

BUSINESS WHITE PAPER

The Storage Reliability Imperative

Evolving Requirements in the Era of Digital Transformation

September 2023

In this era of digital transformation, businesses are increasingly data-driven. Enterprises are capturing, storing, protecting, and analyzing more data than ever before in workflows that are critically dependent on the availability of information technology (IT) infrastructure and the integrity of data. Newer processing technologies like accelerated compute are being combined with artificial intelligence (AI) and machine learning (ML) to use this data to optimize product design, manufacturing efficiencies, market intelligence, revenue generation, and customer experience. When IT infrastructure can't meet the demanding requirements for 24/7 availability, all aspects of a business—from development, production and inventory management to business development, sales and the quality of customer interactions—are potentially affected.

Enterprise storage solutions are right at the heart of this challenge to serve data up as needed to drive key business processes while at the same time ensuring that it is both accurate and recoverable. Today's storage systems have to be very different than they have been in the past. While performance and availability have always been critical purchase criteria, today's systems must meet those requirements while operating at a whole new level of scale. Additionally, high data growth and the fast pace of business require an on-demand agility that was unheard of in the past, and multi-petabyte data stores growing within a few short years into the tens of petabytes are the norm for many enterprises.

To meet the need for heightened reliability in enterprise storage systems, enterprises should consider carefully the interconnection between performance, availability, data durability and resiliency. This white paper will discuss how those topics are inter related in enterprise storage, providing recommendations to help businesses configure high performance, highly available storage infrastructure that can meet today's evolving requirements.

Evolving Storage Reliability Requirements

When it comes to enterprise storage infrastructure, the concept of reliability has implications for performance, availability, data durability, and resiliency. At its core, reliability implies that a system will work as expected across both normal and failure mode operation, meeting very specific performance, availability, and data durability requirements. Performance refers to expected latencies, input/output operations per second (IOPS), and throughput and how that may vary across different failure modes. Availability refers to a storage system's ability to effectively service data requests and is comprised of a combination of both storage system uptime and data availability. Durability focuses on ensuring that the data is not lost or corrupted by normal, failure, defensive, or recovery mode operations. Resiliency describes a system's ability to self-heal, recover and continue operating predictably in the wake of failures, outages, security incidents and other unexpected occurrences.



Performance

While performance during normal operation can be well understood, it can be unpredictable in the wake of failures. Controller failures and disk rebuilds provide two good examples. In an active/passive controller architecture, one controller in a pair handles the load during normal operation. When a failure occurs, the second controller takes over and there is no impact to performance. In an active/active controller architecture, the load is split across the controllers in normal operation. When a failure occurs, the remaining controller may not be able to fully meet performance service level agreements (SLAs) since it will now be running its own load plus whatever was transferred over from the failed controller. The performance of the system in failure mode is very dependent on the workload and can vary significantly.

Performance can also be impacted during disk rebuilds. When a drive fails, data will be pulled from different devices in the system to rebuild the data on a new device. Given a defined data transfer rate, larger devices can take a lot longer to rebuild than smaller ones. This is a particular concern with HDDs which have a data transfer rate for serial writes (<300MB/sec) that is 20 times slower than a commodity off the shelf (COTS) solid-state disk (<6000MB/sec). A full data rebuild on a 10TB HDD would take over 55 hours while it would take less than three hours on a 10TB SSD.

During a disk rebuild, application performance on a system can be impacted with the degree of impact determined by

Availability

For critical enterprise workloads, “five nines plus” (99.999%) availability, which translates to roughly around five minutes of downtime per year, is a very common target requirement. It is interesting to note that public cloud storage services will generally not provide an SLA greater than “four nines”. The ability to meet very high availability requirements like this is complicated by the fact that today there is much more diversity in the fault models that must be taken into account at much higher levels of scale. In addition to potential data corruption in normal operation and component failures, enterprises must also deal with the fallout from ransomware attacks and power failures. The wider diversity

of fault models has demanded innovations in the way that storage infrastructure is architected and implemented. the load on the system and the amount of data that must be moved. When rebuild times are long, enterprises are concerned not only with the duration of the performance impact but also with the possibility that a second device failure might occur during the rebuild. This latter concern is greater with HDDs [because of their higher annual failure rates (AFRs) and lower component-level reliability relative to flash] and becomes even more of an issue as HDDs age (because their AFRs increase in years 4 and 5). Multi parity data protection [like RAID 6 or multi parity erasure coding (EC)] has become popular to address this concern (since they can sustain multiple simultaneous device failures without data loss). It is also why enterprises building large systems with HDDs often opt for smaller device sizes.

From this discussion, it should be clear that performance and availability are linked. To make systems simultaneously high performance and highly available, focus not only on performance during normal operation but also performance in the wake of failures. Systems that can sustain failures without performance degradation drive better value because of the predictability of their operation. N+1 designs like active/passive controller architectures deliver that predictability across a range of failure modes, resulting in a better experience for both end users and IT. If you doubt that, consider this: if a system’s performance drops by 20-30% in the event of a failure, is it still appropriate to consider it available?

of fault models has demanded innovations in the way that storage infrastructure is architected and implemented.

One approach to consider is defining availability zones across system, rack, pod, data center, or power grid boundaries to ensure that a failure at any one of these levels does not adversely impact business operations. This leverages redundancy that already exists within your IT infrastructure for resiliency and can provide more cost-effective approaches to deal with catastrophic failures that go far beyond just a system failure. There is also value in winnowing down the number of fault responses required to comprehensively handle



any types of failures that can occur. As an example, defining less interdependent availability zones that span multiple pods covers many potential failures more simply and less expensively than coming up with a resiliency strategy that operates only at the enclosure level. Fewer recovery workflows also mean there is less testing that must occur as those workflows evolve.

When noting vendor descriptions of achievable availability, understand whether a particular rating for a system includes scheduled downtime exceptions or not. Many storage systems today still require some downtime for device-level firmware upgrades, on-disk data format changes, certain component replacements and multi-generational technology refresh. A system rated at “five nines plus” by the vendor may in fact provide far less availability because of scheduled maintenance downtime that is not incorporated into the calculation. Availability also often depends on how systems are configured so storage managers need to ensure that they use the appropriate system features correctly to hit their availability targets. Just because a storage system *can* be configured to support “five nines plus” doesn’t mean that it always *will*.

Data Durability

Data durability refers to the accuracy of data during the process of storing, protecting, and using it. A variety of well-known techniques such as cyclical redundancy checks (CRC), which use checksums and/or parity to support single error correction/double error detection, have been in use for years to ensure that data is transferred and stored correctly and that its integrity is maintained in use. Approaches like RAID and erasure coding (EC), both of which inject parity and impose some data redundancy as data is distributed and initially stored, ensure that data will be protected during normal operation and despite component failures within a system.

Different RAID and EC algorithms can provide different levels of protection from data loss and/or data corruption in the wake of even multiple simultaneous component failures within a system. Algorithms with more parity bits can provide higher levels of resiliency, but this will also incur more capacity overhead for the protection. Performance and recovery time (i.e. data rebuild) considerations must also be taken into account, since different approaches can have impacts with different risk

A painful system upgrade to next generation technology has been an accepted part of the enterprise storage life cycle for decades. In the “forklift” upgrade a customer brings in an entirely new system, is forced to re-buy their existing storage capacity plus any additional required capacity, re-license their storage software, migrate their data, and then re-provision this new storage to their servers. Multi-generational technology upgrades are expensive, time-consuming and risky.

The ability to migrate data across external networks to the new system while applications continue to run helps to defray some of the pain of the upgrade, but application performance can be impacted during the migration (which may take months depending on how much data must be moved) and it does nothing to reduce the cost. In practice, forklift upgrades often do impact service availability, a factor which must be taken into account since the typical enterprise storage system must be upgraded every four or five years. And they clearly impact administrative productivity, since time spent managing the upgrade means less time available for other management and more strategic activities.

profiles. As an example, fifteen years ago many distributed scale-out software systems were experimenting with EC across different data centers to provide site-level redundancy, but today enterprises seem to prefer to combine local distribution of data with replication to a remote site. The latter approach can provide the same level of resiliency with significantly less performance impact (albeit at generally a slightly higher cost).

All storage vendors provide some form of on-disk data protection (i.e. RAID or EC) to protect data from storage device failures, and many provide several options from which to choose to give storage managers the freedom to choose the one(s) that best meet their needs for various workloads and data sets. Some EC implementations evolve as additional storage devices are added to provide the same level of required resiliency with less capacity overhead. When crafting and deploying a highly durable data environment, evaluate vendor approaches for performance in both normal and failure mode operation, depth of resiliency, flexibility, ability to scale and capacity overhead associated with their approaches.



Significant experimentation has been used to balance the trade-off between resiliency, performance and capacity overhead in on-disk data protection over the last two decades, and there can be large differences between different vendor RAID and/or EC implementations. Given that more and more systems are growing to multi-PB scale and beyond, the ability

of a proprietary algorithm to reduce capacity overhead by only 5-10% without undue performance and/or resiliency impacts can mean a large savings in terms of hardware and energy costs. Larger systems with more individual storage devices provide additional options for protection algorithms by allowing data to be more widely distributed.

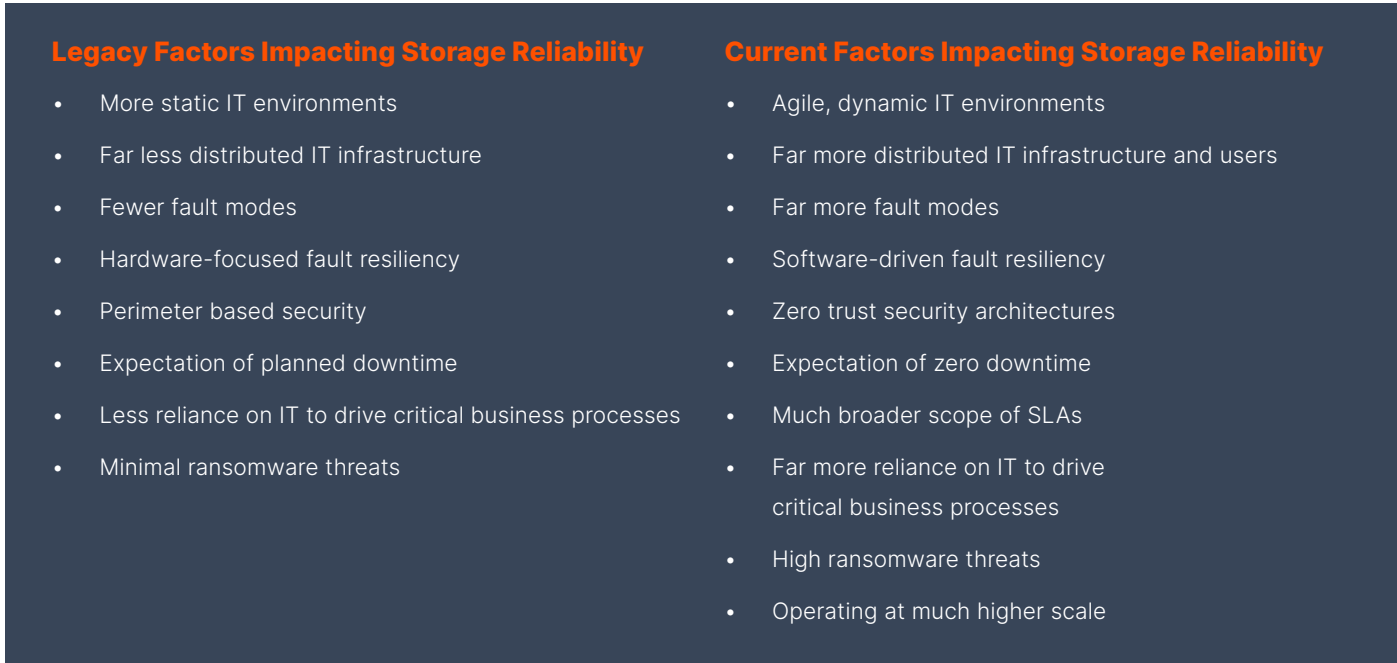


FIGURE 1 The era of digital transformation changes a number of factors impacting storage reliability.

Resiliency

Regardless of how a system is constructed, it is inevitable that failures will occur. High availability storage systems need to be designed to continue to serve accurate data despite failures. Recovery from failures also needs to be rapid and non-disruptive, two additional requirements driving storage system design in today’s business climate. Running two identical systems in parallel was an early approach to achieving high availability, but a variety of innovations over the last two or three decades are now able to meet high availability requirements in far more cost-effective ways. Some of these include dual-ported storage devices, hot pluggable components, multi-pathing with transparent failover, N+1 designs, stateless storage controllers, more granular fault isolation, broader data distribution, and evolution towards more software-defined availability management.

Architectural complexity is a key detractor from system-level reliability. Given a target combination of performance and capacity, the system with fewer components will be simpler, more reliable, easier to deploy and manage, and will likely consume less energy and floor space. Contrasting dual with multi-controller system architectures can help to illuminate this point. The contention that multi-controller systems deliver higher availability (the primary claim that vendors offer) is based on the argument that a multi-controller system can sustain more simultaneous controller failures while still continuing to operate. While theoretically that is true, multi-controller systems require more components, are more complex and pose many more fault modes that must be thoroughly tested with each new release.



The experience of many enterprises is that the higher reliability of these more complex multi-controller systems in practice is in fact *not* true. Dual controller systems have fewer components, are less complex, and have far fewer fault modes that must be tested with new releases. In fact, one interesting innovation is to use the same “failover” process for both planned upgrades and failures to simplify storage operating system code and make that common process more reliable than the different processes that most storage systems use to handle failures and upgrades. As long as a dual controller system delivers the required performance and availability in practice, it will be a simpler, more reliable system with fewer components and very likely more comprehensively tested failure behaviors. Simpler, more reliable systems mean less kit needs to be shipped (driving lower shipping, energy and floor space costs), less unnecessary redundancy to meet target, and availability SLAs (driving lower system costs), and fewer service visits.

Storage managers need to determine the levels of resiliency they require—component, system, rack, pod, data center, power grid, etc. For each level, the goal is to deploy a storage system that could continue to serve accurate data despite a failure in any one of these levels. Many storage vendors provide system architectures that use software to transparently sustain and recover from failures in individual components like storage devices, controllers, network cards, power supplies, fans and cables. Software features like snapshots and replication give customers the tools they need to address larger fault domains like enclosures, entire systems, racks, pods, data centers, and even power grids. Systems designed for easier serviceability simplify any maintenance operations required when components do need to be replaced and/or upgraded.

Designs that reduce human factors during recovery have a positive impact on overall resiliency. Humans are the least reliable “component” in a system, and less human involvement means less opportunity that a technician will issue the wrong command, pull the wrong device, or pull or jostle the wrong cable.

While initially focused more on hardware redundancies, the storage industry’s approach to resiliency has changed in two ways. First, we have moved away from a focus on system-level availability to one that centers on service availability. System-level availability is still a critical input, but modern software architectures provide options to ensure the availability of various services, regardless of whether those are compute, storage or application based, that offer significantly more agility. And second, fault detection, isolation and recovery has become much more software-based and much faster than in the past, providing more agility than was possible with less cost-efficient hardware-based approaches.

With data growing at 30% to 40% per year¹ and most enterprises dealing with at least multi-PB data sets, storage infrastructures tend to have many more devices and are larger than they have been in the past. A system with more components can be less reliable, and vendors building large storage systems have to ensure they have the resiliency to meet higher availability requirements when operating at scale. Flash-based storage devices demonstrate both better device-level reliability and higher densities than hard disk drives (HDDs), and they can transfer data much faster to reduce data rebuild times (a factor which is generally a concern as enterprises consider moving to larger device sizes). The ability to use larger storage device sizes without undue impacts on performance and recovery time results in systems with fewer components and less supporting infrastructure (controllers, enclosures, power supplies, fans, switching infrastructure, etc.) that have a lower manufacturing carbon dioxide equivalent (CO₂e), lower shipping costs, and use less energy, rack and floor space during production operation. While highly resilient storage system architectures can be built using HDDs, it can be far simpler to build them using larger flash-based storage devices that are both more reliable and have far faster data rebuild times.



Pure Storage: All-Flash, Enterprise-Class Storage Infrastructure

Pure Storage® is a \$3 billion enterprise storage vendor with over 12,000 customers. In 2012, we started the move towards all-flash arrays that has come to dominate external storage systems revenues for primary storage workloads. We provide the industry's most efficient all-flash storage solutions based on watts per terabyte (watts/TB) and storage density (TB/U) metrics. We offer two types of unified storage: a scale-up, unified block and file storage system called [FlashArray™](#) and a scale-out, unified file and object storage system called [FlashBlade®](#).

Based on performance we've tracked over our more than a decade in existence using Pure1® (our AI/ML-driven system monitoring and management platform), our systems deliver "five nines plus" availability in single node configurations. Stretch cluster configurations that can deliver "six nines plus" availability are also available. And while we provide all-flash arrays for primary storage workloads, our quad level cell (QLC) NAND flash-based Pure//E™ family of systems can cost-effectively replace all-HDD platforms used in secondary workloads with a raw cost per gigabyte between \$0.15 and \$0.20/GB (including three years of 24/7 support).

Our design philosophy ensures that systems deliver the same performance in both normal and failure mode operation and do not require any downtime (planned or otherwise) during failures, upgrades, system expansion, or multi-generational technology refresh (an approach we refer to as [Evergreen® Storage](#)). The simplicity of our systems (which can be up to 85% more energy efficient and take up to 95% less rack space than competitive systems) is another major contributor to our extreme reliability. We do not require forklift upgrades to move to next generation storage technology while still allowing our customers to take advantage of the latest advances throughout a proven ten-year life cycle. A proof point for that is that 97% of all the systems that Pure Storage has ever sold are still in production operation and feature the latest in storage technologies. An additional proof point is that we have the highest independently validated Net Promoter Score in the industry at 81.4 (on a scale of -100 to +100), a trusted indicator of the excellent customer experience we provide to our customers.

Learn More

If you're interested in moving to highly efficient, all-flash storage that meets the high reliability requirements of the digital era, we'd [like to meet](#) to discuss how we can help move your business forward. We have a compelling reliability story with our storage solutions which, while beyond the scope of this white paper, we'd be happy to discuss with you.

[Contact Us](#)

1 Evaluating Power and Cost-Efficiency Considerations in Enterprise Storage Infrastructure (IDC #US49354722, June 2022)

purestorage.com

800.379.PURE

