



**HPE** Alletra Storage

## **HPE Alletra Storage MP X10000**

Intelligent, enterprise-grade, scale-out object storage  
to accelerate data-intensive applications

Data is the life force of modern business. It powers the insights, innovation, and competitive advantages that move your business forward, faster. Yet while data-driven modernization is a top priority, achieving it requires confronting a host of data storage challenges that slow you down: lack of data intelligence, enterprise performance at scale, and management complexity.

HPE Alletra Storage MP X10000 is a unique, software-defined, scale-out data system that combines automated metadata enrichment services, high performance all-flash object storage, massive capacity, and effortless management. The X10000 accelerates time to value for your data-intensive initiatives, whether you are building active data lakes for analytics, implementing generative AI and large language models (LLMs), or enabling rapid restore from backups. The X10000 takes full advantage of the industry's first disaggregated multiprotocol architecture,<sup>1</sup> allowing it to scale from terabytes to exabytes on the same hardware. Cost efficiency is achieved by the ability to scale capacity and performance independently.

HPE Alletra Storage MP X10000 is managed through the HPE GreenLake cloud, providing simplified and unified cloud management of block, file, and object storage services. This enables enterprises to optimize their hybrid estate to take full advantage of AI.

## Accelerate unstructured data management and insights

The X10000 provides fast and scalable data handling and processing capabilities, including a built-in data intelligence engine. It is a versatile platform that represents a significant leap forward in AI-ready enterprise class object storage. Now you can manage, secure, and store massive amounts of data efficiently, benefit from all-flash based performance, scale across hybrid cloud environments, unlock data value, and achieve faster time to outcomes.

## Vitalize data and accelerate data access

- **Access data across your hybrid estate:** Built to integrate with both on-premises and cloud-native environments, the X10000 is a software-defined, containerized solution utilizing Kubernetes-based orchestration. Leverage a unified data foundation for AI and analytics across hybrid cloud by integrating with the HPE Data Fabric global namespace.
- **Infuse your data with intelligence:** Bring data closer to compute with a built-in inline data intelligence engine that automatically adds custom metadata to objects. This enhances the speed, quality, and relevance of your data for Retrieval-Augmented Generation (RAG), large language models (LLMs), and analytics—all while reducing infrastructure costs.
- **Fortify your data:** Protect and secure data with object versioning, immutability, authentication, data in-place and in-flight encryption, detailed monitoring, auditing, always-on data durability processes, and more.

## Sustain efficient performance at scale

- **Built for data-intensive workloads:** Power your most demanding AI, ML, active data lake, and data protection workloads with a high performance, low latency solution optimized for flash. The X10000 leverages a disaggregated architecture to maximize resource utilization and ensure linear performance at scale.

<sup>1</sup> [“HPE GreenLake for Block Storage built on HPE Alletra Storage MP”](#) HPE, 2024

- **Increase throughput and lower latency:** HPE is collaborating with NVIDIA® to enable a direct datapath for direct memory access (DMA) transfers between GPU memory, system memory and the X10000, which is critical for AI applications. This direct path will further increase system bandwidth, decrease latency and improve GPU utilization.
- **Restore data in a flash:** Modernize your backup environment with one or more data protection accelerator nodes for HPE Alletra Storage MP X10000 to achieve industry-leading up to 1.2 PB/ hour backup storage throughput and 60x data reduction, with rapid restores from outages, disasters, and ransomware attacks.
- **Rightsize with ease to support diverse workloads:** Independently scale performance and capacity using disaggregated storage and a choice of TLC or QLC drives, eliminating overprovisioning and lowering costs for evolving workloads. Rebalance performance and capacity during expansion with minimal performance impact and no need for expensive data movement.
- **Start small, scale big:** Don't be forced to overprovision right from the start. Start with three nodes and expand to hundreds—without disruption—while ensuring storage efficiency and high availability.
- **Store data efficiently and densely:** Cut storage costs, save energy, boost sustainability, and shrink your data footprint with efficient inline data reduction and erasure coding technologies.



## Simplify with an intuitive cloud operational experience

- **Unify management for IT and end users:** Manage, monitor, and protect your global storage environment from a single cloud console, which makes managing hundreds of systems as simple as managing one. An intuitive software-as-a-service (SaaS)-based user experience means everything you need to manage your fleet of data infrastructure is at your fingertips—accessible from anywhere, on any device.
- **Scale simply, non-disruptively, and infinitely:** Seamlessly scale without pausing your business operations based on an elastic cloud-native software platform that scales autonomously with your data infrastructure.
- **Optimize your infrastructure and meet service-level agreements (SLAs):** Take the guesswork out of provisioning decisions. HPE GreenLake Cloud offers improved visibility into the utilization of your current infrastructure, helping to avoid over-saturation scenarios that could compromise SLAs.
- **Simplify onboarding:** Streamlined device deployment can be completed in minutes. New systems are automatically discovered and onboarded.
- **Ensure secure management of global infrastructure:** Secure entry points end-to-end—from identity and access management to back-end hardening and intrusion detection using advanced multilevel capabilities.
- **Cloud-like consumption:** HPE GreenLake Flex allows cloud-like procurement through a pay-as-you-go\*, consumption-based model. Scale storage and compute capacity up and down while paying only for the capacity consumed.

\* May be subject to minimums or reserve capacity may apply





## Capitalize on our expertise and hit the ground running

Trust in a solution developed by HPE, an industry leader in helping organizations accelerate outcomes by unlocking value from their data. Rely on the considerable expertise of our storage technologist team, who will work with you in partnership with our solutions and product experts to ensure a smooth transition to data-first modernization.

## Get started with HPE Alletra Storage MP X10000 today

Now is the time to migrate to a modern, data-first solution to accelerate your AI, data protection, and digital transformation journeys. In today's AI era, you gain more insights by analyzing larger data sets collectively and utilizing architectures that offer superior performance. This approach will let you interact with your data in unprecedented ways. When it comes to creating value from your data, HPE Alletra Storage MP X10000 is a game changer.



## Learn more at

[HPE.com/HPE-Alletra-Storage-MP/X10000](https://HPE.com/HPE-Alletra-Storage-MP/X10000)

Visit [HPE.com](https://HPE.com)



Chat now (sales)

© Copyright 2025 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

NVIDIA is a trademark and/or registered trademark of NVIDIA Corporation in the U.S. and other countries. All third-party marks are property of their respective owners.

a00143713ENW, Rev. 2

