

# Cloudian Expands HyperScale AIDP with Broader Data Reach, End-to-End Security, and Multi-OEM Validation

*Turnkey on-premises AI appliance delivers data sovereignty and 60% cost savings versus public cloud AI*

SAN MATEO, CA, UNITED STATES, June 9, 2026 /EINPresswire.com/ -- [Cloudian](#), Inc. today announced HyperScale® AIDP v1.1, an update to its turnkey on-premises AI appliance that gives enterprises a sovereign, cost-effective alternative to public cloud AI services. Delivered as a pre-integrated system, HyperScale AIDP lets customers run production AI on their own infrastructure — keeping sensitive data under their direct control and saving 60% on cost by avoiding the recurring token, egress, and inference fees that make public cloud AI expensive at scale.



Cloudian HyperScale AIDP Integrated AI Appliance

The v1.1 release introduces three enhancements aimed at the practical barriers enterprises encounter when moving AI projects from pilot to production.

Production-ready NVIDIA Blueprints: HyperScale AIDP v1.1 supports NVIDIA AI Blueprint for [Enterprise Document RAG](#) and NVIDIA [Metropolis Blueprint](#) for video search and summarization (VSS). Enterprise Document RAG transforms an organization's existing volumes of contracts, policies, technical documentation, and reports into a conversational knowledge base that returns precise answers with source citations. Metropolis VSS Blueprint extends the same capability to video, surfacing insight from live and recorded visual sensors, training content, broadcast archives, and inspection video. VSS automatically indexes new footage as it arrives rather than requiring scheduled batch processing, a change that makes the feature practical for security/compliance monitoring, quality control inspection, and other use cases where video accumulates in real time. Both blueprints run natively on HyperScale AIDP's S3-native storage and built-in Milvus vector database, eliminating the integration work that typically stalls enterprise AI initiatives.

Ingest from NFS file and object data sources: Most enterprise unstructured data lives across a mix of file shares and object stores accumulated over years. HyperScale AIDP v1.1 ingests natively from both, eliminating the need for migration or consolidation projects in front of every AI initiative. To enhance storage efficiency, HyperScale AIDP indexes and vectorizes content without making a second copy, storing only the resulting embeddings. Customers add AI capability without doubling their enterprise storage footprint.

End-to-end security: HyperScale AIDP enforces user- and group-level access controls throughout the AI pipeline — ingest, indexing, retrieval, and response — so each query returns only the content the requesting user is authorized to see. Those same controls extend to the vector database itself, blocking unauthorized access to the embeddings — an important safeguard because exposed vectors can be inverted to reconstruct the underlying source data. Sensitive contracts, personnel records, regulated data, and proprietary research remain protected as they become part of active AI workflows, making the platform deployable in regulated industries and sovereign environments without creating new exposure paths.

HyperScale AIDP is validated on NVIDIA-Certified GPU server platforms from both Supermicro and Lenovo, giving customers flexibility to align AI deployments with their existing hardware standards, support relationships, and procurement frameworks. Dual-vendor validation also reduces supply and continuity risk for an infrastructure investment customers expect to scale for years.

“Enterprise AI succeeds or fails on data access,” said Neil Stobart, CTO of Cloudian. “Customers tell us their projects stall not because the models lack capability, but because the data those models need is scattered across file and object silos, or because opening up that data would violate internal access policies. HyperScale AIDP v1.1 addresses both barriers directly. It reaches across the data estates customers already have, and it carries permissions all the way through to inference — so the right people get the right answers, and nothing more.”

“Scaling AI from pilot projects to production requires addressing the fundamental challenges of data access, security and control over distributed data estates,” said Jason Hardy, vice president of storage technology, NVIDIA. “By adopting NVIDIA AI Blueprints and running on NVIDIA-Certified platforms, Cloudian’s HyperScale AIDP offers a solution for enterprises to accelerate their path from data to insight while preserving architectural flexibility.”

#### Availability

HyperScale AIDP v1.1 is available now through Cloudian and its global partner network on NVIDIA-Certified server platforms from Supermicro and Lenovo. For more information, visit [cloudian.com/hyperscale-aidp](https://cloudian.com/hyperscale-aidp).

#### About Cloudian

Cloudian is the leading provider of S3-native object storage, purpose-built for the AI era. HyperStore delivers breakthrough performance, exabyte scalability, and government-verified

security within a unified architecture that eliminates data silos and accelerates the most demanding AI workloads. As an NVIDIA-Certified Storage solution, Cloudian integrates natively within NVIDIA AI Factory deployments—providing the high-throughput, low-latency data foundation for GenAI, RAG, and large-scale model training at any scale. The HyperScale AI Data Platform extends that foundation into a complete enterprise AI solution. From government agencies to global research organizations, the world's most demanding data teams trust Cloudian. For more information, visit [www.cloudian.com](http://www.cloudian.com).

Jon Toor  
Cloudian  
+1 650-227-2380  
[email us here](#)

---

This press release can be viewed online at: <https://www.einpresswire.com/article/917377425>

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.