Cloudian Powers Next-Generation AI Storage with NVIDIA and Lenovo Reference Architecture

Advanced storage platform enables seamless performance and capacity growth from terabytes to exabytes for AI model training and inference workloads.

SAN MATEO, CA, UNITED STATES, March 18, 2025 /EINPresswire.com/ -- [Cloudian](#) announced today an enterprise storage solution for artificial intelligence that combines HyperStore object storage with Lenovo's latest AI-optimized servers in a comprehensive NVIDIA-validated architecture. The platform delivers throughput of 20GB/s per node with linear scaling capabilities, meeting the intense I/O demands of modern AI workloads including large language model training, inference operations, and checkpointing functions.

> This validated architecture unites Lenovo's GPU-optimized systems and flash-optimized storage with Cloudian's enterprise object storage to create a comprehensive AI infrastructure solution."
>
> *Stuart McRae, Lenovo*

As organizations increasingly rely on machine learning and deep learning models, object storage has emerged as a foundational technology for AI infrastructure. Cloudian leads this evolution with full S3 API compatibility, enterprise-grade security features, and enhanced performance through [NVIDIA GPUDirect](#) Storage (GDS) integration, enabling organizations to maximize their AI investments.

The validated architecture demonstrates key capabilities:
⬜ Flash-optimized performance utilizing NVIDIA GPUDirect Storage and RDMA networking with linear throughput scaling
⬜ Enterprise-class features including observability, self-healing capabilities, multi-tenant support, QoS controls, dynamic storage provisioning, and data lifecycle management

Comprehensive security framework with enterprise-grade certifications, built-in ransomware protection, and standard OS kernel implementation

"Today's AI storage requirements demand both extreme performance and enterprise-grade data management," said Amit Rawlani, Sr. Director of Alliances at Cloudian. "Our solution transforms AI infrastructure by combining HPC-level performance with the scalability and efficiency of S3-compatible object storage, while delivering the enterprise features organizations require."

"This validated architecture unites Lenovo's GPU-optimized systems and flash-optimized storage infrastructure with Cloudian's enterprise object storage to create a comprehensive AI infrastructure solution," said Stuart McRae, GM of Data Storage at Lenovo. "This solution is well suited for the dynamic workload requirements of modern enterprises."

The reference architecture enables organizations to move quickly from proof-of-concept to enterprise-wide AI deployments. By delivering native S3 compatibility, enterprise-grade security, and a scalable architecture, the platform adapts to growing AI requirements while maintaining consistent performance and centralized management. Organizations can start with a targeted deployment for specific workloads, then expand capacity and performance as business needs evolve.

Download the reference architecture document here.

About Cloudian
Cloudian is the leader in secure, S3-compatible AI data lake platforms. With enterprise-grade security, limitless scale, and seamless cloud integration, Cloudian delivers plug-and-play interoperability with NVIDIA GPUDirect and with AI tools such as PyTorch, TensorFlow, Kafka, Arrow and more. Our software-defined, on-premises solutions let users simplify and accelerate AI workflows, meet data sovereignty requirements, and cut costs by consolidating information to a single, cloud-like environment. Cloudian's geo-distributed architecture manages and protects object and file data at the edge, core, and in the cloud, for both traditional and modern applications.

Jon Toor
Cloudian
email us here

This press release can be viewed online at: https://www.einpresswire.com/article/793957060