



HPE Storage

Top 10 reasons to choose HPE GreenLake for File Storage for AI

HPE GreenLake for File Storage offers everything you need to help maximize the vast potential of AI; enterprise performance at AI scale that supercharges workloads across every stage of the data pipeline; radical simplicity at AI scale with an intuitive cloud experience; and enhanced efficiency at AI scale for higher ROI.

HPE 
GreenLake

Here are the top 10 reasons to choose HPE GreenLake for File Storage for AI

1

Accelerate AI workloads with enterprise performance at AI scale

HPE GreenLake for File Storage accelerates your most data-intensive AI applications with enterprise performance that spans all the stages of AI—from data aggregation, data preparation, and training and tuning to inferencing. And it's not just performance that reaches a peak at some point in time for a small dataset. Instead, it's fast, sustained performance that spans the entire scale of your data for the most demanding AI workloads, including generative AI (GenAI) and large language models (LLMs). Enterprise performance at AI scale helps you unlock more value from all your aggregated data—delivering faster time to insight and giving you a real-world competitive advantage.

Because HPE GreenLake for File Storage is certified for NVIDIA BasePOD™ and NVIDIA OVX™ storage validation, customers can confidently use HPE GreenLake for File Storage to run AI applications with the quick deployment enabled by the NVIDIA DGX BasePOD architecture. With an integrated AI infrastructure, organizations can now implement the most impactful, mission-critical AI use cases that NVIDIA DGX BasePOD is designed to address with ease—Health and Life Sciences, Financial Services, Energy, Telecommunications, and Large Language Models (LLMs). By leveraging a proven framework that removes the deployment and operational complexity of AI workloads, you can jump-start your AI journey.

2

Scale without compromise with disaggregated, shared-everything architecture

HPE GreenLake for File Storage has a disaggregated, shared-everything, highly resilient modular architecture that allows you to scale performance and capacity independently—and it's designed for exabyte scale. You can supercharge your most data-intensive AI applications with all-NVMe speed for fast, predictable performance and no front-end caching, data movement between media, or tiered data pipelines. No east-west cluster traffic means every node adds a proportionately linear amount of performance to the cluster, enabling seamless scaling of concurrent users, client nodes, and performance across the entire solution.

3

Optimize GPU utilization

With support for optimized GPU utilization through InfiniBand, NVIDIA® GPUDirect, and RDMA, HPE GreenLake for File Storage accelerates AI workloads by boosting performance for model training and tuning with faster checkpointing. Front-end host InfiniBand connectivity to networks, including the NVIDIA Quantum-2-based InfiniBand network platform, provides flexibility.

NVIDIA OVX servers combine high-performance, GPU-accelerated compute with high-speed storage access and low-latency networking to address a range of complex AI and graphics-intensive workloads. As HPE GreenLake for File Storage has storage validation with NVIDIA OVX systems, customers will fully utilize the compute power of their GPUs when running those systems to create the most complex digital twins of factories, buildings, and even entire cities.

4

Go faster with self-service agility

HPE GreenLake for File Storage simplifies on-premises AI storage with the speed and agility of a cloud operational experience, which makes underlying infrastructure invisible while shifting operations to be workload, not infrastructure, centric. Unlike legacy file solutions with their traditional storage management approach, HPE GreenLake for File Storage empowers line-of-business (LOB) application owners and data scientists to effortlessly create file shares to accelerate AI workloads while freeing up IT resources to work on strategic, higher-value initiatives.



5

Manage from anywhere

Unified management enables you to manage, monitor, and protect your global AI file 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) user experience means everything you need to manage your fleet of data infrastructure is at your fingertips—accessible from anywhere, on any device.

6

Enhanced efficiency at AI scale

HPE GreenLake for File Storage can slash your AI storage costs with high-density and low-power-consumption storage and compute nodes—as well as further lower your carbon footprint with industry-leading data reduction, nondisruptive upgrades, and an AI storage as-a-service consumption model that helps eliminate overprovisioning. Customers can scale up to 720 PB of effective capacity (with 3:1 data reduction) for large enterprise-scale AI file data. You can scale performance and capacity independently for higher efficiency at lower cost, and you can help maximize GPU utilization—and therefore GPU ROI—with enterprise performance at AI scale.

7

Store more for less with best-in-class data reduction

HPE GreenLake for File Storage provides improved data reduction with the Similarity algorithm, which, unlike compression and deduplication, reduces data with both a global and fine-grained approach. Savings are 2:1 for life sciences data; 3:1 for pre-reduced backups, pre-compressed log files, and HPC and animation data; and 8:1 for uncompressed time-series data.¹

8

Gain enhanced resilience with low overhead

With next-generation erasure codes that break the cost/resilience trade-off, HPE GreenLake for File Storage brings the cost of error correction down to just 3% while delivering significantly more resilience than traditional erasure codes.² The new data protection codes provide higher resiliency with no rebuild times for controller failures, resulting in a 99.9999% durable architecture.³

9

Get everything as a service

Modernize your data management with a comprehensive suite of hybrid cloud services on the HPE GreenLake edge-to-cloud platform. Get fast and easy access to a large array of storage services in a single platform—spanning block, file, private cloud, storage fabric management, backup and recovery, and disaster recovery—all delivered everywhere with an intuitive cloud experience.

10

Make AI real in your enterprise

To power your AI initiatives and compete in today's marketplace, you need an AI-ready file storage solution that can deliver enterprise performance, simplicity, and enhanced efficiency—all at AI scale. Truly accelerating your AI workloads requires all three measures of AI scalability, as each one is essential for AI success, and HPE GreenLake for File Storage is designed to deliver across each of them. It's a single, flexible solution that spans not just all the stages of AI, but the full scale of your data. HPE GreenLake for File Storage helps you achieve success wherever you are in your journey to tap the power of AI and make it a reality in your business.



Contact your Connection Account Team for more information.

Business Solutions
1.800.800.0014

Enterprise Solutions
1.800.369.1047

Public Sector Solutions
1.800.800.0019

www.connection.com/HPE

**Hewlett Packard
Enterprise**

© Copyright 2024 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, GPUDirect, and NVIDIA DGX 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.

a50010298ENW, Rev. 1

2746661-0824

^{1, 2, 3} [“HPE Storage substantiation,”](#) HPE, 2023