



Dragon Slayer Consulting

Marc Staimer

W H I T E P A P E R

Ignore the Impending RAID Catastrophe at Your Own Risk

Ignore the Impending RAID Catastrophe at Your Own Risk

RAID is no longer the answer to data resilience, but has become the problem

Marc Staimer, President & CDS of Dragon Slayer Consulting
marcstaimer@mac.com 503-579-3763

Executive Summary

In the current age of stored data exponential growth (or “current age of exponential data growth?”), RAID (redundant array of independent disks) technology is quickly hitting its limitations. Data loss events are becoming more frequent. More disturbing is that the size of the losses is expanding. There is a direct correlation to RAID’s shortcomings. This paper explains why and how RAID no longer is capable of doing the job it was designed to do. It will additionally explore how RAID:

- Creates difficult operational headaches for IT organizations;
- Devastates IT budgets with expensive workarounds that defers instead of solving the problem;
- Sets up IT for catastrophic data failures.

Finally, it will show can be done right now to solve these problems and prevent unintended and unwanted data loss catastrophes from occurring.

Introduction

RAID has become so ubiquitous when it comes to data storage that it is often forgotten that its purpose is to prevent data loss in the event one or more hard disk drive (HDD) fails. This is because HDDs fail and fail relatively frequently. It’s simply a fact of life.

The HDD is an electro-mechanical device that has the highest probability of a failure and the lowest mean time between failures (MTBF) of any component within a storage system. As the number of HDDs increase within a storage system, so statistically do the number of failures. Failures can be failed disk sectors, or un-recoverable bit errors (UBE). The manufacturers’ Enterprise SAS or FC drives rated MTBF is approximately 1 failure in every 1.5M hours. For SATA drives it is 1 failure in every 600K hours. Wow, if that seems counter to most IT pros’ real-world experience, it is. There are two reasons for this. The first is that manufacturer calculated MTBF must be reduced by the number of drives in the system. If there are 240 drives, then the MTBF is divided by 240. This means one of those 240 SAS/FC drives will fail every 6,250 hours or one drive failure every 8 to 9 months, which is a pretty comforting .8% failure rate. For SATA that drops to approximately one drive failure every 104 days. These numbers probably feel a bit counter-intuitive and should based on real-world results.

Numerous published studies expose a much lower practical MTBF for HDDs. The one most often cited is "[Disk Failures in the Real World: What Does an MTTF of 1,000,000 Hours Mean to you](#)" by Bianca Schroeder and Garth Gibson of Carnegie Mellon University in Pittsburgh.

Adding in that extraordinarily rapid annual storage growth of between 50 and 62% per year (per the analysts IDC, Gartner, 451, Forrester, DSC, and others), projects into a serious IT problem. That problem is unintended, unwanted data loss, which is never ever a good thing. Most storage pros are thinking right now that this cannot be much of a problem because that’s what RAID solves. After all, that is the IT conventional wisdom. Regrettably, as is

Summary of "Disk Failures in the Real World: What Does an MTTF of 1,000,000 Hours Mean to you"

The study tracked roughly 100,000 drives used at large-scale storage sites of high-performance computing and Web servers. The data suggests HDD reliability is considerably overstated. This study clearly determined that a typical FC, SAS, or SATA HDD (it did not matter the drive type or RPM), is more likely to have a realistic MTBF of approximately 6 years or 52,560 hours. This translates into a 16.67% percent annual HDD failure rate, or more than 20 times worse than published MTBFs of Enterprise class FC and SAS. Little cited and perhaps more interesting was that they found failure rates occurred much earlier than expected. HDDs demonstrated significant early onset of wear degradation that grew constantly with age that conventional wisdom believes does not occur until after a nominal lifetime of 5 years. Meaning, the percentage of HDD failures started higher and continued to increase each year.

often the case, conventional wisdom is unequivocally wrong.

IT RAID operational problems


The first rule of storage is similar to doctors Hippocratic oath: “Do no harm”. RAID is no longer meeting this very basic principle. RAID combines multiple HDDs into a single logical unit. When a HDD failure occurs, the data from that failed HDD or multiple HDDs is recreated from parity or copied from a mirror, depending on RAID set type (see Appendix A on RAID sets). It doesn’t matter which type of RAID or parity is being utilized, a HDD failure requires the data be rebuilt in some manner. And rebuilds have become quite ugly.

When HDD density was 9GB, rebuilds occurred in minutes. Storage system performance impacts were minimal. Risk of data loss was minimal. With current HDD densities at 3TB and growing to 4TB by the end of 2011, rebuilds take quite a bit more time, storage system performance hits are devastating, and data loss risk is exponentially higher. Take the example of a 2TB HDD in a common RAID-5 group. Instead of minutes, the rebuild takes approximately 50 to 60 hours if it’s set as a high priority. But, high prioritization reduces storage system performance by as much as 50% or even more. Few IT organizations can tolerate that much storage system performance reduction for 2 and ½ days. Therefore, rebuild prioritization is often set as a background task reducing that performance penalty to more tolerable levels. While this eliminates most of the negative performance impact on the storage system, it also lengthens that rebuild time by as much as 7X or 700% to more than 2 weeks. That’s more than 2 weeks when a second HDD failure or un-recoverable read error (URE) in that RAID group will cause complete data loss in the group.

Unrecoverable Read Errors
 Unrecoverable read errors are based on the HDD’s bit error rate. Those old 9GB HDDs had a bit error rate of approximately 1 in 10E¹⁴ bits. HDD capacity density today has increased more than 333 times (3TB) while the projected bit error rate has improved only 10 times for Enterprise class SATA (10E¹⁵) and 100 times for Enterprise class SAS or FC (10E¹⁶). Sadly, just as real-world HDD failure rates are much high than the projected MTBF, so too are the bit error rates. HDD bit error rates increase rapidly when they’re placed in drive enclosures where variables such as density, vertical (vs. horizontal) placement, vibration, heat, and cooling all increase those bit errors. URE is the primary cause of second or third HDD failures in a RAID group.

The potential of a 2nd HDD failure or URE is an extraordinarily high risk based on the industry’s historical experience that when an HDD fails, it is quite common for a second HDD to fail or URE to occur within that same RAID group in a relatively short period of time. It’s easy to understand why. The HDDs are usually the same age. So when one HDD fails it is more likely that another HDD will also fail. During the rebuild, every HDD within the RAID group is stressed as the data is being rebuilt from those drives also increasing the likelihood of another HDD failure or URE. UREs increase over time because of a phenomenon called bit rot. Bit rot is when HDDs acquire latent defects over time from background radiation, wear, dust, etc. Bit rot causes data rebuilds to fail. Many storage systems have implemented some variation background scrubbing that reads, verifies, and corrects “bit rot” before it can cause UREs. However, scrubbing takes system resources that are not immune to the capacity and performance dynamics of the HDD. More capacity by necessity requires more time to scrub. Again, something that took minutes or hours in the past is now measured in days, weeks, even months (depending again on system prioritization). This universe is a closed loop system where there are always tradeoffs.

No one likes to lose data, which is why RAID-6 with double parity has become increasingly popular. Double parity prevents data loss in the event of a second HDD failure or URE during a rebuild. The second parity does not come for free. The cost comes from additional overhead that reduces storage system performance as well as less usable storage (RAID-5 typically reduces usable storage by 20 to 25% whereas RAID-6 reduces usable storage by 25 to 35%). It does not stop there. The additional parity reduces storage system write performance. First failed HDD rebuilds have the same or worse impact as RAID-5 on the storage system’s performance. And when that second HDD has to be rebuilt concurrently, storage system performance



plummets even further while rebuild times grow ever longer increasing the risk of a 3rd HDD failure or URE. And should that 3rd HDD in the RAID group fail or suffer a URE, once again data is lost leading some vendors to advocate striped, triple, or tiered parity RAID. Each parity stripe adds additional overhead and consumes more storage resources while only marginally decreasing the risk of data loss.

For IT professionals this parity Band-Aid on top of parity Band-Aid becomes layers of complexity that end up a complicated operational nightmare and with a whole lot less usable storage.

○ **Life of a failed HDD**

Examining the operational process when one or more HDDs fail is like peeling back the layers of an onion in revealing the layered complexity. First, the HDD must be replaced manually (pulled from the storage system) or replaced electronically from an idle pool of HDDs. If electronically replaced, the failed HDD must eventually be manually physically replaced at some point in the future. Regardless of when that HDD is manually replaced, there is a high probability of human error when the wrong HDD is pulled. How often is the wrong HDD replaced? It happens more frequently than expected. When it occurs, there will be data loss in a RAID-5 group or a double HDD rebuild in a RAID-6 or better, crippling system performance.

Once the failed HDD is pulled a chain of ownership must be documented from the time it's removed from the storage system to the point in time when it is either destroyed or securely wiped. The documentation must include:

- What data is on the HDD;
- Who handled that HDD and for how long;
- How the HDD was secured or not secured for every minute it was outside the storage system;
- Who had access to the HDD and for how long even if they did not directly access it;
- And the process in which the HDD was destroyed or reconditioned.

This is a tedious manually intensive task that's a bit less severe if the HDD is encrypted.

But after all of that painstaking effort it turns out that much of it is for naught. This is because the vast preponderance of failed HDDs sent back to the factory for analysis or erasure (somewhere between 67 and 90% per Seagate and WD) report that the HDD is good. In other words, their bench testing found no failure. Lamentably, this discovery happens after the storage system failed the HDD, had the HDD pulled (assuming the service tech did not accidentally pull the wrong HDD causing more issues), reconstructed the data, and documented the chain of ownership. That's a lot of operational headaches for 'no trouble found.'

○ **Stored Data is Scaling Rapidly, Storage Admins are Not**

Whereas stored data is growing at the previously mentioned range of 50 to 62% the number of storage admins managing that data is growing very slowly or not at all. Unlike storage systems, people do not scale. People have hard fixed limits on the number of discreet tasks they can complete in a given period of time. IDC conservatively pegs the total amount of stored digital data in 2020 will be > 35 zettabytes, which is 35 Billion terabytes. The number of storage admins, they peg at being only 40% more than there were in 2009. That's a problem .

The IT coping mechanism of implementing bigger and bigger HDDs is running out of gas. The growth in HDD capacity technology is slowing and no longer doubling year over year. In 2009 the average large capacity drive was 1TB; in 2010 2TB; in 2011 it's looking more like 3TB. That's not a doubling of capacity; it is a 50% increase. The next generation of HDDs is projecting to be 4TB or only a 33% increase. It appears the timeframes between HDD capacity increases is lengthening beyond the current 12-month

Simple 2 PB example today using the 62% growth rate clearly illustrates the increasing failure rate
<i>62% for simplicity's sake since storage then doubles every 18 mos</i>
<i>2PB today utilizing 2TB HDDs = 1,000 HDDs</i>
<i>Using the realistic HDD failure rate from the field studies of 52,560 hours / 1,000 drives calculates an estimated failed HDD every 52.56 hours or every 2 days, 4 hours, 33 min, and 34 seconds</i>
<i>3 yrs grows 2PB into 8PB</i>
<i>Using 4TB HDDs, calculates an est failed HDD every 26.28 hrs or 1 day, 2 hrs, 16 min, 48 secs</i>
<i>3 more yrs grows storage capacity into 32PB</i>
<i>Using 8TB HDDs = 4,000 HDDs; calculates an est HDD failure rate of 1 per 13.14 hrs or 13 hrs, 8 min, 24 secs</i>

window. This means more HDDs to meet application storage capacity growth demands. More HDDs = more:


- HDD failures, subsequent rebuilds, and data loss;
- Rack space and floor space;
- Power and cooling;
- CapEx and OpEx.

HDDs will be failing more quickly than they can possibly be rebuilt with traditional RAID. System performance will rapidly decline dramatically and data loss will be a consistent reality. And it is only getting worse.

○ Typical Workarounds and Why They Fail

Storage admins are not going to sit back and wait for data to be lost if they can help it. They will implement workarounds. The first is to mirror the data. This provides a full copy on another drive or RAID set that can be copied back when a HDD fails. This greatly speeds up rebuilds; however, there is still the risk that if a second HDD fails that it could be the mirror and data loss.


That leads to another workaround implementation of “Snapshots”. Snapshots provide a point-in-time copy of the data that can be copied directly. The frequency of snapshots is dependent on the amount of data that can be lost in the event of a failure. Copy of the data is misleading. Snapshots today are



primarily copy-on-write (COW) or redirect-on-write (ROW). Both types do not actually copy the data on snapshot. They copy the metadata and point to the original data until the data changes. This enables snapshots to be almost instantaneous while looking and feeling and acting like a real copy, but it really is not. For copy-on-write as the data changes, it copies the original data that's

being changed to a new volume, but not until the data changes. For redirect-on-write, only the changes are copied and it's copied on the same volume. Neither form of snapshot can guarantee a copy of data that is lost.

The next layer of workaround is to replicate the data to another volume or storage system periodically. The most common form is to replicate one or more of the snapshots. Replication requires that the data actually be copied, not just the pointers. Which finally provides a decent amount of protection against data loss, except at a very steep cost.



Each layer of protection reduces the amount of usable data within the storage system or systems requiring more HDDs, racks, infrastructure, storage systems, network switches, cables, connectors, floor space, power, cooling and cost. And in the end, there is still too much risk of data loss that leads to the next line of defense, backup, recovery, and restore (BARR) technology.

BARR backs up the data either through the application servers, desktops, and laptops creating the data, or directly from NAS systems via NDMP. BARR is the last line of defense allowing data to be recovered and restored from a backup. This type of recovery and restore is manually intensive and quite time consuming. Additionally, if it is stored on tape, historical data suggests there is still a 30% chance that the data will never be recovered or restored.

○ What About SSDs as an Alternative to HDDs?

The advantage of Flash SSDs is performance. Random IOPS and throughput are multiple orders of magnitude faster than HDDs. IOPS are measured in tens of thousands of random IOPS vs. hundreds for HDDs. The downside to Flash SSDs is cost. If the issue is performance, SSDs are a solid solution. If the issue is capacity, as it is rapidly becoming for unstructured data, rich media, etc., SSDs are much too costly to be a capacity play. And Flash SSD reliability is still being determined. Preliminary data (and there is not

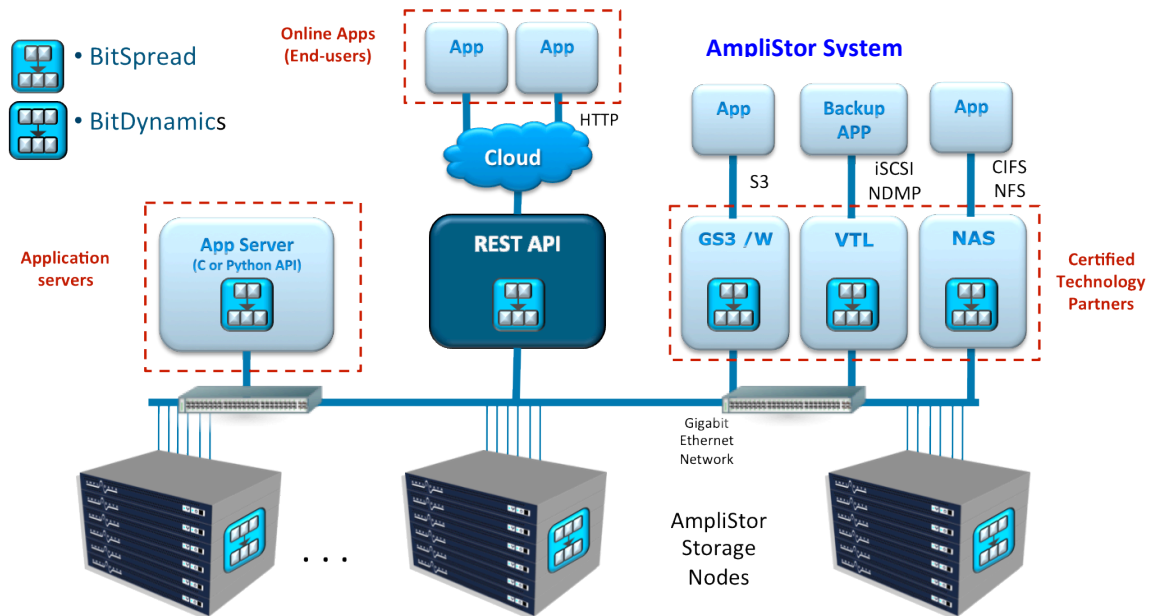
enough data to be definitive) indicates a similar reliability to HDDs, which means similar failures and rebuilds. The good news is that rebuilds are measurably faster. The bad news is that organizations deploying SSDs tend to use more high capacity HDDs to make up for the capacity gap once again increasing data loss issues.

There has to be a better way to eliminate or at least mitigate this accelerating shortening of meantime to data loss (MTDL), extensive recoveries and restores, and un-recoverable data. Fortunately, there is from Amplidata known as AmpliStor.

Stopping the Impending RAID Catastrophe at its Root

AmpliStor is an object storage system from Amplidata (www.amplistor.com). Object storage combines data with rich metadata in order to preserve information about both the context and the content of the data. By storing the data, metadata, and index together as a single object eliminates scaling bottlenecks and makes object storage far more scalable than NAS or SAN storage.

What makes AmpliStor different from most other storage systems (NAS, SAN, or object) is that it does not use RAID or require any form of replication. It utilizes a completely different data protection technology called Erasure Codes. Erasure codes (a.k.a. forward error correction code or FEC) efficiently stores and protects data. Erasure codes transform a data object (files or blocks) into a longer message (typically a linear equation) with “*n*” data blobs such that the original object can be recovered from a subset of the *n* data blobs.



AmpliStor leverages Erasure Code expertise to deliver two unique technologies: BitSpread and BitDynamics in clustered nodes called storage bricks.

○ BitSpread

BitSpread distributes and stores data redundantly across a large number of disks locally or geographically distributed. Each data object is divided in a number of chunks. These chunks are converted to a larger number of check blocks. BitSpread reads in real-time the original data object from a subset of check blocks as soon as the designated sufficient numbers of them are retrieved. And it doesn't matter the order of retrieval, or which check blocks were retrieved, or even from which of the disks or nodes the data was retrieve from, it is still readable in real-time. A good and still imperfect analogy would be a Sudoku puzzle. When a sufficient number of fields are filled, the remaining fields can be recalculated without having to read all data. The beauty of this technology is its failure tolerance flexibility is user determined by data set.

- The minimum number of disks required in the data spread
- The number of simultaneous failures in this spread that data needs to be able to survive

- The geographical spread rules

AmpliStor leverages its built-in expertise to optimize data availability, performance, and accessibility based on these parameters. The system selects which disks data is stored; optimal data spread across the available disks and their locations, while minimizing the negative impact of a component failure. This safeguards data availability. When there is a HDD failure, AmpliStor automatically reconfigures itself to store new incoming data according to an alternative spread of available disks, as defined in the users data set policies.

BitSpread is significantly more HDD failure resistant than any form of RAID. When BitSpread is set up to tolerate up through four (4) HDD failures it is 10,000 times less likely to lose data than RAID-6. That's four orders of magnitude better. HDD failures are never urgent. Even if a technician accidentally pulls the wrong HDD during replacement, it is a non-event.

BitSpread is in fact so resilient to data loss that Amplidata calls it "Unbreakable Storage".

○ BitDynamics

BitDynamics is a piece of software or agent that runs on each and every AmpliStor node (<http://www.amplistor.com/Storage+Node>). It enables AmpliStor scalability from terabytes to petabytes to exabytes even zettabytes. BitDynamics monitors the available storage space on each node. As HDDs are consumed with data, the BitSpread agent will automatically generate new spreads on disks with available capacity, transparently, without requiring any virtual disk reconfiguration.

BitDynamics has unique advanced automated out-of-band storage management functions. It will continuously assess stored data object integrity verification through CRC checksums. When data becomes corrupted from a write error, bit rot or tampering, the BitDynamics agent will detect that error and proactively correct it before it would become an issue to the user. When there is a disk or node failure, the agents on the storage nodes will generate additional data blocks replacing lost or damaged data. It does this processing completely out-of-band causing no impact on system or application performance. Rebuilding the data is not urgent as long as the data object has not exceeded the maximum number of failures. And because multiple BitDynamics agents work in parallel to heal the data, the heal time is lightening quick and far shorter than what comes from RAID based system.

BitDynamics additionally removes obsolete data such as is left when a snapshot or clone has been deleted. As with other BitDynamics functions, this scrubbing is performed out-of-band on the storage nodes having no impact of users system performance. BitDynamics optimizes I/O performance towards the BitLog clients by efficiently organizing the data blocks.

The key is that BitDynamics keeps the AmpliStor storage system healthy and optimized without any manual human interventions.

Conclusion

RAID is becoming a serious and painful problem for IT organizations. Data loss events are growing at an alarming rate. Layers of complexity are similar to putting Band-Aids on a hemorrhage. It will fail and large amounts of data will be lost. The frequency of these events and amounts of un-recoverable data are growing faster than human beings can manage.

AmpliStor Unbreakable Storage is a new scalable resilient storage paradigm that solves this problem. It reduces the risk of data loss to such a low probability (.0073% or once every 14,000 years) that for all intents and purposes it's eliminated. AmpliStor statistically provides an unheard of ten nines of data resilience and durability.

The problem is difficult. AmpliStor makes the solution simple.

For more detailed information, please contact:

info@amplidata.com or go to www.amplidata.com.

About the author: Marc Staimer is the founder, senior analyst, and CDS of Dragon Slayer Consulting in Beaverton, OR. The consulting practice of 13 years has focused in the areas of strategic planning, product development, and market development. With over 31 years of marketing, sales and business experience in infrastructure, storage, server, software, and virtualization, he's considered one of the industry's leading experts. Marc can be reached at marcstaimer@mac.com.

Appendix A: RAID Set Definitions			
RAID-Set	Description	# HDDs protect	Standard or Proprietary
RAID-0	Data striped across HDDs for maximum write performance. Unlike other RAID sets, there is no actual data protection.	0	Std
RAID-1	Synchronously mirrors all data from each HDD to an exact duplicate HDD. No data is lost if either HDD has a fault or failure. Typically the highest-performing RAID level at the expense of lower usable capacity.	1	Std
RAID-2	Data protected by ECC (error correcting codes similar to those found in DRAM). Parity HDD requirements proportional to the log of HDD number. Somewhat inflexible & considerably less efficient than RAID-5 or RAID-6 with less performance & reliability. Not very popular.	1	Std
RAID-3	Data is protected against the failure of any HDD in a group of N+. This is similar to RAID-5; however, the blocks are spread across the HDDs. The parity is bitwise vs. RAID-5 block. Parity resides on a single HDD rather than being distributed between all disks. Random write performance is quite poor with random read performance fair at best.	1	Std
RAID-3 Double Parity	RAID-3 with a second byte level parity disk. Protects data in the event of a second HDD failure or loss of the parity HDD. Performance is marginally slower than standard RAID-3. However, system performance can degrade noticeably if dual HDDs are rebuilding concurrently.	2	Prop
RAID-4	Similar to RAID-3, it stripes data across many HDDs in blocks instead of RAID-3 bytes. This improves random access performance over RAID-3. Data protection is provided and by a dedicated parity HDD. RAID-4 is also similar to RAID-5 except that instead of distributed parity it uses dedicated parity. The dedicated parity HDD remains a bottleneck, especially for random write performance.	1	Std
RAID-5	Most common form of RAID. Provides RAID-0 similar performance with more economical redundancy. It stripes block data across several HDDs while distributing parity among the HDDs. No single disk is devoted to parity. RAID-5 utilizes HDDs more efficiently providing overlapped read and write operations speeding up small writes in a multiprocessor system, while providing greater usable storage than RAID level 1 or 10 (the redundancy storage penalty is approximately 20% instead of 50%). Data protection comes from parity information that's used to reconstruct data if a drive in the RAID group has a fault or fails. RAID-5 issues include a minimum of 3 and usually 5 HDDs per RAID group, lower performance of the storage system while a HDD is being rebuilt, and potential of total RAID group data loss if a second drive faults or fails during rebuild. Read performance also tends to be lower than other RAID types because parity data is distributed on each HDD.	1	Std
RAID-6	RAID-6 is similar to RAID-5 but includes a second parity scheme distributed across the HDDs of the RAID group with P & Q being the 2 parity groups. Dual parity protects data against loss of a second HDD in the RAID group. RAID-6 tends to have lower storage system performance than RAID-5 and can plummet if dual HDD rebuilds are occurring.	2	Std
RAID-10	RAID 1 striped. Improves write performance closer to RAID-1.	1	Std

RAID-50	RAID 5 striped. Improves write performance closer to RAID-1.	1	Std
RAID-60	RAID 6 striped. Improves write performance closer to RAID-1.	2	Std
RAID-6 EVENODD	EVENODD utilizes only two additional redundant HDDs and consists of simple exclusive-OR computations. The advantage of EVENODD is that it only requires parity hardware, which is typically present in standard RAID-5 controllers. This reduces the number of exclusive-OR operations over the more common Reed-Solomon computations at approximately 50% (based on 15 drives). EVENODD has similar performance issues of RAID-6 when dual HDD rebuilds are occurring.	2	Prop
RAID-DP or Row-diagonal Parity	Stores row parity across the HDDs in a RAID 4 group, the additional parity HDD stores diagonal parity across the HDD in a RDP group. The 2 RDP parity stripes provides data protection in the event of two HDD failures occurring within the RAID group. Performance is nearly equal to single parity RAID-4 or RAID-5. Higher performance than standard RAID-6 (Reed-Solomon) but with similar performance issues when concurrently rebuild 2 HDDs.	2	Prop
RAID-TM or RAID Triple Mirror	RAID-TM provides the very high speed of RAID-1 while providing the high reliability and double HDD fault protection of RAID-6. RAID-TM writes the data simultaneously to three separate HDDs. Even if there are two HDD faults or unrecoverable read errors in the same mirror, the application still has access to its data with no degradation in performance even as the drives are rebuilt.	2	Prop
RAID-X	RAID-X doesn't require a spare HDD just spare capacity on existing HDDs being used throughout the entire storage system to hold data. The objects can be mirrored between any two types of HDDs so there's no need to match HDD type (size and rotational speed), which is a definite advantage over other RAID. Rebuild performance is extremely fast since data is mirrored. This is a variation of RAID-1 or RAID-10 but with the added protection of random distribution. A second drive failure of the same data can result in data loss, which can only be mitigated with additional mirroring. Usable storage can be severely restricted depending on the number of mirrors (minimally half).	2 plus, depending on mirror #	Prop
RAID-Z	RAID-Z comes from ZFS and is primarily RAID 5 without the write hole vulnerability. Write-hole vulnerability is when there is a system failure while there are active writes, the parity of a stripe may become inconsistent with the data. If this is not detected and repaired before a disk or block fails, data loss may ensue as incorrect parity will be used to reconstruct the missing block in that stripe. RAID-Z2 is RAID 6 equivalent; RAID-Z3 is RAID 6 triple parity, and any nested combination of those like 1+0.	1, 2, or 3	Prop
RAID: Tiered Parity	RAID: Tiered parity has 3 components: horizontal parity stripes data across an array of disks and calculates one or more parity blocks across that stripe analogous to other multi-parity RAID; vertical parity protects against disk media sector issues; network parity provides end-to-end parity from initiator to storage.	1, 2, or 3	Prop