

Software-defined storage extends beyond file- and object-based options for unstructured data to block-based alternatives designed to expand capacity for structured data on standard servers.

Growth of Data-Intensive Workloads Drives New and Enhanced Scale-Out, Software-Defined Storage

December 2024

Written by: Carol Sliwa, Research Director, Storage and Converged Systems

Introduction

Two of the top trends influencing enterprise storage requirements today are the shift to hybrid cloud infrastructure and the growing use of data-intensive applications such as AI and analytics as organizations pursue digital business initiatives. IDC's *Enterprise Infrastructure Pulse Storage Survey* shows that 70% of organizations now take a hybrid cloud or hybrid multicloud approach for data storage. The recent survey further indicates the primary drivers for 2024 increases in storage spending are AI and generative AI (GenAI) workloads followed by new and existing applications. IDC's workloads forecast projects that AI life cycle, structured databases, digital services, and data analytics will be the top applications fueling spending on enterprise infrastructure in 2028. However, the traditional scale-up storage area networks (SANs) that enterprises have historically favored to deliver high performance for their mission-critical structured databases may not be the optimal choice for modern data-intensive workloads that require higher levels of storage capacity and data throughput, sometimes on short notice.

AI, analytics, and other data-heavy workloads are driving interest in scale-out, software-defined storage (SDS) alternatives that can run on general-purpose server hardware and help ease scaling through the addition of nodes on an as-needed basis. New disaggregated SDS architectures offer organizations greater flexibility to independently scale compute and storage and focus on the resources they need. Expanding capacity on commodity hardware offers the potential for cost savings over traditional scale-up SAN and network-attached storage (NAS) systems that scale compute and storage in tandem on proprietary hardware at fixed ratios.

AT A GLANCE

KEY STATS

According to IDC survey data:

- » Most organizations take a hybrid multicloud (44%) or hybrid cloud (26%) approach for data storage infrastructure.
- » The top workloads driving spending on enterprise infrastructure through 2028 will be AI life cycle, structured databases, digital services, and data analytics.

KEY TAKEAWAYS


- » Scale-out, software-defined storage can help ease capacity and performance expansion on standard server hardware with data-intensive workloads.
- » On-premises storage as a service offers an alternative to public clouds to speed procurement and provisioning and ease operation and management.

To date, the market for scale-out SDS has largely consisted of file- and object-based systems that target unstructured data, and both are growing at higher rates than scale-up storage. IDC forecasts that shipped capacity will increase at a compound annual growth rate (CAGR) of 26.6% for scale-out file storage and 27.3% for object stores through 2028 across on-premises and public-cloud environments. New scale-out block-based options are also emerging to ease storage capacity expansion for structured data with suitable workloads. IDC's *Enterprise Infrastructure Pulse Storage Survey* shows that organizations currently use block, file, and object storage for AI life cycle, the top workload generating spending on enterprise infrastructure.

The use of storage-as-a-service alternatives also continues to grow, especially for organizations under a time constraint to provision systems for high-priority digital business initiatives. Public cloud storage has been a popular choice for years, and newer on-premises storage-as-a-service (STaaS) options are also drawing considerable interest for select workloads. IDC's *Enterprise Infrastructure Pulse Storage Survey* shows that 86% of enterprises use on-premises STaaS to varying degrees, with 10% of respondents using it for all or most of their storage workloads and the rest using it for 50% or less. Top reasons for using on-premises STaaS include speeding storage deployment, providing storage for AI workloads, easing storage procurement, and offloading storage operations, maintenance, and management (see Figure 1).

FIGURE 1: **On-Premises Storage as a Service**

Q What are the top 3 reasons your organization uses or plans to use on-premises storage as a service that provides a cloud-like model for procurement, deployments, and management of infrastructure?



n = 774

Source: IDC's *Enterprise Infrastructure Pulse Storage Survey*, October 2024

Definitions

- » **Scale-up storage** is the architectural approach that vendors of block- and file-based systems have historically taken. Typical scale-up storage systems have dual controllers for high availability and increase capacity through the addition of expansion shelves of HDDs and/or SSDs. When a scale-up system reaches its maximum capacity or performance level, the user must add or cluster a new dual-controller system to expand resources.
- » **Scale-out storage** is a software-driven system designed to expand capacity and performance through the addition of server hardware nodes to a cluster. Scale-out storage has historically stored unstructured data in discrete file- and object-based systems, and some new options make use of a common key value store to enable data access via file and object protocols. Block-based scale-out storage options are also emerging for structured data.
- » **Software-defined storage** separates the software responsible for provisioning capacity, protecting data, and controlling data placement from the underlying physical storage hardware. SDS systems are designed to run on standard, general-purpose server hardware, in contrast to proprietary storage hardware in use with traditional SAN and NAS systems. They typically expand capacity through the addition of nodes to a storage pool or cluster.
- » **Non-volatile memory express (NVMe)** is a host controller interface and storage protocol designed to accelerate data transfer between servers and storage devices. NVMe-based SSDs connect to a host system's central processing unit (CPU) via a peripheral component interconnect express (PCIe) bus to enable higher throughput and lower latency than storage drives based on the small computer systems interface (SCSI), which require adapters to communicate with the host.
- » **NVMe over Fabrics (NVMe-oF)** enables storage systems to support low-latency access to data by facilitating the transmission of NVMe commands over switched fabrics, including Ethernet-based transmission control protocol (TCP), fibre channel, and remote direct memory access (RDMA).
- » **Quad-level cell (QLC)** NAND flash stores four bits per cell to enable greater storage density than triple-level cell (TLC) NAND flash, which stores three bits per cell. QLC-based SSDs and flash modules cost less on a price-per-GB basis and store more data in a smaller footprint, while TLC-based flash drives support higher performance and endurance.

Benefits

Software-defined storage systems can help ease and speed the scaling of capacity and/or performance through the addition of standard hardware nodes to a server cluster. Historically, capacity planning and procurement took months with traditional SAN and NAS systems, and deployment, operation, and management often required IT staff with special skill sets. Organizations also commonly purchased excess storage resources to meet future performance and capacity requirements, and some resources went unused. Although clustering or federating systems can enable block storage to scale up to significant levels of performance and capacity, the arrays ultimately reach limits, at which point users must purchase a new system and address the risk of creating data silos. Scale-out SDS offers an alternative that can be helpful for organizations confronting rapidly growing volumes of data or deploying new data-intensive workloads, such as AI and analytics, that may need to expand capacity on demand or on short notice. Disaggregated SDS systems that enable independent scaling of compute and storage further allow users to focus on the resources they need and potentially save

money through expansion on commodity server hardware. SDS also holds out the potential for organizations in hybrid cloud environments to run the same system on premises and in the public cloud if their suppliers support it.

Deploying SDS through a storage-as-a-service option can ease storage procurement, provisioning, operation, and management for organizations pursuing digital business initiatives on tight schedules or lacking sufficient in-house storage skills. On-premises STaaS offers an alternative to the public cloud for enterprises facing high data egress charges or confronting regulatory compliance, security, or performance requirements. Options are generally available for colocation facilities and hosted sites in addition to onsite datacenters. Like the public cloud model, on-premises STaaS can enable enterprises to expand or reduce capacity on demand, pay for only the storage they use, and shift infrastructure purchases from a capital expense to a predictable operational expense model. In some cases, on-premises STaaS also offers self-service capabilities to provision infrastructure, a unified console to manage disparate systems, and AI-based monitoring capabilities to detect potential problematic issues.

Considering Hewlett Packard Enterprise's New Block Storage Approach

Hewlett Packard Enterprise (HPE) is a \$30 billion technology vendor with a broad portfolio that includes compute, networking, and storage infrastructure and cloud services. The company ranks among the top 5 worldwide revenue leaders in external enterprise storage systems, anchored by the HPE Alletra Storage line that it launched in 2021 to unify its primary block storage offerings under a common brand name. The HPE Alletra 5000 Hybrid and HPE Alletra 6000 All Flash series are built on HPE's Nimble Storage architecture, and the HPE Alletra 9000 is based on code from the company's high-end HPE 3PAR and HPE Primera storage arrays. HPE GreenLake for Block Storage added storage-as-a-service options for the HPE Alletra arrays in 2022. The HPE GreenLake cloud, which debuted in 2017, offers an expansive range of services for on-premises datacenters, edge sites, and colocation facilities through a pay-per-use, public cloud-like model, backed by a unified control plane to deploy and manage resources that span private and public cloud environments. HPE GreenLake includes services for managing compute, storage, networking, AI inferencing, and vertical business applications.

In addition to selling traditional scale-up storage, HPE has expanded its portfolio with new scale-out multi-protocol (MP) offerings based on the HPE Alletra Storage MP hardware architecture, with AMD EPYC embedded processors, that was introduced in 2023. The company designed the disaggregated, shared-everything HPE Alletra Storage MP architecture to enable independent scaling of standard compute/controller nodes and NVMe-based flash storage expansion shelves, interconnected by a high-speed NVMe-oF back-end switch fabric, with unified management through a cloud console. HPE expects the modular, software-defined MP architecture will not only help alleviate the scaling challenges that customers face with traditional dual-controller systems but also facilitate cost-effective storage expansion for a wide range of workloads, especially those with unpredictable capacity requirements. Announced in April 2023, scale-out HPE GreenLake for Block Storage, based on the company's intellectual property, and HPE GreenLake for File Storage, leveraging software from partner VAST Data, were the first products designed to use the HPE Alletra Storage MP architecture. High-performance, flash-based HPE Alletra Storage MP X10000 object storage followed in November 2024.

Also in November 2024, HPE announced streamlined branding for the storage systems based on the HPE Alletra Storage MP architecture. HPE Alletra Storage MP B10000 is the new brand name for "HPE GreenLake for Block Storage built on HPE Alletra Storage MP." The "B" stands for block, and the "10000" signifies the next generation of the HPE Alletra 5000, 6000, and 9000 portfolio. HPE's scale-out block storage is available through traditional capital spending and service-based

operational expense models, as are the new object store and HPE GreenLake for File Storage. The product brand name for HPE GreenLake for File Storage was not changed.

New offerings and capabilities that became available in 2024 for HPE block storage include:

- » **HPE Alletra Storage MP B10000:** The May 2024 scale-out, block-based HPE Alletra Storage MP release doubled the system's maximum raw capacity to 5.9PB, with each of the 16 JBOF expansion shelves fully loaded with 24 NVMe-based TLC SSDs at 15.4TB. HPE recently announced options for denser 30.7TB or 15.4TB QLC SSDs, raising the system's raw capacity limit to 11.8PB, and its capacity could potentially go higher than that after data deduplication and compression. Organizations can start with a two-node system and scale up and out to multinode configurations, upgrading storage capacity in two-SSD increments and controllers in one-node increments for greater flexibility over traditional paired-node storage architectures. The B10000 uses AMD EPYC processors and offers 32-core, 16-core, and 8-core configuration options, with up to 128 Peripheral Component Interconnect Express (PCIe) 5.0 lanes for high throughput. The B10000's newly added three-site replication provides an additional option for data protection, and HPE's Zerto Cyber Resilience Vault now integrates with HPE Alletra Storage MP to enable immutable data copies and air-gapped recovery to address the threat of cyberattacks.

On the management front, HPE Alletra Storage MP B10000 now integrates HPE Infosight capabilities to enable AI-driven, cross-stack analytics capabilities designed to detect and mitigate performance issues across virtual machine (VM), storage, network, compute, and cloud infrastructure. The new cross-stack analytics capabilities combine with AI-driven workload intelligence offered through HPE OpsRamp to provide greater visibility across the infrastructure stack. For instance, HPE says the new cross-stack analytics can provide a topology map of the resources used by every VM connected to an HPE Alletra Storage MP B10000 array to help determine if a VM latency issue stems from the storage, the host application server, or the network by correlating VM performance in a data store with insights on hypervisor and host resource constraints, such as the CPU, memory, or network. The AI operations (AIOps) capabilities could, in turn, help users address "noisy-neighbor" VMs and reclaim space from underutilized VMs. Another new management feature enables HPE Alletra Storage MP B10000 customers to observe energy consumption and carbon emission trends through the HPE Sustainability Insight Center.

- » **HPE Alletra Block Storage for Amazon Web Services (AWS):** This new public-cloud alternative based on the Alletra Storage MP architecture is designed to enable organizations to run the same software-defined HPE block storage on premises and in AWS and move data, workloads, and backups with greater ease and at lower cost. The HPE GreenLake Cloud Console aims to facilitate unified management across hybrid environments, enabling users to maintain consistent policies for data protection and security. Target use cases for HPE Alletra Block Storage for AWS include disaster recovery (DR), with an on-premises HPE Alletra Storage MP array able to replicate to a software-defined HPE MP-based system in AWS as a DR site, and test/development, giving organizations the ability to build applications on premises or in the cloud, then test and deploy them using storage at either site to minimize disruption in production environments.
- » **HPE Timeless Program:** This new storage investment protection program qualifies participating HPE Alletra Storage MP B10000 customers for a nondisruptive, data-in-place upgrade to a next-generation controller with a subscription renewal. HPE claims a customer could save 30% on storage TCO by avoiding the forklift upgrade inherent with a traditional scale-up system. The new HPE Timeless Program includes a 30-day satisfaction

guarantee, access to all software and operating system features, the company's long-standing 100% data availability guarantee, and the HPE StoreMore 4:1 data reduction guarantee.

Challenges

HPE Alletra Storage MP B10000 represents a new approach to block storage over the traditional SANs that companies have used for years to meet the requirements of their mission-critical databases and other high-performance applications. The B10000 is built on the HPE Alletra Storage MP architecture that HPE designed to give organizations greater flexibility to scale out compute/controller nodes separately from NVMe-based flash storage expansion shelves, based on their needs for increased performance or capacity. HPE may need to educate customers on the new MP-based block storage and its management and monitoring capabilities to ease the transition to the new system. The HPE Alletra Storage MP B10000's 100% data availability guarantee should help address potential concerns about using the new architecture. However, as with any new technology approach, HPE's scale-out, disaggregated block storage may require considerable testing to give customers accustomed to using scale-up SANs the comfort level and confidence that the new system will meet the performance and other requirements of their mission-critical workloads. Proven field results will be especially important since HPE plans to eventually shift its entire block storage portfolio to a common software stack running on standard HPE Alletra Storage MP hardware over the coming years in the highly competitive market for all-flash block storage.

The new HPE Alletra Storage MP B10000 is available through traditional licensing models and as an HPE GreenLake service for customers that prefer to offload the provisioning, operation, and maintenance of the system. The company faces considerable competition in the on-premises STaaS market, so keeping HPE GreenLake prices in check will be important. HPE's use of standard hardware components with the HPE Alletra Storage MP architecture should help in that respect.

Conclusion

Growth of data-intensive workloads such as AI and analytics is driving greater interest in scale-out, software-defined storage systems that can ease the expansion of capacity and performance on standard server hardware. SDS systems can in some cases enable the independent scaling of compute and storage resources in hybrid cloud environments. Organizations that turned to public cloud services to speed the provisioning and deployment of storage are also showing high levels of interest in on-premises storage as a service as they confront high data-egress charges and face regulatory compliance, security, and performance requirements. Available through traditional purchasing models and as an HPE GreenLake service, HPE Alletra Storage MP B10000 is a flash-based software-defined storage alternative designed to provide greater flexibility to scale compute nodes separately from NVMe-based storage expansion shelves to expand performance and capacity. IDC expects the markets for scale-out SDS on commodity server hardware and for on-premises storage as a service will continue to grow, and as a well-established player in the enterprise storage systems market, HPE has a significant opportunity to help organizations meet their modern storage requirements, if it can address the challenges outlined in this paper.

Growth of data-intensive workloads such as AI and analytics is driving greater interest in scale-out, software-defined storage systems that can ease the expansion of capacity and performance on standard server hardware.

About the Analyst



Carol Sliwa, Research Director, Storage and Converged Systems

Carol Sliwa is a research director in the Storage and Converged Systems practice of IDC's Worldwide Infrastructure Research organization. She is the lead analyst for IDC's storage infrastructure coverage, and her core research spans block, file, and object storage, with a special focus on unstructured data. Carol's coverage also includes flash storage media and scale-out storage systems for performance- and data-intensive workloads. With more than 25 years of experience as a technology journalist, including 13 years covering enterprise storage, Carol gained extensive insight on the ways in which the IT industry has developed technologies, platforms, and systems over time to address the evolving needs of IT organizations.

MESSAGE FROM THE SPONSOR

HPE Alletra Storage MP B10000 delivers the industry's only disaggregated, scale-out block storage with a 100% data availability guarantee. Built on HPE Alletra Storage MP hardware and managed via the HPE GreenLake cloud, this unique block storage offering brings an AI-driven cloud experience, efficient scale, and extreme resiliency and performance to mission-critical apps — from large databases to modern cloud native apps to consolidated mixed workloads.

Key Benefits

- » Scale performance and capacity independently for greater efficiency and **up to 40%** lower costs
- » Free non-disruptive controller refresh reduces TCO by **30%** over traditional forklift upgrades
- » Simplified management via HPE GreenLake cuts operational time savings by **up to 98%**
- » Industry-leading AI-Ops prevent **86%** of time-consuming disruptions before they occur
- » Eliminate costly downtime with a **100% data availability guarantee**
- » Slash storage costs, power usage and carbon/rack footprint with **4:1 total savings guarantee**
- » Utilizing the AMD EPYC processors' combination of high core density, energy efficiency, and 128 PCIe lane support to optimize performance, scalability, and efficiency

Find out more at hpe.com/us/en/alletra-storage-mp-b10000



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

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

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.