

Growing use of GenAI, AI, and analytics applications in hybrid cloud and multicloud environments presents opportunities to consider workload consolidation through a unified data storage and management platform.

Unifying Data Storage and Management Offers Potential Advantages for Modern Workloads in Hybrid Cloud Environments

April 2024

Questions posed by: NetApp

Answers by: Carol Sliwa, Research Director, Infrastructure Systems, Platforms, and Technologies Group

Q. What are the most significant trends influencing enterprise requirements for data storage infrastructure today?

A. The growing use of data-intensive analytics, artificial intelligence (AI), and machine learning (ML) applications in connection with digital transformation is having a notable impact on enterprise requirements for data storage infrastructure. Those workloads can hinge on large amounts of structured and unstructured data from a diverse range of sources, including block-, file-, and/or object-based storage systems that, in some cases, span on-premises and public cloud sites. Combining and managing such disparate data sets can be a challenging proposition with storage systems that offer limited support for differing data types and access protocols. Many enterprises need to modernize their data storage infrastructure in tandem with digital transformation efforts involving analytics, AI/ML and, increasingly, generative AI (GenAI) applications that demand flexible data access, high performance, and potentially massive scalability. IDC's *Future Enterprise Resiliency and Spending Survey* shows that organizations with plans to increase IT spending for GenAI in 2024 estimate that 24% will go to dedicated server and storage infrastructure in on-premises, colocation, and remote edge sites, and 21% will be allocated to public cloud infrastructure such as compute and storage cloud services.

A related trend that continues to impact enterprise storage requirements is the shift to hybrid cloud infrastructure that combines on-premises systems and services from one or more cloud providers. IDC's 2023 *IT Infrastructure for Storage and Data Management Survey* showed that 57% of organizations take a hybrid multicloud approach to data storage, and 27% have adopted a hybrid cloud strategy. Many IT departments turned to public cloud infrastructure to speed provisioning as business leaders sought to gain greater insight and value from their data assets through AI and analytics workloads. Organizations also looked to the cloud to offload operations, management, and maintenance with general business applications and ease budget planning by shifting infrastructure from a capital to operational expense model.

However, many learned public cloud storage may not be optimal for all workloads given the data egress charges or performance, security, and compliance requirements, and most enterprises now favor a hybrid cloud or hybrid multicloud approach for data storage infrastructure.

IDC's 2023 *IT Infrastructure for Storage and Data Management Survey* also showed the top driver for storage budget increases was greater data protection needs, as organizations prioritize security to combat the ever-present threat of ransomware and other cyberattacks. The stakes become especially high with sensitive data in use with AI/GenAI/ML workstreams.

Q. How can a unified approach to data storage and management benefit organizations in hybrid cloud environments?

A. Relying on different suppliers in a hybrid cloud or multicloud storage environment may eliminate the potential perils of vendor lock-in, but it can also add a taxing level of complexity for the IT personnel tasked with overseeing the disparate systems and managing the potentially large volumes of data they store. Data silos can be especially problematic with modern analytics, AI, and ML applications that need to draw data from structured and unstructured data sources that may span on-premises and public cloud sites. Consolidating storage where possible can ease administration, increase efficiency, and help organizations avoid costly and cumbersome data migration projects. Unifying storage infrastructure also provides a consistent control plane for data protection, governance, and security to guard against the threat of cyberattacks.

Centralized, web-based consoles facilitate a unified view of storage and data assets that may extend across multiple geographies. Increasingly sophisticated AI-driven management tools with automation capabilities can help staffing-challenged IT organizations improve productivity, simplify provisioning and monitoring, and predict and proactively address technical issues. Software-defined storage with an operating system designed to run on general-purpose servers or certified hardware can reduce costs and ease infrastructure scaling across on-premises datacenters and public cloud sites through a single, unified system. Many software-driven storage products support data access through multiple file-, object-, and/or block-based protocols and interfaces to give organizations the flexibility to use the systems with an expanding range of workloads. Enterprises that want to take additional steps to simplify their infrastructure will find increasing options for on-premises storage as a service that mimics public cloud alternatives, offering storage on demand through a pay-as-you-go operational expense model.

Q. How suitable is unified data storage for hybrid cloud and hybrid multicloud environments?

A. Unified storage has changed over the years and can vary by vendor, but it shares the basic characteristics of storing and managing data through a single system designed to support multiple data formats, structural types, and/or access mechanisms. One evolving direction for modern unified data storage is adapting to the needs of enterprises operating in hybrid cloud and hybrid multicloud environments. Decoupling the operating system and software from proprietary hardware can enable unified data storage to run on general-purpose or certified server hardware in any location, from

on-premises, colocated, and hosted datacenters to remote edge sites and public clouds. Using the same operating system and software can give enterprises a consistent model for data management, protection, security, and governance no matter where the storage and data resides.

Software-defined storage can also ease scaling across hybrid cloud and multicloud environments through the addition of server nodes and enable important data services such as compression, deduplication, snapshots, replication, multi-tenancy, and quality of service to which enterprises have grown accustomed with traditional storage systems.

Q. What types of modern workloads could benefit most from unified data storage in hybrid cloud and multicloud environments?

A. While unified data storage and management can provide benefits for a broad spectrum of enterprise workloads, the advantages can be particularly significant for applications that require access to multiple types of data through more than one interface or protocol. For instance, analytics, AI, GenAI, and ML that require large data sets to provide optimal insights may need to combine structured and unstructured data from block-, file-, and/or object-based systems that may extend across on-premises and cloud environments. Having a common model for storage and management through a unified platform can reduce complexity and address concerns such as data protection, security, and compliance.

Other modern workloads that could benefit from unified data storage and management include cloud-native, container-based applications that DevOps teams build and test as part of digital business initiatives. Those applications often require the flexibility to move across on-premises and cloud environments, and unified data storage designed for hybrid multicloud use can provide mobility advantages.

Unified data storage can also help deliver a consistent level of performance to a group of similar workloads that enterprises want to consolidate onto a single platform. Technologies that have removed historical barriers to storage workload consolidation include high-performance flash and nonvolatile memory express (NVMe) technologies that can enable a unified data storage system to lower latency, increase I/O per second, and boost throughput to address a wider range of applications with greater reliability and efficiency. Organizations may not want or be able to consolidate all workloads onto a unified data storage platform, but they can decrease the disparate systems they need to support and potentially realize savings in total cost of ownership without putting performance, availability, security, and other key capabilities at risk.

Q. What are the key selection criteria to consider when evaluating unified data storage infrastructure to meet modern enterprise requirements?

A. Workload requirements are the most important consideration when selecting any new data storage system. Historically, performance was often the critical determinant for enterprises deciding between legacy block-, file-, and/or object-based systems. While performance remains a key criterion, the increasing use of high-performance, NVMe-based flash storage has made it possible to at least consider any type of storage for a more extensive range of workloads than ever before. Workloads will dictate the level of priority for other key evaluation criteria, including data access options, scalability,

availability, data protection, security features, and storage capabilities such as compression, deduplication, snapshots, and replication.

Organizations should do a careful assessment of their applications, prioritize requirements, and set overall storage objectives, such as reducing complexity, containing costs, lowering energy consumption, and sustaining a hybrid multicloud environment. Once organizations identify the most important features to meet their objectives, they can seek out vendors with unified data storage systems that offer those capabilities. Since some workloads may not be conducive to consolidation onto a unified storage platform, grouping application candidates based on workload characteristics may make sense. Enterprises hoping to consolidate workloads on unified data storage will find that some systems support more storage protocols and data access methods than others, and products with software deployment options may offer advantages in hybrid cloud and multicloud environments. Also, some vendors may not offer the performance, scalability, or comprehensive range of features with their unified systems that they do with their more traditional SAN, NAS, and/or object storage products. However, given the ongoing developments and improvements, IT managers would be wise to keep an eye out for opportunities to consolidate workloads where possible to gain the benefits that unified data storage and management architecture can bring.

About the Analyst



Carol Sliwa, Research Director, Infrastructure Systems, Platforms, and Technologies Group

Carol Sliwa is a research director in IDC's Enterprise Infrastructure Practice. Her core area of focus is enterprise storage, including block-, file-, and object-based systems in flash-, hybrid-, and disk-based configurations. With 15 years of experience covering enterprise storage technology, Carol has gained extensive insight into the ways in which the industry has adapted systems over time to address the evolving needs of IT customers.

MESSAGE FROM THE SPONSOR

NetApp is the intelligent data infrastructure company, combining unified data storage, integrated data services, and CloudOps solutions to turn a world of disruption into opportunity for every customer. NetApp creates silo-free infrastructure, harnessing observability and AI to enable comprehensive data management. As the only enterprise-grade storage service natively embedded in the world's biggest clouds, our data storage delivers seamless flexibility. In addition, our data services create a data advantage through superior cyber resilience, governance, and application agility. Our CloudOps solutions provide continuous optimization of performance and efficiency through observability and AI. No matter the data type, workload, or environment, with NetApp you can transform your data infrastructure to realize your business possibilities.

Learn more at www.netapp.com or follow us on [X](#), [LinkedIn](#), [Facebook](#), and [Instagram](#).

**IDC Research, Inc.**

140 Kendrick Street
Building B
Needham, MA 02494
T 508.872.8200
F 508.935.4015
Twitter @IDC
idc-insights-community.com
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 2024 IDC. Reproduction without written permission is completely forbidden.

