

HPE Solutions with WEKA



What's new

- HPE Solutions with WEKA now offers the HPE Alletra Storage Server 4210: a 1U 2P all-NVMe server that delivers exceptional performance, high capacity, and thermal optimizations for AI workloads.
- The HPE Alletra Storage Server 4210 delivers performance-enhancing PCIe Gen5 and Enterprise and Data Standard Form Factor (EDSFF) drives with up to 20 E3.S drives per 1U enclosure.
- Support for hybrid NVMe configurations with QLC or pTLC drives for performant and economical all NVMe clusters.
- New tiers for remote WEKA software installation aligned with cluster size and customer requirements, delivered by WEKA Professional Services.
- Get PCIe Gen5 performance and ultra-dense 20 E3.S NVMe in a 1U 1P enclosure with the HPE ProLiant DL325 Gen11, part of the data server family holding seven world records in AI/ML performance.
- Support multiple virtual clusters on the same hardware with multi-tenancy via composable clusters, each fully isolated at the hardware

Overview

Looking for the fastest file performance that can unleash the full power of your GPUs? Discover extreme speed and scalability with NeuralMesh™ by WEKA®, designed to simplify high-performance storage management. Transform stagnant data silos into dynamic pipelines, delivering ultra-fast results wherever speed matters most, across on-prem, cloud, or hybrid environments.

HPE Solutions with WEKA avoids the bottlenecks common with unstructured file and object data through a zero-tuning, zero-copy, highly parallelized architecture. Effortlessly handle even the most complex and diverse AI, machine learning, and high performance computing workloads without IT expertise. Increase GPU performance with an all-NVMe solution that exceeds the demanding performance standards of NVIDIA Cloud Partners. Reach I/O speeds faster than local storage, driving productivity and ROI to new heights.

No compromises—just pure, accelerated potential for your business.

layer with dedicated encryption rules, drives, cores, and memory.

Features

Unlock Maximum Speed with AI-Native Storage

Increase utilization of your valuable GPUs by avoiding the storage bottleneck. HPE Solutions with WEKA delivers scalable, AI-native storage with high throughput, IOPS, and ultra-low latency that meets the rigorous standards of NVIDIA Cloud Partner (NCP) certification.

Slash AI job runtimes with the world's fastest shared file system (SPEC Storage 2020, SPEC SFS 2014, IO-500, and STAC). Process billions of files at unbeatable speeds, whether handling small files or mixed workloads. Reach peak efficiency and I/O with full NVIDIA Magnum IO™ software suite support.

Make data copying and storage silos a thing of the past. Run any part of your pipeline on a single system, regardless of data size, type, or volume. The unique zero-copy, parallel architecture distributes metadata across nodes, making data as fast as local across all types and patterns.

Unleash WEKA's ultra-fast file performance on HPE's world record-holding[1] PCIe Gen5 all-NVMe servers. Optimized for AI/ML scenarios, these servers are powered by the latest processors from AMD and Intel®, including NCP-qualified configurations.

Don't let the network slow down your AI training and inferencing. Get high-speed connectivity of up to 400 Gbps Ethernet or InfiniBand network cards and HPE's fast network switch offerings from NVIDIA and others.

Gain Transformative Scale Wherever Your Data Pipeline Lives

Scale your AI storage to hundreds of petabytes—or even exabytes—across tiers, on-prem, cloud, or hybrid environments. This future-driven, multi-tenant, multi-workload, multi-location platform effortlessly handles billions of files in a single namespace.

Mitigate management complexity with a unified platform where all data sets live in a single namespace. Automatically managed data tiering streamlines file management, sharing, and access, while GPUs gain rapid access to data needed for training.

Easily scale to 512 PB of flash. Extend file system namespace to scale to exabytes with an economical object tier through HPE Solutions with Scality or any S3-compatible object store, public, or private.

Seamless multi-tenancy enables easy management of multiple customers, teams, or projects at desired privacy, security, and performance levels with composable clusters on shared hardware. Scale as priorities evolve without impacting the entire cluster, optimizing costs and management overhead.

Deploy flexibly across multiple environments (bare-metal, containers, native in cloud) with built-in multi-protocol access at no extra charge. Simultaneous NVIDIA GPUDirect Storage (GDS), POSIX, S3, NFS, SMB, and HDFS support enable a single platform to handle all your AI data lifecycle needs.

Simplify Operations with Fast File Storage Built for Enterprise

Take the complexity out of fast file storage. Management is made easy, requiring no tuning or storage specialists. Modernize with AI/ML-optimized hardware offering superior quality and security. With a full best-in-class solution, intensive administration is history—just set it and forget it.

Auto-adapt to any workload instantly. Built-in intelligent data management helps ensure peak performance without manual tuning. Effortlessly configure, monitor, and administer clusters from one place with WEKA Home's full-featured, intuitive console—no IT expert needed.

Enable fast, secure data access for multiuser enterprises. Protect against rogue actors with end-to-end encryption (at-rest and in-transit), authentication, and automatic cloud backups. The unique snap-to-object feature for disaster recovery safeguards your data without performance loss.

Protect your AI workloads with best-in-class servers and availability. High-load servers with built-in redundancy and failovers reduces downtime, while silicon root of trust anchors firmware. Advanced monitoring and behavioral analytics deliver robust end-to-end IT protection.

Technical specifications	HPE Solutions with WEKA
Protocol supported	WEKA client software is implemented using a POSIX file system driver, which creates access similar to using local storage for each client for AWS S3, object storage, HDFS, GPUDirect Storage, and NFS v3 SMB.

[1] [intel.com/content/www/us/en/gaming/resources/what-is-pcie-4-and-why-does-it-matter.html#:~:text=Each%20generation%20of%20PCIe%20is%20twice%20as, speed%20before%20encoding%E2%80%94realized%20speeds%20may%20be%20slower](https://www.intel.com/content/www/us/en/gaming/resources/what-is-pcie-4-and-why-does-it-matter.html#:~:text=Each%20generation%20of%20PCIe%20is%20twice%20as, speed%20before%20encoding%E2%80%94realized%20speeds%20may%20be%20slower)

HPE GreenLake

[HPE GreenLake edge-to-cloud platform](#) is HPE's market-leading as-a-Service offering that brings the cloud experience to apps and data everywhere – data centers, multi-clouds, and edges – with one unified operating model, on premises, fully managed in a pay per use model.

If you are looking for more services, like **IT financing solutions**, please [explore them here](#).

[For additional technical information, available models and options, please reference the QuickSpecs](#)

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

[Chat now](#)

© Copyright 2026 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.

Parts and Materials: HPE will provide HPE-supported replacement parts and materials required to maintain the covered hardware.

Parts and components that have reached their maximum supported lifetime and/or the maximum usage limitations as set forth in the manufacturer's operating manual, product quick-specs, or the technical product data sheet will not be provided, repaired, or replaced as part of these services.

AMD is a trademark of Advanced Micro Devices, Inc. Intel is a trademark of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. GPUDirect and NVIDIA are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Red Hat is a registered trademark of Red Hat Inc. in the United States and other countries. All third-party marks are property of their respective owners.

Image may differ from the actual product.

[PSN1013033917ROEN](#), March, 2026.

HEWLETT PACKARD ENTERPRISE

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

