

Cloudian Launches High-Performance S3-Compatible AI Infrastructure Built on NVIDIA AI Data Platform

Advanced AI-ready solution simplifies deployment of full-stack platforms that bring accelerated computing and storage networking to S3-compatible storage

SAN MATEO, CA, UNITED STATES,

September 30, 2025 /

EINPresswire.com/ -- [Cloudian](#), the

leader in enterprise object storage solutions, today announced the launch

of the Cloudian HyperScale® AI Data

Platform, a revolutionary AI-ready solution built on S3 API storage that

maintains full compatibility with existing data sets and applications while seamlessly scaling to exabyte capacity. The platform supports the S3 storage protocol over remote direct memory access (RDMA) to deliver enhanced object storage performance and efficiency for AI workloads.



“

Cloudian's HyperScale AI Data Platform integrates NVIDIA AI data platform technologies with high-performance S3-compatible storage to help enterprises unlock the value of their unstructured data.”

*Justin Boitano, vice president,
enterprise AI products at
NVIDIA*

The Cloudian HyperScale AI Data Platform addresses the critical challenge enterprises face as they embrace agentic and generative AI: making their vast repositories of unstructured data instantly accessible and actionable for AI applications. Built on the [NVIDIA](#) RTX PRO 6000 Blackwell Server Edition GPU and aligned with the NVIDIA AI Data Platform reference design, the HyperScale AI Data Platform enables organizations to tap into the 90% of enterprise data that is unstructured—the reports, manuals, presentations, and multimedia content that represent institutional knowledge.

Eliminates Complex Infrastructure Requirements

Traditional AI implementations require enterprises to build

complex, dedicated file system structures and separate vector databases to achieve optimal

performance—a costly and time-intensive process that demands specialized skills many organizations lack. Many enterprise IT teams don't have the expertise to build and maintain these separate infrastructure layers, forcing them to either invest heavily in new talent or accept suboptimal AI performance.

Cloudian eliminates these barriers by enabling enterprises to deploy AI applications directly on their native S3-compatible data sources without requiring additional file system layers or separate vector database infrastructure. This approach dramatically reduces complexity, accelerates deployment timelines, and allows organizations to leverage their existing IT expertise and infrastructure investments.

Transforms Enterprise Data into AI-Ready Intelligence

The solution brings AI compute directly to enterprise data stores, automatically transforming storage into AI-ready infrastructure. With integrated vector database capabilities built into the storage layer, the Cloudian HyperScale AI Data Platform automatically ingests, embeds, and indexes multimodal unstructured content, making it instantly searchable through vector search interfaces and immediately available for retrieval-augmented generation (RAG) workflows.

This intelligent approach ensures AI agents throughout the organization can access, search, and analyze information in near real-time, thus eliminating traditional data silos and enabling businesses to excel in the agentic AI era where artificial intelligence systems require instant access to complete organizational knowledge.

The platform's S3-over-RDMA capability, developed in collaboration with NVIDIA, further enhances this performance by delivering increased read and write throughput while optimizing AI storage access efficiency.

"The HyperScale AI Data Platform represents a fundamental shift in how enterprises approach AI readiness," said Neil Stobart, CTO at Cloudian. "Rather than forcing enterprises to build complex, separate infrastructure for AI—which requires skills most don't possess—we've engineered a platform that automatically transforms existing data into actionable intelligence at the storage layer. This enables enterprises to unlock the full potential of their institutional knowledge while maintaining the security and governance standards that modern businesses demand."

"Enterprise AI success depends on eliminating data movement bottlenecks and bringing compute directly to where data resides," said Justin Boitano, vice president, enterprise AI products at NVIDIA. "Cloudian's HyperScale AI Data Platform integrates NVIDIA AI data platform technologies with high-performance S3-compatible storage to help enterprises unlock the value of their unstructured data and deploy advanced agentic AI solutions at scale."

Key Solution Capabilities

Eliminates Infrastructure Complexity: Built-in vector database capabilities eliminate the need for enterprises to architect and manage separate file systems and vector database infrastructure, removing barriers that typically require specialized expertise.

Automated AI Readiness: Data is instantly classified, enriched with real-time metadata, and vectorized as it's stored, eliminating manual data preparation processes and complex infrastructure provisioning.

Native S3 API: Full interoperability with the vast ecosystem of AI tools written for the S3 API and immediate access to data already stored in the S3 format, the de facto standard for cloud-based information.

Performance Excellence: RDMA networking provides S3 object storage with direct data paths between storage, system memory, and GPU memory for ultra-low latency and up to 35 GB/s per node (reads), scalable to TBs per second.

Proven Scale: Built on Cloudbian's proven exabyte-scale platform, the solution supports massive vector datasets with consistent high-performance access, enabling organizations to grow their AI capabilities without architectural limitations.

Data Sovereignty: On-premises infrastructure maintains complete enterprise control over sensitive data with robust security, enterprise-grade access controls, and governance policies. Organizations retain full ownership and control of their data while enabling AI innovation.

Unified Architecture: Eliminates the complexity and cost of managing separate data, metadata, and vector database storage systems by integrating all capabilities into a single, enterprise-ready platform.

Enterprise-Ready AI Infrastructure for the Agentic Era

By bringing accelerated computing and integrated vector database capabilities directly to enterprise data storage, the HyperScale AI Data Platform eliminates the need for enterprises to build and manage complex, separate AI infrastructure. This unified approach removes the technical barriers that have prevented organizations from deploying AI at scale, enabling rapid deployment using existing IT skills and infrastructure.

Key solution components include NVIDIA RTX PRO 6000 Blackwell GPUs, NVIDIA BlueField DPUs, and NVIDIA AI Enterprise software — including NVIDIA NIM and NVIDIA NeMo Retriever microservices. NVIDIA Spectrum-X Ethernet networking delivers breakthrough performance with RDMA that provides direct data paths between storage, system memory, and GPU memory. This architecture achieves ultra-low latency and delivers up to 35 GB/s per node (reads), scalable to terabytes per second.

The solution enables enterprises to deploy AI agents that can immediately understand, query, and reason over their complete business knowledge base, positioning them for success as agentic AI becomes central to enterprise operations. Organizations can now build AI applications that securely access enterprise data with reduced overhead and faster time to deployment, all

while maintaining full control over their most valuable asset—their data.

Jon Toor
Cloudian
[email us here](#)

This press release can be viewed online at: <https://www.einpresswire.com/article/852873371>

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.

© 1995-2025 Newsmatics Inc. All Right Reserved.