

Modern storage infrastructure needs to provide performance, scalability, reliability, and efficiency to maximize data value and minimize time to insights for organizations in the digital business era.

Balance Storage Performance, Resilience, and Efficiency to Drive Digital Business Outcomes

May 2023

Written by: Dave Pearson, Research Vice President, Infrastructure Systems, Platforms, and Technologies Group

Introduction

Digital transformation (DX) has given way to the advent of digital business. Businesses have increasingly embraced digital tools, platforms, and strategies to optimize their operations, improve customer experiences, and drive growth. IDC research has shown that despite uncertain economic conditions, organizations intend to increase spending on datacenter technologies including storage systems, converged and hyperconverged infrastructure, and storage software.

Workloads continue to proliferate as new applications are developed and deployed at a staggering rate. Additionally, workloads such as high-performance computing (HPC) and artificial intelligence (AI) are entering mainstream enterprise environments and are increasingly demanding — in terms of both performance and capacity. IDC's Global DataSphere tells us that enterprises generated nearly 60% of the over 100 zettabytes of data created last year — a figure that will grow to over 260 zettabytes by 2027. Unstructured data — such as data in active archives, machine learning environments, and media repositories — will account for 80% of that growth. This unstructured data contains massive volumes of file and object data requiring increases in scalability and efficiency to derive maximum value from it.

As business processes become more digital and applications drive much of an organization's competitive advantage, the value of data increases commensurately. Downtime is unacceptable; customers, employees, and partners depend on the availability of an organization's data and the resilience of an organization's infrastructure. Protecting that data from internal and external threats, user error, equipment failure, and other unplanned disruptions is critical to ensure the digital resiliency of an organization and a crucial factor in driving business outcomes from digital assets.

AT A GLANCE

WHAT'S IMPORTANT

- » Data protection requirements, as well as the capacity and performance demands of modern workloads, have pushed storage buyers to consider more modern architectures to support business objectives in the digital business era.
- » Evaluating storage platform capabilities is critical to providing applications with the infrastructure required to derive deep insights in a timely manner. Hardware acceleration, scalability, software capabilities, increased availability and resilience, and efficiency are all key criteria for vendor selection.

Definitions

- » **GPU-accelerated storage:** Systems that offload computationally intensive tasks such as erasure coding, encryption, and other data processing from traditional CPUs and ASICs to GPUs in order to take advantage of the capabilities of massively multicore processing
- » **Erasure coding:** A parity-based data protection technique in which data is broken into encoded fragments that can be distributed anywhere within an organization's storage environment, providing robust data protection with rapid recovery but also with increased processing costs

GPU-Accelerated Architectures

Storage performance is consistently a top criterion for vendor selection in IDC buyer perception studies. Storage buyers are seeking even incremental performance gains to keep up with greater demands from current-generation and next-generation workloads in enterprise datacenters, increased consolidation, and greater application density and capacity requirements for storage systems. For certain computationally intensive tasks, GPUs offer performance that is an order of magnitude better than that of CPUs and at prices equivalent to those of current CPUs.

GPU-accelerated storage architectures can improve erasure coding performance, especially in petabyte-scale and block storage environments where computational overhead makes these environments less performant and attractive to buyers. GPU-accelerated storage can provide the capabilities required for the mixed block, file, and object environments of the next generation of enterprise workloads, including denser consolidation, the use of large capacity low-cost HDDs, lower power and cooling requirements, and higher levels of resiliency, availability, and data protection while ameliorating some of the cost, complexity, and computing requirements of CPU-only erasure coding implementations.

Scalability

The explosion of data growth, the proliferation of new applications and new sources of data in the enterprise, and the need to control datacenter complexity, costs, and management overhead have made scalability a critical criterion for storage system selection among buyers. Both scale-up and scale-out capabilities are sought after for a variety of reasons.

Increasing performance through scale-up capabilities can help organizations increase consolidation, as well as efficiency and utilization, while reducing complexity and administration overhead in the datacenter. Scale-out systems can increase capacity, availability, and flexibility while providing cost-effective performance.

The key to effective scaling, however, is to avoid disruptions in operation. Adding capacity or replacing existing media with newer or higher-density drives or modules, as well as acquiring new capabilities through next-generation hardware, software, or OS patching or upgrades, must be managed nondisruptively. Scheduled downtime is still downtime, and operational efficiency depends on reducing any and every instance.

Consolidation and Simplicity

The number of new applications entering the enterprise and the performance attributes associated with each of them can lead organizations to seek best-of-breed solutions on a "per app" basis. This can lead to silos of data, communication, and innovation. Modern storage architectures are able to address more of the spectrum of workload requirements, allowing companies to provide denser consolidation of these applications on suitable storage platforms.

This has multiple positive effects: fewer technical and operational silos in the datacenter; easier management and reduced overhead for storage administrators; reduced complexity impacting security, compliance, and governance teams; and a smaller infrastructure footprint, which can improve efficiency and utilization while lowering space requirements and power and cooling costs. Care must be taken to avoid impacts on applications, from "noisy neighbor" problems to attribute mismatch (e.g., a policy requirement for a particular workload that impacts others on the platform due to a lack of granular management may negatively impact workload performance or capacity demands).

Rearchitecting for Resiliency

As the events of the past few years drove home lessons in how digital businesses can thrive during economic and geopolitical upheaval and uncertainty, organizations came to realize that digital resilience was at least as important as innovation during this period. According to IDC surveys, preventing disruptions, as well as recovering quickly and gracefully from them when they do occur, has become a top business priority.

Resiliency needs to be a priority throughout the entire technology stack, and efficient, GPU-driven erasure coding can provide increased data integrity, availability, and resilience through a greater resistance to drive failures than older technologies. Maximizing the number of drives that can safely fail before data loss while minimizing the performance hit due to drive failure is a key benefit. Erasure coding improves data resiliency by creating redundancy that can withstand multiple disk failures or data corruption. This makes it ideal for large-scale data storage systems, where the risk of data loss due to hardware failures is high.

Erasure coding can also be more cost effective and can provide faster data reconstruction times than traditional data protection methods. This is because only a subset of the parity blocks needs to be read to reconstruct the lost or corrupted data. Operational resilience is improved at the same time, as the urgency of drive replacements is reduced, even in the event of multiple drive failures.

Pushing Storage Efficiency for ROI, ESG, and Innovation

Limited IT budgets and economic uncertainty combined with high data growth rates have already made improving storage efficiency a key consideration as IT infrastructures evolve. Doing more with less is the mantra of nearly every IT shop, but there are even wider considerations today. Environmental concerns — especially with regard to the cost and availability of power, but extending through the entire life cycle of hardware deployments — are growing among organizations in many geographies.

Consolidating workloads, implementing highly efficient infrastructure to improve utilization rates, and using the latest large capacity storage devices, whether spinning disk or flash based, are all ways to lower the footprint of storage infrastructure in the datacenter as well as reduce power and cooling requirements. While these initiatives are tied to positive environmental, social, and governance (ESG) outcomes, they also increase operational efficiency by removing silos and reducing complexity in storage infrastructure. The increased efficiency compounds the return on these initiatives by reducing skills and staff required, lowering administrative overhead, and allowing administrators and resources to return to value-added innovative activities for the business.

Considerations for Storage Buyers

The proliferation of storage vendors, technologies, and deployment methodologies is a double-edged sword for buyers today — options exist for every workload, but choosing the appropriate technology during the planning process is key because upgrades and migrations can be complex and costly and locking into the wrong deployment modality can be an expensive mistake for enterprises.

Examining total cost of ownership (TCO) for a new storage solution, from acquisition costs to operating costs to the availability of skills and resources, is key for buyers. Matching application requirements while removing complexity from the storage environment can pay dividends as operational requirements, resourcing, and migrating workloads to appropriate infrastructure are all risks to TCO for enterprises. Data protection, data growth, and higher-performance storage are the top budget drivers for this year and next (see Table 1).

TABLE 1: **Top Budget Drivers for New Storage**

	% of Respondents
Greater data protection needs	51
Significant data growth	48
Need for higher-performance storage	47
New storage needs for existing applications/workloads	41
New applications/workloads	38
Storage infrastructure for edge workloads	36
General refresh	16

n = 600

Source: IDC's 4Q22 Infrastructure (Storage) Worldwide Survey

Conclusion

Better decision making in the digital business era requires better access to data, the tools and applications required to extract insight from the data, and the infrastructure to enable those applications and speed time to insight. We continue to see the demand for storage grow along the workload spectrum, from the highest performance, lowest latency, and greatest throughput to the massive scale of unstructured file and object data. Satisfying the workload requirements and storage demands of the modern digital enterprise requires similarly modern infrastructure planning, procurement, and operations.

Buyers should seek the following capabilities from modern storage architectures:

» Performance

- Throughput, latency, and IOPS characteristics that match application requirements
- Scalability to meet data growth expectations
- Nondisruptive upgrades, migrations, and operations

» Resilience

- High reliability and graceful recovery from failure
- Data protection and data loss prevention
- Consistency and data integrity capabilities transparent to the user experience

» Efficiency

- Dense consolidation
- High utilization
- Cost-effective media
- Simplified management and low administration overhead

Satisfying the workload requirements and storage demands of the modern digital enterprise requires similarly modern infrastructure planning, procurement, and operations.

About the Analyst



Dave Pearson, Research Vice President, Infrastructure Systems, Platforms, and Technologies Group

Dave Pearson is Research Vice President for Storage and Converged Systems practice within IDC's worldwide infrastructure research organization. He also oversees IDC Canada's Infrastructure Solutions research practice. He and his team are responsible for IDC's coverage of storage, integrated, hyperconverged, and composable systems and platforms. This includes storage for performance-intensive use cases like high-performance computing, artificial intelligence, and analytics.

MESSAGE FROM THE SPONSOR

About Nyriad

Unstructured object and file data continues to grow at a breakneck pace driven by demanding applications and workloads across segments such as media and entertainment, backup & recovery, HPC, active archive and rich media video and imaging. These applications & workloads continue to proliferate, becoming increasingly more prevalent in mainstream enterprise environments.

To grow, adapt and stay competitive, businesses must, at the same rate, adopt and deploy innovative cost-effective storage solutions to support performance, resiliency, and productivity requirements, while bolstering sustainability goals.

Nyriad[®] has unleashed the power and speed of GPUs combined with CPUs to revolutionize how data is stored, accessed, and managed. Nyriad's UltraIO[™] software-defined data storage system enables massive amounts of data to be safely and intelligently managed by a single storage platform with a low TCO, while delivering exceptional performance, resiliency, and efficiency. With Nyriad, organizations can enhance agility, accelerate innovation, and gain the competitive advantage necessary to drive business growth and success.

To learn more, please visit: <https://www.nyriad.io/>.

 **IDC Custom Solutions**

The content in this paper was adapted from existing IDC research published on www.idc.com.

This publication was produced by IDC Custom Solutions. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2023 IDC. Reproduction without written permission is completely forbidden.

IDC Research, Inc.
140 Kendrick Street
Building B
Needham, MA 02494, USA
T 508.872.8200
F 508.935.4015
Twitter @IDC
idc-insights-community.com
www.idc.com