components such as disk drives would not be credible.

These models may well be the best tools available to evaluate different techniques for preventing data loss, but they aren’t good enough to answer our question. We need to know the *maximum* rate at which data will be lost. The models assume things, such as uncorrelated errors and bug-free software, that all real-world studies show are false. The models exclude most of the threats to which stored data is subject. In those cases where similar claims, such as those for disk reliability,^{30,35} have been tested, they have been shown to be optimistic. The models thus provide an estimate of the *minimum* data loss rate to be expected.

Metrics

Even if we believed the models, the MTTDL number does not tell us how much data was lost in the average data-loss event. Is petabyte system A with an MTTDL of 10^6 years better than a similar-size system B with an MTTDL of 10^3 years? If the average data-loss event in system A loses the entire petabyte, where the average data-loss event in system B loses a kilobyte, it would be easy to argue that system B was 10^9 times better.

Mean time to data loss is not a useful metric for how well a system stores bits through time, because it relates to time but not to bits. Nor is the UBER (unrecoverable bit error rate) typically quoted by disk manufacturers; it is the probability that a bit will be read incorrectly regardless of how long it has been sitting on the disk. It relates to bits but not to time. Thus, we see that we lack even the metric we would need to answer our question.

Let us oversimplify the problem to

get a clearer picture. Suppose we had eliminated all possible sources of correlated data loss, from operator error to excess heat. All that remained would be *bit rot*, a process that randomly flips the bits the system stores with a constant small probability per unit time. In this model we can treat bits as radioactive atoms, so that the time after which there is a 50% probability that a bit will have flipped is the *bit half-life*.

The requirement of a 50% chance that a petabyte will survive for a century translates into a bit half-life of 8×10^{17} years. The current estimate of the age of the universe is 1.4×10^{10} years, so this is a bit half-life approximately 6×10^7 times the age of the universe.

This bit half-life requirement clearly shows the high degree of difficulty of the problem we have set for ourselves. Suppose we want to know whether a system we are thinking of buying is good enough to meet the 50% chance of keeping a petabyte for a century. Even if we are sublimely confident that every source of data loss other than bit rot has been totally eliminated, we still have to run a benchmark of the system's bit half-life to confirm it is longer than 6×10^7 times the age of the universe. And this benchmark has to be complete in a year or so; it can't take a century.

So we take 10^3 systems just like the one we want to buy, write a petabyte of data into each so we have an exabyte of data altogether, wait a year, read the exabyte back, and check it. If the system is just good enough, we might see five bit flips. Or, because bit rot is a random process, we might see more, or less. We would need either a lot more than an exabyte of data or a lot more than a year to be reasonably sure the bit half-life was long enough for the job. But even an exabyte of data for a year costs 10 times as much as the system we want to buy.

What this thought-experiment tells us is we are now dealing with such large numbers of bits for such a long time that we are never going to know whether the systems we use are good enough:

- The known causes of data loss are too various and too highly correlated for models to produce credible projections.
- Even if we ignore all those causes,

Our inability to compute how many backup copies we need to achieve a reliability target is something we are just going to have to live with. We are not going to have enough backup copies, and stuff will get lost or damaged.

the experiments that would be needed to be reasonably sure random bit rot is not significant are too expensive, or take too long, or both.

Measuring Failures

It wasn't until 2007 that researchers started publishing studies of the reliability that actual large-scale storage systems were delivering in practice. Enterprises such as Google⁹ and institutions such as the Sloan Digital Sky Survey³⁷ and the Large Hadron Collider⁸ were collecting petabytes of data with long-term value that had to remain online to be useful. The annual cost of keeping a petabyte online was more than \$1 million.²⁷ It is easy to see why questions of the economics and reliability of storage systems became the focus of researchers' attention.

Papers at the 2007 File and Storage Technologies (FAST) conference used data from NetApp³⁵ and Google³⁰ to study disk-replacement rates in large storage farms. They showed that the manufacturer's MTTF numbers were optimistic. Subsequent analysis of the NetApp data¹⁷ showed that all other components contributed to the storage system failures and that:

"Interestingly, [the earlier studies] found disks are replaced much more frequently (2–4 times) than vendor-specified [replacement rates]. But as this study indicates, there are other storage subsystem failures besides disk failures that are treated as disk faults and lead to unnecessary disk replacements."¹⁷

Two studies, one at CERN (European Organization for Nuclear Research)¹⁸ and one using data from NetApp,⁵ greatly improved on earlier work using data from the Internet Archive.^{6,36} They studied *silent data corruption*—events in which the content of a file in storage changes with no explanation or recorded errors—in state-of-the-art storage systems.

The NetApp study looked at the incidence of silent storage corruption in individual disks in RAID arrays. The data was collected over 41 months from NetApp's filers in the field, covering more than 1.5×10^6 drives. The study found more than 4×10^5 silent corruption incidents. More than 3×10^4 of them were not detected until RAID restoration and could thus have caused

data loss despite the replication and auditing provided by NetApp's row-diagonal parity RAID.¹¹

The CERN study used a program that wrote large files into CERN's various data stores, which represent a broad range of state-of-the-art enterprise storage systems (mostly RAID arrays), and checked them over a period of six months. A total of about 9.7×10^{16} bytes was written and about 1.92×10^8 bytes were found to have suffered silent corruption, of which about two-thirds were persistent; rereading did not return good data. In other words, about 1.2×10^{-9} of the data written to CERN's storage was permanently corrupted within six months. We can place an upper bound on the bit half-life in this sample of current storage systems by assuming the data was written instantly at the start of the six months and checked instantly at the end; the result is 2×10^8 or about 10^{-2} times the age of the universe. Thus, to reach the petabyte for a century requirement we would need to improve the performance of current enterprise storage systems by a factor of at least 10^9 .

Tolerating Failures

Despite manufacturers' claims, current research shows that state-of-the-art storage systems fall so many orders of magnitude below our bit preservation requirements that we cannot expect even dramatic improvements in technology to fill the gap. Maintaining a single replica in a single storage system is not an adequate solution to the bit preservation problem.

Practical digital preservation systems must therefore:

- Maintain more than one copy by *replicating* their data on multiple, ideally different, storage systems.
- Audit or (*scrub*) the replicas to detect damage, and repair it by overwriting the known-bad copy with data from another.

The more replicas and the more frequently they are audited and repaired, the longer the bit half-life we can expect. This is, after all, the basis for the backups and checksums technique in common use. In fact, current storage systems already use such techniques internally—for example, in the form of RAID.²⁹ Despite this, the bit half-life they deliver is inadequate. Unfortu-

nately, adding the necessary inter-storage-system replication and scrubbing is expensive.

Cost figures from the San Diego Supercomputer Center^c in 2008 show that maintaining a single online copy of a petabyte for a year costs about $\$1.05 \times 10^6$. A single near-line copy on tape costs about $\$4.2 \times 10^5$ a year. These costs decrease with time, albeit not as fast as raw disk costs. The British Library estimates a 30% per annum decrease. Assuming this rate continues for at least a decade, if you can afford about 3.3 times the first year's cost to

we need to be, and thus the cost of the necessary replication. At small scales the response to this uncertainty is to add more replicas, but as the scale increases this rapidly becomes unaffordable.

Replicating among identical systems is much less effective than replicating among diverse systems. Identical systems are subject to common mode failures—for example, those caused by a software bug in all the systems damaging the same data in each. On the other hand, purchasing and operating a number of identical systems

store an extra replica for a decade, you can afford to store it indefinitely. So, adding a second replica of a petabyte on disk would cost about $\$3.5 \times 10^6$ and on tape about $\$1.4 \times 10^6$. Adding cost to a preservation effort to increase reliability in this way is a two-edged sword: doing so necessarily increases the risk that preservation will fail for economic reasons.

Further, without detailed understanding of the rates at which different mechanisms cause loss and damage, it still is not possible to answer the question we started with and to know how many replicas would make us as safe as

will be considerably cheaper than operating a set of diverse systems.

Each replica is vulnerable to loss and damage. Unless they are regularly audited they contribute little to increasing bit half-life. The bandwidth and processing capacity needed to scrub the data are both costly, and adding these costs increases the risk of failure. Custom hardware²⁵ could compute the SHA-1²⁸ checksum of a petabyte of data in a month, but doing so requires impressive bandwidth—the equivalent of three gigabit Ethernet interfaces running at full speed the entire month. User access to data in long-term storage is typically infrequent; it is therefore rarely architected to pro-

c Figures for 2007 are in Moore et al.²⁷

vide such high-bandwidth read access. System cost increases rapidly with I/O bandwidth, and the additional accesses to the data (whether on disk or on tape) needed for scrubbing themselves potentially increase the risk of failure.

The point of writing software that reads and verifies the data-systems store in this way is to detect damage and exploit replication among systems to repair it, thereby increasing bit half-life. How well can we do this? RAID is an example of a software technique of this type applied to disks. In practice, the CERN study¹⁸ looking at real RAID

probabilities, but to illustrate the advantage of using a model checker, and discuss potential trade-offs between different protection schemes.”²⁰

Thus, although intersystem replication and scrubbing are capable of decreasing the incidence of data loss, they cannot eliminate it completely. And the replication and scrubbing software itself will contain bugs that can cause data loss. It must be doubtful that we can implement these techniques well enough to increase the bit half-life of systems with an affordable number of replicas by 10^9 .

of 2×10^5 stone DVDs. A lot can happen to a stack that big in 100 years. Truly magic media that are utterly reliable would make the problems better, but they would not make them go away completely.

I remember magnetic bubble memory, so I have a feeling of *déjà vu*, but it is starting to look possible that flash memory, or possibly more exotic solid-state technologies such as memristors or phase-change memory, may supplant disks. There is a lot to like about these technologies for long-term storage, but will they improve storage reliability?

Again, we don’t know the answer yet. Despite flash memory’s ubiquity, it is not even clear yet how to measure its UBER:

“UBER values can be much better than 10^{-15} but UBER is a strong function of program/erase cycling and subsequent retention time, so UBER specifications must be coupled with maximum specifications for these quantities.”²⁶

In other words, it depends how you use it, which does not appear to be the case for disk. Flash memory used for long-term data storage, which is written once and read infrequently, should in principle perform very well. And the system-level effects of switching from hard disk to flash can be impressive:

“FAWN [fast array of wimpy nodes] couples low-power embedded CPUs to small amounts of local flash storage, and balances computation and I/O capabilities to enable efficient, massively parallel access to data. ...FAWN clusters can handle roughly 350 key-value queries per joule of energy—two orders of magnitude more than a disk-based system.”³

Fast CPUs, fast RAM, and fast disks all use lots of power, so the low power draw of FAWN is not a surprise. But the high performance comes from another aspect of disk evolution. Table 1 shows how long it would take to read the whole of a state-of-the-art disk of various generations.

Disks have been getting bigger but they have not been getting equivalently faster. This is to be expected; the data rate depends on the inverse of the diameter of a bit, but the capacity depends on the inverse of the area of a bit. FAWN nodes can read their entire con-

systems from the outside showed a significant rate of silent data corruption, and the NetApp study⁵ looking at them from the inside showed a significant rate of silent disk errors that would lead to silent data corruption. A study²⁰ of the full range of current algorithms used to implement RAID found flaws leading to potential data loss in all of them. Both this study, and another from IBM,¹⁶ propose improvements to the RAID algorithms but neither claims it can eliminate silent corruption, or even accurately predict its incidence:

“While we attempt to use as realistic probability numbers as possible, the goal is not to provide precise data-loss

Magic Media

Considering the difficulties facing disk-drive technology,¹² the reliability storage systems achieve is astonishing, but it clearly isn’t enough. News sites regularly feature stories reporting claims that some new storage medium has solved the problem of long-term data storage. Synthetic stone DVDs²³ claimed to last 1,000 years were a recent example. These claims should be treated as skeptically as those of Sun and other storage system manufacturers. It may well be that the media in question are more reliable than their competitors, but as we have seen, raw media reliability is only a part of the story. Our petabyte would be a stack

tents very quickly, useful for scrubbing, as well as answering queries.

This is all encouraging, but once it became possible to study the behavior of disk storage at a large scale, it became clear that system-level reliability fell far short of the media specifications. This should make us cautious about predicting a revolution from flash or any other new storage technology.

Economics

Ever since Clayton Christensen published *The Innovator's Dilemma*¹⁰ it has been common knowledge that disk-drive cost per byte halves every two years. So you might argue that you don't need to know how many copies you need to keep your data safe for the long term, you just need to know how many you need to keep it safe for the next few years. After that, you can keep more copies.

In fact, what has happened is the capacity at constant cost has been doubling every two years, which is not quite the same thing. As long as this exponential grows faster than you generate new data, adding copies through time is a feasible strategy.

Alas, exponential curves can be deceiving. Moore's Law has continued to deliver smaller and smaller transistors. A few years ago, however, it effectively ceased delivering faster and faster CPU clock rates. It turned out that, from a business perspective, there were more important things to spend the extra transistors on than making a single CPU faster. Like putting multiple CPUs on a chip.

At a recent Library of Congress meeting, Dave Anderson of Seagate warned⁴ that something similar is about to happen to hard disks. Technologies—HAMR (heat-assisted magnetic recording) and BPM (bit pattern media)—are in place to deliver the

Disks have been getting bigger but they have not been getting equivalently faster. This is to be expected; the data rate depends on the inverse of the diameter of a bit, but the capacity depends on the inverse of the area of a bit.

2013 disk generation (that is, a consumer 3.5-inch drive holding 8TB). But the business case for building it is weak. The cost of the transition to BPM in particular is daunting.²⁴ Laptops, netbooks, and now tablets are destroying the market for the desktop boxes that 3.5-inch drives go into. And very few consumers fill up the 2009 2TB disk generation, so what value does having an 8TB drive add? Let alone the problem of how to back up an 8TB drive on your desk!

What is likely to happen—indeed, is already happening—is that the consumer market will transition rather quickly to 2.5-inch drives. This will eliminate the high-capacity \$100 3.5-inch drive, since it will no longer be produced in consumer quantities. Consumers will still buy \$100 drives, but they will be 2.5 inches and have perhaps one-third the capacity. For a while the \$/byte curve will at best flatten, and more likely go up. The problem this poses is that large-scale disk farms are currently built from consumer 3.5-inch drives. The existing players in the market have bet heavily on the exponential cost decrease continuing; if they're wrong, it will be disruptive.

The Bigger Picture

Our inability to compute how many backup copies we need to achieve a reliability target is something we are just going to have to live with. In practice, we are not going to have enough backup copies, and stuff will get lost or damaged. This should not be a surprise, but somehow it is. The fact that bits can be copied correctly leads to an expectation that they always *will* be copied correctly, and then to an expectation that digital storage will be reliable. There is an odd cognitive dissonance between this and people's actual experience of digital storage, which is that loss and damage are routine occurrences.²²

The fact that storage is not reliable enough to allow us to ignore the problem of failures is just one aspect of a much bigger problem looming over computing as it continues to scale up. Current long-running petascale high-performance computer applications require complex and expensive checkpoint and restart schemes, because the probability of a failure during execution is so high that restarting from

Table 1. The time to read an entire disk of various generations.

1990	240
2000	720
2006	6450
2009	8000
2013	12800

Table 2. The four cases of message digest comparison.

Digest	Match	No Match
Unchanged	Data OK	Data bad
Changed	Deliberate alteration	Data and/or digest bad

scratch is infeasible. This approach will not scale to the forthcoming generation:

“...it is anticipated that exascale systems will experience various kinds of faults many times per day. It is also anticipated that the current approach for resilience, which relies on automatic or application-level checkpoint-restart, will not work because the time for checkpointing and restarting will exceed the mean time to failure of a full system. ...

“Some projections estimate that, with the current technique, the time to checkpoint and restart may exceed the mean time to interrupt of top supercomputers before 2015. This not only means that a computation will do little progress; it also means that fault-handling protocols have to handle multiple errors—current solutions are often designed to handle single errors.”⁷

Just as with storage, the numbers of components and interconnections are so large that the incidence of failures is significant, and the available bandwidths are relatively so low that recovering from the failures is time consuming enough that multiple failure situations have to be handled. There is no practical, affordable way to mask the failures from the applications. Application programmers will need to pay much more attention to detecting and recovering from errors in their environment. To do so they will need both the APIs and the system environments implementing them to become much more failure-aware.

API Enhancements

Storage APIs are starting to move in this direction. Recent interfaces to storage services² allow the application's write call to provide not just a pointer to the data and a length, but also, optionally, the application's message digest of the data. This allows the storage system to detect whether the data was damaged during its journey from the ap-

plication to the device, or while it was sitting in the storage device, or being copied back to the application. Recent research has shown the memory buffers⁴⁴ and data paths¹⁷ between the application and the storage devices contribute substantially to errors.

Let's take the Amazon S3 (Simple Storage Service) REST API² as an example to show that, while these developments are welcome, they are far from a panacea. The PUT request supports an optional (and recommended) Content-MD5 (Message-Digest algorithm 5) header containing the application's digest of the data. The responses to most requests, including PUT, include an ETag (entity tag) header with the service's MD5 of the object. The application can remember the digest it computed before the PUT and, when the PUT returns, verify that the service's digest matches.

Doing so is a wise precaution, but all it really tells the application is that the service knows what the application thinks is the correct digest. The service knows this digest, not because it computed the digest of the correct data, but because the application provided it in the Content-MD5 header. A malign or malfunctioning service could respond correctly to PUT and HEAD requests by remembering the application's digest, without ever storing the data or computing its digest.

The application could try to detect a malign or malfunctioning service by using a GET to obtain the stored data, computing the digest (a) of the returned data, and comparing that with (b) either the digest in the response's ETag header or the digest it computed before the original PUT and remembered (which should be the same). It might seem that there are two cases: if the two message digests match, then the data is OK;^d otherwise it isn't. There are actually four

cases, as shown in Table 2, depending on whether the digest (b) is unchanged or not. The four cases illustrate two problems:

- The bits forming the digest are no different from the bits forming the data; neither is magically incorruptible. A malign or malfunctioning service could return bad data with a digest in the ETag header that matched the data but was not the digest originally computed. Applications need to know whether the digest has been changed. A system for doing so without incorruptible storage is described in Haber et al.¹⁵

- Given the pricing structure for cloud storage services such as Amazon S3, it is too expensive to extract the entire data at intervals to confirm it is being stored correctly. Some method in which the service computes the digest of the data is needed, but simply asking the service to return the digest of a stored object is not an adequate check.³³ The service must be challenged to prove its object is good. The simplest way to do this is to ask the service to compute the digest of a nonce (a random string of bits) and the object; because the service cannot predict the nonce, a correct response requires access to the data after the request is received. Systems using this technique are described in Maniatis et al.²¹ and Shah et al.³⁸

Early detection is a good thing: the shorter the time between detection and repair, the smaller the risk that a second error will compromise the repair. But detection is only part of the solution; the system also has to be able to repair the damaged data. It can do so only if it has replicated the data elsewhere—and some deduplication layer has not optimized away this replication.

Conclusion

It would be nice to end on an upbeat note, describing some technological fix that would allow applications to ignore the possibility of failures in their environment, and specifically in long-term storage. Unfortunately, in the real world, failures are inevitable. As systems scale up, failures become more frequent. Even throwing money at the problem can only reduce the incidence of failures, not exclude them entirely.

^d Assuming the digest algorithm has not been broken, not a safe assumption for MD5.¹⁹

Applications in the future will need to be much more aware of, and careful in responding to, failures.

The high-performance computing community accurately describes what needs to be done:

“We already mentioned the lack of coordination between software layers with regards to errors and fault management. Currently, when a software layer or component detects a fault it does not inform the other parts of the software running on the system in a consistent manner. As a consequence, fault-handling actions taken by this software component are hidden to the rest of the system. ...In an ideal world, if a software component detects a potential error, then the information should propagate to other components that may be affected by the error or that control resources that may be responsible for the error.”⁷

In particular, as regards storage, APIs should copy Amazon’s S3 by providing optional data-integrity capabilities that allow applications to perform end-to-end checks. These APIs should be enhanced to allow the application to provide an optional nonce that is prepended to the object data before the message digest reported to the application is computed. This would allow applications to exclude the possibility that the reported digest has been remembered rather than recomputed.

Acknowledgments

Grateful thanks are due to Eric Allman, Kirk McKusick, Jim Gettys, Tom Lipkis, Mark Compton, Petros Maniatis, Mary Baker, Fran Berman, Tsutomu Shimomura, and the late Jim Gray. Some of this material was originally presented at iPRES 2008 and subsequently published in the *International Journal of Digital Curation*.³²

This work was funded by the member libraries of the LOCKSS Alliance, and the Library of Congress’ National Digital Information Infrastructure and Preservation Program. Errors and omissions are the author’s own.

Related articles on queue.acm.org

Triple-Parity RAID and Beyond

Adam Leventhal

<http://queue.acm.org/detail.cfm?id=1670144>

Hard Disk Drives: The Good, the Bad and the Ugly!

Jon Elerath

<http://queue.acm.org/detail.cfm?id=1317403>

You Don't Know Jack about Disks

Dave Anderson

<http://queue.acm.org/detail.cfm?id=864058>

References

- Adams, D. *The Hitchhiker's Guide to the Galaxy*. British Broadcasting Corp., 1978.
- Amazon. Amazon S3 API Reference (Mar. 2006); <http://docs.amazonwebservices.com/AmazonS3/latest/API/>.
- Andersen, D.G., Franklin, J., Kaminsky, M., Phanishayee, A., Tan, L., Vasudevan, V. FAWN: A fast array of wimpy nodes. In *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles* (2009), 1–14.
- Anderson, D. Hard drive directions (Sept. 2009); http://www.digitalpreservation.gov/news/events/other_meetings/storage09/docs/2-4_Anderson-seagate-v3_HDTrends.pdf.
- Bairavasundaram, L., Goodson, G., Schroeder, B., Arpaci-Dusseau, A.C., Arpaci-Dusseau, R.H. An analysis of data corruption in the storage stack. In *Proceedings of 6th Usenix Conference on File and Storage Technologies*, (2008).
- Baker, M., Shah, M., Rosenthal, D.S.H., Roussopoulos, M., Maniatis, P., Giuli, T.J., Bungale, P. A fresh look at the reliability of long-term digital storage. In *Proceedings of EuroSys2006*, (Apr. 2006).
- Cappello, F., Geist, A., Gropp, B., Kale, S., Kramer, B., Snir, M. Toward exascale resilience. Technical Report TR-JLPC-09-01. INRIA-Illinois Joint Laboratory on Petascale Computing, (July 2009).
- CERN. Worldwide LHC Computing Grid, 2008; <http://lcg.web.cern.ch/LCG/>.
- Chang, F., Dean, J., Ghemawat, S., Hsieh, W. C., Wallach, D. A., Burrows, M., Chandra, T., Fikes, A., Grube, R.E. Bigtable: A distributed storage system for structured data. In *Proceedings of the 7th Usenix Symposium on Operating System Design and Implementation*, (2006), 205–218.
- Christensen, C.M. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Harvard Business School Press (June 1997), Cambridge, MA.
- Corbett, P., English, B., Goel, A., Grcanac, T., Kleiman, S., Leong, J., Sankar, S. Row-diagonal parity for double disk failure correction. In *3rd Usenix Conference on File and Storage Technologies* (Mar. 2004).
- Elerath, J. Hard-disk drives: The good, the bad, and the ugly. *Commun. ACM* 52, 6 (June 2009).
- Elerath, J.G., Pecht, M. Enhanced reliability modeling of RAID storage systems. In *Proceedings of the 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, (2007), 175–184.
- Engler, D. A system's hackers crash course: techniques that find lots of bugs in real (storage) system code. In *Proceedings of 5th Usenix Conference on File and Storage Technologies*, (Feb. 2007).
- Haber, S., Stornetta, W.S. How to timestamp a digital document. *Journal of Cryptology* 3, 2 (1991), 99–111.
- Hafner, J.L., Deenadhyayan, V., Belluomini, W., Rao, K. Undetected disk errors in RAID arrays. *IBM Journal of Research & Development* 52, 4/5, (2008).
- Jiang, W., Hu, C., Zhou, Y., Kanevsky, A. Are disks the dominant contributor for storage failures? A comprehensive study of storage subsystem failure characteristics. In *Proceedings of 6th Usenix Conference on File and Storage Technologies*, (2008).
- Kellemen, P. Silent corruptions. In *8th Annual Workshop on Linux Clusters for Super Computing*, (2007).
- Klima, V. Finding MD5 collisions—A toy for a notebook. Cryptology ePrint Archive, Report 2005/075; <http://eprint.iacr.org/2005/075>.
- Kriukov, A., Bairavasundaram, L.N., Goodson, G.R., Srinivasan, K., Thelen, R., Arpaci-Dusseau, A.C., Arpaci-Dusseau, R.H. Parity lost and parity regained. In *Proceedings of 6th Usenix Conference on File and Storage Technologies*, (2008).
- Maniatis, P., Roussopoulos, M., Giuli, T.J., Rosenthal, D.S.H., Baker, M., Muliadi, Y. Preserving peer replicas by rate-limited sampled voting. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles*, (Oct. 2003), 44–59.
- Marshall, C. “It’s like a fire. You just have to move on.” Rethinking personal digital archiving. In *6th Usenix Conference on File and Storage Technologies*, (2008).
- Mearian, L. Start-up claims its DVDs last 1,000 years. *Computerworld*, (Nov. 2009).
- Mellor, C. Drive suppliers hit capacity increase difficulties. *The Register*, (July 2010).
- Michail, H.E., Kakarountas, A.P., Theodoridis, G., Goutis, C.E. A low-power and high-throughput implementation of the SHA-1 hash function. In *Proceedings of the 9th WSEAS International Conference on Computers*, (2005).
- Mielke, N., Marquart, T., Wu, N., Kessenich, J., Belgal, H., Schares, E., Trivedi, F., Goodness, E., Nevill, L.R. Bit error rate in NAND flash memories. In *46th Annual International Reliability Physics Symposium*, (2008), 9–19.
- Moore, R. L., D'Aoust, J., McDonald, R. H., Minor, D. Disk and tape storage cost models. In *Archiving 2007*.
- National Institute of Standards and Technology (NIST). *Federal Information Processing Standard Publication 180-1: Secure Hash Standard*, (Apr. 1995).
- Patterson, D. A., Gibson, G., Katz, R.H. A case for redundant arrays of inexpensive disks (RAID). In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, (June 1988), 109–116.
- Pinheiro, E., Weber, W.-D., Barroso, L. A. Failure trends in a large disk drive population. In *Proceedings of 5th Usenix Conference on File and Storage Technologies*, (Feb. 2007).
- Prabhakaran, V., Agrawal, N., Bairavasundaram, L., Gunawi, H., Arpaci-Dusseau, A.C., Arpaci-Dusseau, R.H. IRON file systems. In *Proceedings of the 20th Symposium on Operating Systems Principles*, (2005).
- Rosenthal, D.S.H. Bit preservation: A solved problem? *International Journal of Digital Curation* 1, 5 (2010).
- Rosenthal, D.S.H. LOCKSS: Lots of copies keep stuff safe. In *NIST Digital Preservation Interoperability Framework Workshop*, (Mar. 2010).
- Rosenthal, D.S.H., Robertson, T.S., Lipkis, T., Reich, V., Morabito, S. Requirements for digital preservation systems: a bottom-up approach. *D-Lib Magazine* 11, 11 (2005).
- Schroeder, B., Gibson, G. Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you? In *Proceedings of 5th Usenix Conference on File and Storage Technologies* (Feb. 2007).
- Schwarz, T., Baker, M., Bassi, S., Baumgart, B., Flagg, W., van Imngen, C., Joste, K., Manasse, M., Shah, M. Disk failure investigations at the Internet Archive. In *Work-in-Progress Session, NASA/IEEE Conference on Mass Storage Systems and Technologies*, (2006).
- SDSS (Sloan Digital Sky Survey), 2008; <http://www.sdss.org/>.
- Shah, M.A., Baker, M., Mogul, J.C., Swaminathan, R. Auditing to keep online storage services honest. In *11th Workshop on Hot Topics in Operating Systems*, (May 2007).
- Storer, M.W., Greenan, K. M., Miller, E.L., Voruganti, K. Pergamum: replacing tape with energy-efficient, reliable, disk-based archival storage. In *Proceedings of 6th Usenix Conference on File and Storage Technologies*, (2008).
- Sun Microsystems. Sales Terms and Conditions, Section 11.2, (Dec. 2006); http://store.sun.com/CMTemplate/docs/legal_terms/TnE.jsp#11.
- Sun Microsystems. ST5800 presentation. Sun PASIG Meeting, (June 2008).
- Talagala, N. *Characterizing large storage systems: error behavior and performance benchmarks*. Ph.D. thesis, Computer Science Division, University of California at Berkeley, (Oct. 1999).
- Williams, P., Rosenthal, D. S. H., Roussopoulos, M., Georgis, S. Predicting the archival life of removable hard disk drives. In *Archiving 2008*, (June 2008).
- Zhang, Y., Rajimwale, A., Arpaci-Dusseau, A.C., Arpaci-Dusseau, R.H. End-to-end data integrity for file systems: A ZFS case study. In *8th Usenix Conference on File and Storage Technologies*, (2010).

David S.H. Rosenthal has been an engineer in Silicon Valley for a quarter of a century, including as a Distinguished Engineer at Sun Microsystems and employee #4 at NVIDIA. For the last decade he has been working on the problems of long-term digital preservation under the auspices of the Stanford Library.

© 2010 Author 0001-0782/10/1100 \$10.00