



GET FILE STORAGE PERFORMANCE AT AI SCALE

HPE GreenLake for File Storage gives you
scalable performance



Enterprise performance at scale is essential for modern, data-intensive applications such as artificial intelligence (AI), machine learning (ML), Big Data, deep learning (DL), high performance computing (HPC), large data lakes, and high-speed access to data archives. These applications span a wide range of industry verticals — including life sciences, financial analytics, and media and entertainment — and process enormous volumes of file data.

To achieve performance at AI scale, file storage must have an architecture that breaks through performance bottlenecks. HPE GreenLake for File Storage has just such an architecture.

HPE demonstrates leadership in storage for AI with NVIDIA DGX BasePOD™ certification and NVIDIA OVX™ storage validation for HPE GreenLake for File Storage. NVIDIA DGX BasePOD is a fully validated reference architecture with a proven server, networking, and storage framework. This reduces the deployment and operational complexity of AI workloads, providing a fast track to AI success.

With this certification, customers can confidently use HPE GreenLake for File Storage as the storage component of the NVIDIA DGX BasePOD reference architecture to run the most impactful, mission-critical AI use cases. These include Health and Life Sciences, Financial Services, Energy, Telecommunications, and Large Language Models (LLMs).

HPE GreenLake for File Storage accelerates AI, GenAI, and GPU-intensive workloads at scale with NVIDIA® AI Computing by HPE and is included in the HPE Private Cloud AI portfolio.



Enterprise performance at AI scale

When HPE announced HPE GreenLake for File Storage, it was to address the need for enterprise performance at scale, which is essential for AI workloads because they are the most data-intensive. Enterprise performance at AI scale breaks through the limits of current systems and goes way beyond where they begin to fall short.

Consider the following analogy of two different vehicles. One is a race car, which is a complex, highly customized system driven by a solitary, highly skilled driver. The car can set a record by reaching a breakthrough speed for a few seconds but has no practical, day-to-day transportation value.

On the other hand, a vehicle that is also fast and can easily transport many people is more effective, even if it doesn't set a land speed record. In this analogy, the passengers represent the data that's our real concern because of its value. Enterprise performance at AI scale isn't simply flat-out performance that reaches an unprecedented instantaneous peak for a small data set. Instead, it's a fast and sustained performance that spans the entire scale of your data and all the stages of AI to support workloads that process extremely high volumes of data.

Another way to understand enterprise performance at scale is to imagine a sprinter who can run a marathon at a world-class 100-meter pace. It's a sustained, high-speed performance over a long stretch. And this is just what HPE GreenLake for File Storage delivers. Where legacy NAS infrastructure hits a limit and can't keep

going, HPE GreenLake for File Storage continues to scale and sustain performance even as it processes huge amounts of data.


An architecture for enterprise performance at AI scale

What allows HPE GreenLake for File Storage to achieve enterprise performance at AI scale? It's a solution designed for exabyte scale with two key components: VAST Data software with a Disaggregated Shared Everything (DASE™) architecture and HPE Alletra Storage MP modular, resilient hardware, offering independent scaling of performance and capacity.

The HPE Alletra Storage MP compute nodes are connected to HPE Alletra Storage MP storage nodes over an NVMe fabric. The storage nodes contain ultra-efficient, all-NVMe storage for blazing-fast performance, and every compute node can access all storage nodes.

HPE GreenLake for File Storage helps eliminate the drawbacks of legacy, scale-out NAS, including crosstalk, rebuilds, and interdependencies that only increase geometrically with cluster size. As the architecture is designed to grow to exabyte capacities, it future-proofs your storage investment. HPE GreenLake for File Storage delivers linear performance scaling while overhead remains flat. The result is more performance scaling than most organizations can ever consume. Independent scaling of performance and capacity provides flexibility and efficiency in meeting exactly the performance and capacity points you need.

Shared-nothing architecture



Cross talks, rebuilds, and interdependencies increase geometrically with cluster size

*Per cluster

- ✗ Designed for exabyte scale ✓
- ✗ Flat overhead/Linear performance scaling ✓
- ✗ ... Independent scaling of performance and capacity designed for infinite lifecycle (no migration) ✓
- ✗ Global namespace at scale* ✓
- ✗ Unlimited controller node failures allowed (N-1) ✓
- ✗ No controller resource contention ✓
- ✗ No rebuilds on server failure ✓

Disaggregated, shared-everything cluster architecture




Figure 1. Traditional shared nothing vs. DASE architectures

All of this means you can easily scale up and scale out. You can add compute and storage resources as needed instead of getting another box or appliance that sacrifices flexible and efficient scaling. The architecture of HPE GreenLake for File Storage reduces migration issues and is designed for an infinite data lifecycle.

With HPE GreenLake for File Storage, there is no controller resource contention. Controller nodes are stateless, as all the metadata is kept in the storage nodes in persistent, fast storage class memory. All storage nodes have access to all the metadata.

You have a global namespace for your storage system for data sharing, and you can have up to N-1 controllers fail and still keep going. There are no rebuilds on server failure, and for data protection, a new erasure code algorithm can rebuild a failed SSD without reading the entire large stripe that the system uses. This results in a low overhead of around 2.7%.

These capabilities are not possible with shared-nothing systems. Figure 1 shows contrasting summaries of the two different architectures.

With shared nothing, the system is limited by:

- Ownership of partial data
- Fast performance only if directly attached to the servicing node. Even so, at scale, there will likely be contention that incurs significant overhead
- Significantly slower performance in the event of a node failure

With shared everything, the system provides:

- Access to all data for all compute nodes
- Performance scaling without adding overhead from east-west node traffic
- Continuous operations with even up to N-1 compute node failures (you just need one functional node)
- Load balancing without resource contention



Supercharge AI and other data-intensive applications

Combine all the DASE advantages with the HPE GreenLake cloud, and you have a truly compelling file storage solution. Along with blazing-fast performance, you also get an intuitive cloud experience; over 50 other cloud services spanning storage, networking, compute, and workloads; and the security, control, performance, and cost advantages of an on-prem infrastructure.

What's more, you can experience streamlined deployment, easy file share creation, unified storage management with a single cloud console, and automated, nondisruptive upgrades. And, with up to 2.3x the capacity density of Dell PowerScale and Pure Storage FlashBlade,¹ you get a file storage architecture designed for the enterprise performance at AI scale that is essential for today's AI and other modern workloads.

HPE GreenLake for File Storage is an ideal file storage solution for today—and tomorrow—as you continue to amass, process, and extract value from data. So, seize your competitive advantage by supercharging your most demanding AI and other data-intensive applications with the enterprise performance at AI scale of HPE GreenLake for File Storage.

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

Learn more at

[HPE.com/us/en/HPE-GreenLake-file-storage.html](https://www.hpe.com/us/en/HPE-GreenLake-file-storage.html)

¹ Comparisons are based on raw capacity for capacity density with Dell PowerScale F900 and Pure Storage FlashBlade.

[Chat now](#)

© Copyright 2025 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties 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 DGX and NVIDIA are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. All third-party marks are property of their respective owners.

a50009538ENW, Rev. 3

HEWLETT PACKARD ENTERPRISE

[hpe.com](https://www.hpe.com)

