


ZettaLane Systems Delivers World-Class AI Storage Performance in MLPerf® Storage v3.0 Benchmark on Standard Cloud VMs

MayaNAS delivers AI/HPC-scale performance with NAS simplicity, running Lustre on cloud object storage for the full AI pipeline, sovereign in your own cloud.

SANTA CLARA, CA, UNITED STATES, September 1, 2026 /EINPresswire.com/ -- ZettaLane Systems, LLC — a provider of software-defined file and block storage for AI, deployed in the customer's own cloud — today published the results of its first [MLPerf® Storage v3.0](#) submission and introduced MayaNAS, a groundbreaking storage engine that unifies a familiar enterprise NAS with a high-performance parallel file system, running directly on cost-efficient cloud object storage. Among all submissions in the round, MayaNAS was the only one to run the [Lustre parallel file system](#) directly on a cloud object-storage tier.

For years, enterprise file storage and HPC-class parallel storage have lived in separate worlds — separate products, separate infrastructure, separate budgets. MayaNAS collapses that divide. From a single software engine it serves the everyday protocols teams rely on — NFS and SMB with Active Directory — alongside the parallel file systems that feed GPUs, Lustre and pNFS Flex Files, over the very same data. It is the NAS teams already know, that also delivers GPU-class parallel throughput — with bulk capacity written straight to cloud object storage, so that performance arrives at object-storage economics rather than premium block-storage prices.

In MLPerf® Storage v3.0, run entirely on standard Google Cloud virtual machines, ZettaLane submitted CLOSED-division results across the full AI data pipeline with two cloud-native engines — MayaNAS, the object-backed parallel file system, and MayaScale, a companion NVMe-over-TCP block engine. Every result ran on a deliberately compact footprint driven to full saturation: with 200 Gbps client networking fully utilized and no cores or bandwidth left idle, the numbers reflect efficiency per client and per node.



ZettaLane announces MLPerf® Storage v3.0 benchmark results: MayaNAS runs the Lustre parallel file system directly on cloud object storage — the only such submission — alongside the companion MayaScale block engine, on standard Google Cloud VMs.



MayaNAS erases the line where traditional cloud NAS stops and HPC begins — our innovation fuses the Lustre parallel file system with object storage, in your own cloud."

Supramani Sammandam

Highlights:

- Parallel throughput on object storage: MayaNAS sustained 32.42 GB/s checkpoint write and 21.20 GB/s read on Llama 3 70B with just two clients — the file system's entire data path on cloud object storage, with no local scratch tier.
- Full saturation from a single client: one MayaScale client, on 200 Gbps networking, fully fed 72 B200-class accelerators on RetinaNet at 86.68% utilization — a high

per-client accelerator density that saturated the client network with no idle capacity.

- GPUs kept fed, not waiting: 3D U-Net training held approximately 92% accelerator utilization on MayaNAS and 93.31% on MayaScale — storage supplied data continuously rather than starving the accelerators.
- One substrate, full pipeline: the same platform also served a Llama 3.1 8B inference cache at 524.65 tokens per second, and checkpointed Llama 3 8B at 14.43 GB/s write / 10.40 GB/s read from a single client.

For MayaNAS, each OST is an [OpenZFS](#) dataset derived from regional, Standard-class Google Cloud Storage buckets — bulk data read and written directly to object storage, with a small NVMe device holding only metadata — so its entire data path runs on cloud object storage. Together, MayaNAS and MayaScale cover the full AI data pipeline from one vendor: file at object-storage economics, high-performance block on local NVMe, both on standard cloud VMs.

"Traditional cloud NAS stops exactly where HPC and AI begin, and teams have had to buy a second storage system to cross that line," said Supramani Sammandam, Founder of ZettaLane Systems. "MayaNAS erases that line by fusing the parallel file system with object storage — Lustre's HPC performance running directly on cloud object storage, inside your own cloud. MLPerf® Storage v3.0 proves it holds up under real AI workloads."

MLPerf® Storage v3.0 introduced support for an S3 object-storage access layer alongside its established POSIX-compliant layer, a clear signal of object storage's rising role in the AI data pipeline. Rather than treat object storage as an alternative to the parallel file system, MayaNAS fuses the two — a complete POSIX parallel file system (Lustre) whose data path runs directly on object storage, so existing applications need not be rewritten.

"MLPerf® Storage gains depth and relevance every time a new organization joins the community," said David Kanter, Head of MLPerf® at MLCommons®. "ZettaLane's first-time submission to v3.0 adds fresh data points that help paint a fuller picture of storage performance under real-world training workloads."

Both engines run entirely inside the customer's own cloud account — the file system, buckets, and encryption keys never leave it. That makes MayaNAS and MayaScale sovereign by design: AI-scale storage that meets the strictest data-residency and sovereignty requirements, at object-storage economics, with the data and the AI pipeline built on it staying in the customer's own cloud. Both run across Google Cloud, Microsoft Azure, and Amazon Web Services, and deploy as code through the open-lustre-cloud project on GitHub. It is part of ZettaLane's storage family alongside MayaScale, its high-performance NVMe-over-TCP block engine.

About ZettaLane Systems, LLC

ZettaLane Systems builds storage that works the way AI teams already do: parallel file (Lustre and NAS) on object storage for training, NVMe-over-TCP block for databases and Kubernetes, and branchable Postgres for AI agents — all deployed with Terraform into the customer's own account on Google Cloud, AWS, and Microsoft Azure, with active-active high availability. ZettaLane is a Google Cloud Select Technology Partner and a member of the NVIDIA Inception Program.

MLPerf® Storage v3.0, Closed division; submitted by ZettaLane Systems on Google Cloud. Retrieved from <https://mlcommons.org/benchmarks/storage/> on September 1, 2026; entries 3.0-0136 through 3.0-0141. Result verified by MLCommons Association. The MLPerf name and logo are registered and unregistered trademarks of MLCommons Association in the United States and other countries. All rights reserved. Unauthorized use strictly prohibited. See www.mlcommons.org for more information. Google Cloud is a trademark of Google LLC.

ZettaLane Public Relations

ZettaLane Systems, LLC

+1 408-732-8000

pr@zettalane.com

This press release can be viewed online at: <https://www.einpresswire.com/article/938591689>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.