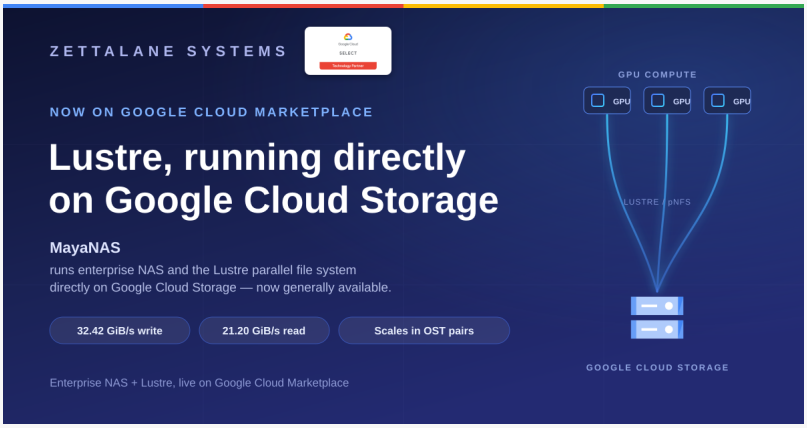
ZettaLane Systems Launches MayaNAS on Google Cloud Marketplace: Lustre Running Directly on Google Cloud Storage

MayaNAS runs enterprise NAS and the Lustre parallel file system from one engine on standard Google Cloud infrastructure, with bulk data in Google Cloud Storage.



MayaNAS brings the Lustre parallel file system directly to Google Cloud Storage, now generally available on Google Cloud Marketplace

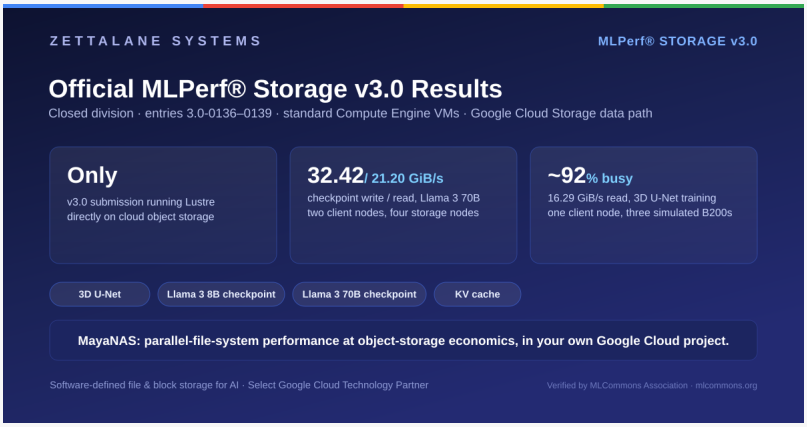
SANTA CLARA, CA, UNITED STATES, October 14, 2026 /EINPresswire.com/ -- ZettaLane Systems, LLC, a Select Google Cloud Technology Partner, today announced that MayaNAS, the only MLPerf® Storage v3.0 submission verified running the Lustre parallel file system directly on cloud object storage, is available on Google Cloud Marketplace. MayaNAS presents standard enterprise NAS (NFS, and SMB with Active Directory) alongside the parallel file systems Lustre and pNFS, with bulk data stored in Google Cloud Storage.

> ZettaLane can now securely scale and support enterprises using MayaNAS to deliver high-performance, shared storage for demanding workloads, including AI training."
>
> *Dai Vu, Managing Director, Marketplace & ISV GTM Programs, Google Cloud*

In MLPerf® Storage v3.0, run entirely on standard Compute Engine instances with Google Cloud Storage, MayaNAS sustained 32.42 GiB/s checkpoint write and 21.20 GiB/s read on the Llama 3 70B workload with two client nodes and four storage nodes. On the 3D U-Net training workload, a single client node kept three simulated NVIDIA B200 accelerators about 92% busy. The result shows that one file system on standard Google Cloud infrastructure can serve both enterprise file workloads and the most demanding parallel ones, with Google Cloud Storage as the storage layer.

The approach has drawn independent attention in the storage press. Blocks & Files wrote:

"Here's an interesting idea; ZettaLane has built MayaNAS, a Lustre parallel filesystem, with NFS, pNFS and SMB support, based on cloud object storage." Storage Newsletter called it a "pretty innovative approach."



Official MLPerf® Storage v3.0 results for MayaNAS on Google Cloud, verified by MLCommons Association.

Most cloud NAS stops where HPC and AI begin. On Google Cloud, MayaNAS lets teams scale from everyday file shares to GPU-class parallel throughput without deploying a second storage system. Bulk data lives permanently in Google Cloud Storage, so there are no tiering or staging jobs between a fast tier and the bucket. Capacity grows with the bucket, and teams pay for performance only while they need it: serving nodes scale up for a run and back down afterward while the data stays in place. The file system, buckets and encryption keys remain inside the customer's own project, supporting data-residency and sovereignty requirements. MayaNAS runs active-active for high availability and uses the standard open-source clients already in Linux, with no proprietary software to install.

Beyond AI, the same parallel throughput serves genomics and life-sciences pipelines, electronic design automation, media rendering, financial risk analysis and seismic processing. MayaNAS is built for enterprise IT, research institutions, regulated industries and sovereign AI initiatives that need parallel file performance at object-storage scale while keeping full control of their data.

"Cloud object storage scales without limits and handles durability on its own. We built MayaNAS to use it directly as the storage layer beneath a parallel file system, so teams get the enterprise file services they already run and the parallel Lustre and pNFS throughput their GPUs need, with no disks, RAID or separate fast tier to manage," said Supramani Sammandam, Founder of ZettaLane Systems. "MLPerf® Storage v3.0 showed that standard Compute Engine instances and Google Cloud Storage make a powerful combination for serving the AI data pipeline, from training and checkpointing to inference cache, with the data staying in the customer's own project."

"Bringing ZettaLane Systems to Google Cloud Marketplace will help customers quickly deploy, manage, and grow the company's MayaNAS storage platform on Google Cloud's trusted, global infrastructure," said Dai Vu, Managing Director, Marketplace & ISV GTM Programs at Google Cloud. "ZettaLane can now securely scale and support enterprises using MayaNAS to deliver high-performance, shared storage for demanding workloads, including AI training."

Customers can adopt MayaNAS the way they already buy cloud: pay-as-you-go on Google Cloud Marketplace with a free trial, deployed in minutes into their own VPC and billed through their

Google Cloud account, with eligible purchases counting toward existing Google Cloud commitments; through private offers for negotiated terms, directly or via Google Cloud reseller partners; or as code, using ZettaLane's open-source Terraform modules, which deploy the same Marketplace image and also cover pNFS, for teams that run their infrastructure as code. To start an evaluation or plan a deployment, contact sales@zettalane.com. MayaNAS is part of ZettaLane's storage family for AI data and agentic workloads. Learn more at zettalane.com.

About ZettaLane Systems, LLC
ZettaLane Systems builds software-defined storage that brings parallel-file-system performance to cloud object storage. Its MayaNAS and MayaScale products deliver multi-protocol file (NFS, SMB), high-performance block (NVMe-over-TCP), and object-backed storage across Google Cloud, AWS, and Microsoft Azure, with active-active high availability and deployment in minutes. Zettabranch gives every pull request, preview deployment and coding agent its own production-like Postgres branch in the customer's own cloud (zettabranch.ai). ZettaLane is a Select Google Cloud 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 3.0-0139 (MayaNAS: 3D U-Net, Llama 3 8B and 70B checkpoint, KV cache). 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, Google Cloud Storage and Compute Engine are trademarks of Google LLC. NVIDIA and B200 are trademarks of NVIDIA Corporation. All other trademarks are the property of their respective owners.

ZettaLane Public Relations
ZettaLane Systems, LLC
+1 408-732-8000
pr@zettalane.com
Visit us on social media:
LinkedIn
X

This press release can be viewed online at: https://www.einpresswire.com/article/948595979

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.