

Data-intensive workloads are driving the growth of scale-out storage systems designed to ease capacity and performance expansion. Storage as a service can help simplify deployment, operation, and management at datacenter, colocation, and hosting sites.

On-Premises Scale-Out Storage Services for Data-Intensive Workloads

May 2023

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

Introduction

Recent IDC survey data shows the top reasons that enterprises plan to increase storage spending this year are significant data growth, requirements for higher-performing systems, and the need for greater data protection. Historically, such objectives might have led to months of planning for the purchase of a new higher-performing storage system and expansion shelves. This option often required special networking gear and skilled IT staff to operate and manage it. Fears of hitting capacity or performance limits drive many organizations to buy excess storage infrastructure, and some of it commonly goes unused. Clustering systems can enable enterprises to scale up to substantial levels of capacity, but the storage infrastructure can be complicated to set up, operate, and manage.

Today's modern data-intensive applications are driving the growth of software-led storage architectures designed to ease the scaling of capacity and performance through the addition of nodes to a cluster on an as-needed basis, store exabytes of data with no theoretical limits, and deliver high throughput to speed queries. In particular, workloads that stand to benefit from the newer scale-out storage designs include analytics and artificial intelligence (AI) applications that enterprises are increasingly pursuing to gain greater business benefit from their data.

Current scale-out storage options largely consist of file- and object-based systems that target ever-growing volumes of unstructured data. IDC projects that shipped capacity of scale-out file and object storage across on-premises and public cloud sites will increase at a 25% compound annual growth rate (CAGR) from 2021 to 2026, while spending grows at a 13% CAGR. By contrast, IDC forecasts that shipped capacity for scale-up file storage will expand at a more modest 9% rate, while spending remains flat. Forecasts are not available for the emerging scale-out block storage systems, but IDC projects shipped capacity for the overall block storage market will grow at a 16% CAGR from 2021 to 2026, and revenue for the mature market will increase by 3%. Technologies paving the way for new scale-out block storage systems designed to target structured data with database workloads include nonvolatile memory express (NVMe) solid state drives (SSDs) and NVMe over Fabrics (NVMe-oF) networking that can speed throughput and reduce latency.

AT A GLANCE

KEY STATS

According to IDC survey data:

- » 79% of enterprises either use or plan to use on-premises storage as a service (STaaS) in 2023.
- » Primary reasons for STaaS use include the need to burst storage capacity on demand and the need to speed storage deployment.

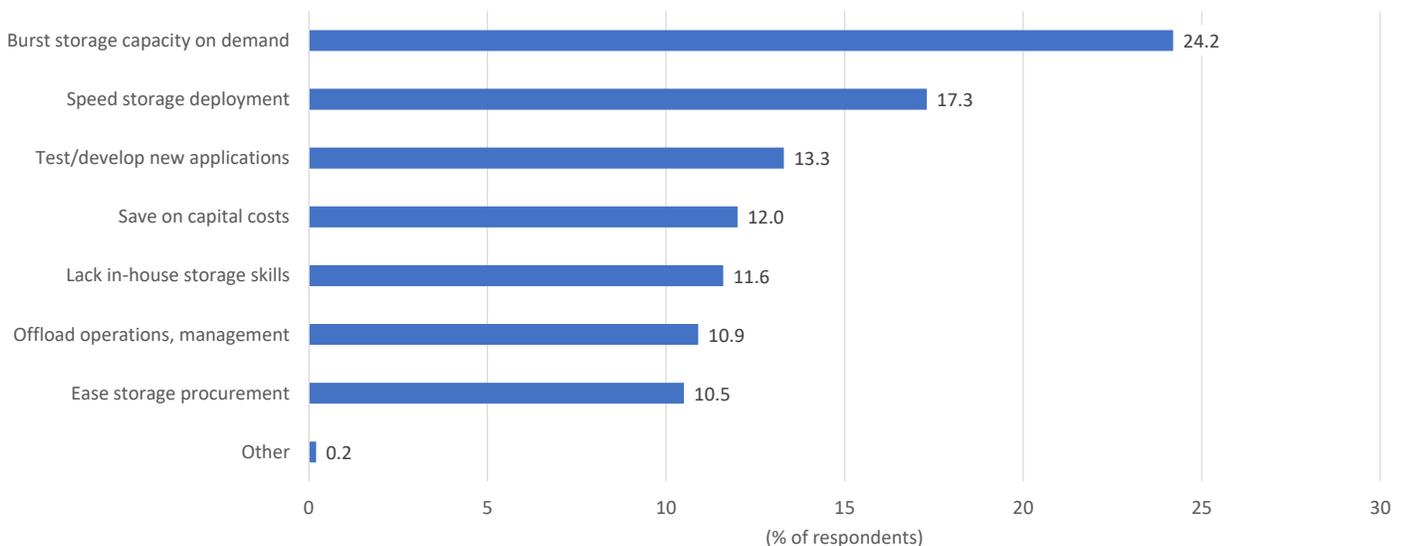
KEY TAKEAWAYS

- » On-premises STaaS aims to simplify the procurement, operation, deployment, and management of storage infrastructure.
- » Scale-out storage is designed to ease capacity and performance expansion with data-intensive workloads.

The modern data-intensive workloads that digitally transforming enterprises increasingly deploy are putting pressure on IT organizations to deliver storage on tighter schedules than they typically did in the past, driving many to pursue public cloud storage alternatives. Storage-as-a-service (STaaS) options introduced in the past few years by traditional storage vendors offer an on-premises alternative for organizations that want to keep their data in local datacenters, colocation facilities, or hosted sites due to concerns ranging from security and regulatory compliance to latency and cloud egress fees. Recent data from IDC's *IT Infrastructure for Storage and Data Management Survey* shows that 47% of organizations use on-premises STaaS offerings from leading storage vendors. Another 32% plan to use on-premises STaaS in 2023, and 17% will consider it in the future. Organizations cited the need to burst storage capacity on demand and the need to speed storage deployment as the primary reasons for using STaaS (see Figure 1). IDC also predicts that backup as a service will experience significant growth, at a 19.1% CAGR, and 55% of organizations will adopt a cloud-centric data protection strategy by 2025.

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

Q What is the primary reason that your organization uses or plans to use storage-as-a-service offerings that provide a cloudlike model for procurement, deployment, and management of infrastructure?



n = 475

Source: IDC's *IT Infrastructure for Storage and Data Management Survey*, January 2023

Definitions

Storage as a service enables customers to procure and provision storage capacity on demand, pay for only what they use as an operational expense rather than an up-front capital purchase, and offload the operation, management, and maintenance of the infrastructure to the provider. STaaS options are available from public cloud providers and on-premises storage vendors with varying levels of self-service capabilities, access and storage types, performance and service levels, administration features, and pricing options.

Scale-up storage is the traditional architectural approach in use with block- and file-based systems. Typical scale-up storage systems have dual controllers for high availability and increase capacity through the addition of expansion shelves of hard disk drives (HDDs) and/or SSDs. When a scale-up system reaches its maximum capacity or performance, the user must add or cluster another dual controller system to expand resources.

Scale-out storage is a software-led system designed to expand capacity and performance through the addition of server hardware nodes to a cluster. Scale-out systems most commonly store unstructured data in file- and object-based storage, but scale-out block-based options are emerging to target structured data.

NVMe is a host controller interface and storage protocol designed to accelerate data transfer between servers and storage devices. NVMe-based SSDs connect to the host system's central processing unit (CPU) via a Peripheral Component Interconnect Express (PCIe) bus to enable higher throughput and lower latency than Small Computer Systems Interface (SCSI)-based storage drives that require adapters to communicate with the host.

NVMe-oF enables storage systems to support low-latency access to data similar to the way NVMe-based PCIe SSDs do in a server. NVMe-oF facilitates the transmission of NVMe commands over switched fabrics, including Ethernet-based Transmission Control Protocol (TCP), Fibre Channel, and Remote Direct Memory Access (RDMA).

Benefits

On-premises storage as a service aims to simplify the procurement, deployment, operation, management, and maintenance of IT infrastructure for organizations that want or need to keep their data in local datacenters, colocation facilities, or hosted sites for security, performance, regulatory compliance, or other reasons. Shifting to a public cloud-like model can enable organizations to expand or reduce capacity on demand, pay for only the storage they use, and shift their infrastructure purchases from a capital expense to a more predictable operational expense, if they choose. The most advanced on-premises STaaS options further seek to mimic the public cloud experience with self-service capabilities to provision infrastructure, manage disparate systems through a unified interface, monitor storage with AI-based technologies, and nondisruptively upgrade systems without manual intervention. STaaS can be especially helpful for enterprises lacking adequate in-house storage experts to cope with growing volumes of data or new data-intensive workloads they want to test and put into production on tight schedules.

Whether deployed as a service or through traditional means, scale-out storage systems have the potential to help organizations address data growth and data-intensive workloads such as analytics, AI, machine learning, deep learning, high-performance computing (HPC), and performance-oriented data lakes. The systems are designed to ease and speed the scaling of capacity and/or performance through the addition of standard server hardware nodes to a cluster, minimizing the need for a lengthy capacity planning process.

Considering HPE GreenLake for Block and File Storage and Data Protection

Hewlett Packard Enterprise (HPE) is a \$28.5 billion technology vendor with a broad portfolio that includes cloud services, compute and storage infrastructure, software, and HPC and AI solutions designed for use in datacenters and edge environments. It reported \$4.7 billion in storage revenue in fiscal 2022, anchored by the HPE Alletra, HPE Primera, and HPE Nimble storage array lines. Exiting its 2022 fiscal year on October 31, HPE noted more than \$8.3 billion in total contract value for the HPE GreenLake edge-to-cloud platform that has become a major area of focus for the company. HPE introduced GreenLake in November 2017 to provide a public cloud-like "pay per use" model for customers to acquire on-premises

infrastructure and facilitate infrastructure modernization. Initial offerings focused on big data, backup, open database, SAP HANA, and edge computing. The HPE GreenLake portfolio has since expanded to encompass more than 50 services, with options for unified management of hybrid environments that span public and private cloud sites.

The March 2022 launch of HPE GreenLake for Block Storage added mission-critical, business-critical, and general-purpose storage services built on the HPE Alletra 9000, HPE Alletra 6000, and HPE Alletra 5000 arrays, with a cloud operational and management model. With HPE GreenLake for Block Storage, the company sought to address challenges in the burgeoning STaaS market by offering a 100% data availability guarantee for designated mission-critical workloads and "six-nines" of availability for all other business-critical or general-purpose applications. In April 2023, HPE attempted to fill additional market gaps with the launch of its HPE GreenLake scale-out block and file storage services, which will also be available through traditional capital purchasing mechanisms.

The latest HPE GreenLake for Block Storage and new HPE GreenLake for File Storage services are built on a modular HPE Alletra Storage MP hardware architecture that can independently scale compute and storage resources, pairing different combinations of controller nodes and capacity nodes to suit workload requirements. The company expects the new disaggregated architecture will deliver cost savings over scale-up systems that scale performance and capacity together. The RDMA-based 100 Gigabit Ethernet fabric is designed to reduce latency across the compute and storage nodes on the backplane, and the system supports Ethernet-based NVMe over TCP to streamline external communication. HPE Alletra Storage MP supports multicore AMD CPUs, PCIe 4.0 (with plans for PCIe 5.0), and NVMe-based SSDs to facilitate parallel throughput for data-intensive workloads. HPE expects the use of common HPE Alletra Storage MP hardware components across the block and file storage systems will help speed the delivery of new features and capabilities, streamline its supply chain, lower costs through economies of scale, and ease customer support.

Although the new scale-out block and file offerings use the same underlying HPE Alletra Storage MP hardware, they are discrete services that rely on different software to manage the storage. HPE GreenLake for Block Storage is built on the company's own technology. HPE GreenLake for File Storage leverages software from partner VAST Data, a fast-growing start-up noted for high-performance, exabyte-scale unified file and object storage. The new HPE GreenLake service supports only VAST Data file storage, and VAST updated its software stack to work with HPE hardware and the HPE GreenLake cloud management console. VAST designed its Disaggregated Shared Everything (DASE) architecture to separate the storage media from the CPUs that provide storage services and manage the drives. Although VAST supports NVMe SSDs with quad-level (QLC) NAND flash memory to help reduce costs and storage footprint, HPE GreenLake for File Storage currently uses less dense triple-level cell (TLC)-based SSDs that offer higher endurance and performance. The company confirmed plans to add options for dual-ported NVMe-based QLC SSDs in the future, and the system already supports storage class memory (SCM) to help reduce the wear on SSDs and speed data processing.

HPE GreenLake for Block Storage is positioned as mission-critical storage at midrange economics. The new block storage consists of a switchless 2U chassis that can scale to a raw capacity of 368TB with 15TB SSDs, with the same 100% data availability guarantee that is offered with other mission-critical arrays. HPE expects the first release will be especially attractive for HPE 3PAR customers and offers built-in tools to help with the migration. The company also plans to extend its HPE Alletra Storage MP-based scale-out block storage architecture to larger 32-core, multinode switched systems by the end of 2023 and to the rest of its HPE Alletra, HPE Primera, and HPE Nimble storage systems over the coming years. The HPE GreenLake for Block Storage operating system includes technologies from HPE's Alletra 9000/Primera/3PAR line, such as active-active access to synchronous replicated volumes (HPE Active Peer Persistence), as well as HPE Alletra 6000/HPE Nimble Storage, including

compression heuristics and redirect-on-write snapshots. Although performance figures were not available for the new HPE GreenLake for Block Storage offering, HPE claims throughput, IOPS, and latency will improve with its new scale-out block storage over its comparable scale-up block arrays.

Customers will also have access to integrated hybrid cloud data protection services through the HPE GreenLake platform. A new HPE GreenLake for Disaster Recovery option, based on HPE's Zerto acquisition, is due in June to complement the HPE GreenLake Backup and Recovery service, which protects on-premises and cloud workloads and consolidates workload visibility to ease management. Additional technologies that may be helpful for storage customers include the AI-driven HPE InfoSight to manage and optimize infrastructure and HPE CloudPhysics to tune performance and data protection through analysis of IT metadata.

An AI-driven HPE GreenLake Data Services Cloud Console enables customers to manage and monitor their potentially geographically dispersed block, file, and data protection services and consolidate tasks such as billing, user management, and device inventory through a unified interface. HPE designed a Setup Service to ease the installation and deployment of storage arrays and minimize manual intervention. A Data Ops Manager allows customers to apply software updates across their entire installed base at once rather than maintaining systems in piecemeal fashion. Additional services offer AI-driven recommendations to assist with provisioning and protecting systems and to facilitate automation through prebuilt policies that users can customize.

Challenges

Scale-out block storage built on the standard hardware building blocks of HPE's new Alletra Storage MP architecture represents an innovative development with the potential to benefit the company and its customers in many ways, including faster delivery of new features, a streamlined supply chain, and lower costs. HPE is wisely beginning the rollout of the scale-out block storage in relatively small configurations before ramping up to large systems, with 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. HPE GreenLake for Block Storage offers a 100% data availability guarantee that can help address potential customer concerns about using the new architecture. However, as with any new technology, HPE's scale-out block storage may require significant testing and proven field results to give customers that are accustomed to traditional scale-up block storage the confidence that the scale-out block approach will meet their mission-critical workload requirements, including from a performance standpoint.

HPE plans to offer scale-out block storage through traditional licensing models and as a service through its HPE GreenLake platform. This service-based option may give customers a greater comfort level to use the new scale-out block storage, with HPE responsible for installation, operation, and maintenance. HPE may also want to educate customers that may be hesitant to have their mission-critical data out of their direct control about the policy-based management capabilities that it makes available for them to dictate how their data is stored, protected, and accessed. The company faces considerable competition in the on-premises STaaS market from storage vendors, including Dell, NetApp, and Pure Storage, as well as Amazon Web Services with its Outposts offerings. Keeping HPE GreenLake prices in check will be important, and HPE's use of standard hardware components with the Alletra Storage MP architecture should help in that regard.

Conclusion

The explosive growth of data and the increasing use of data-intensive workloads such as analytics and AI are fueling scale-out storage architectural approaches designed to ease the expansion of capacity and performance. Enterprises that turned to public cloud services to speed the deployment of storage and burst capacity on demand now have more options for on-premises services to ease the procurement, deployment, operation, management, and expansion of storage infrastructure. New HPE GreenLake for Block Storage and HPE GreenLake for File Storage offer customers innovative service-based alternatives to store structured and unstructured data in scale-out storage systems located in their own datacenters, colocation facilities, and hosting sites, with a cloud operational and management model. The storage services use a common high-performance hardware platform designed for data-intensive workloads, and HPE also gives customers the choice of traditional storage procurement and licensing for the HPE GreenLake "pay per use" services. New HPE GreenLake for Disaster Recovery and previously released HPE GreenLake Backup and Recovery extend the comprehensive storage service portfolio to data protection. IDC expects the markets for on-premises storage, backup, and disaster recovery as a service, as well as scale-out storage, will grow, and to the extent that HPE can address the challenges cited in this paper, it has a substantial opportunity to help enterprises meet their storage needs.

The explosive growth of data and the increasing use of data-intensive workloads such as analytics and AI are fueling scale-out storage architectural approaches designed to ease the expansion of capacity and performance.

About the Analyst



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

Carol Sliwa is a Research Director for Storage Systems in IDC's Enterprise Infrastructure Practice. Her core research area spans block, file, and object storage, with a special focus on the storage of unstructured data. With more than 25 years of experience as a technology journalist, including 13 years covering enterprise storage, Carol 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

HPE GreenLake for Block Storage provides mission-critical resiliency at midrange economics with the industry's first disaggregated, scale-out block storage with a 100% data availability guarantee. This enhanced cloud data service brings an intuitive cloud experience, efficient scale, and extreme resiliency and performance to mission-critical apps – at a midrange price point.

HPE GreenLake for File Storage delivers enterprise-grade, scale-out file storage to accelerate rapidly growing, data-intensive workloads. This new cloud service offers enterprise performance at scale, simple file data management, and faster time to insights.

HPE GreenLake for Block Storage and HPE GreenLake for File Storage run on common HPE Alletra Storage MP hardware and are managed through the HPE GreenLake platform with a unified cloud experience, eliminating silos and complexity and improving TCO/ROI.

Learn how to accelerate data-first modernization by radically simplifying data management with the industry's most comprehensive solution to store, manage, and protect data across hybrid cloud.

[Start today](#)



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