

Part 2:

Data Storage Solutions for Enterprise AI

Kashish Karnick
Sr. Product Manager | Asia Pacific
Lenovo Data Storage Solutions

Accelerating AI Adoption with Lenovo AI Express

Building a successful AI proof-of-concept is only the first step.

Turning that proof-of-concept into a production-ready AI platform often introduces challenges around infrastructure design, deployment timelines, software integration, and operational complexity.

To simplify this journey, Lenovo introduced **Lenovo AI Express**, a portfolio of validated AI Factory offerings designed to make enterprise AI easier to start, easier to price, and faster to deploy. AI Express provides curated hardware, software, services, and support options that help organizations move from concept to production more quickly.

Lenovo AI Express is built around three simple principles:

- Easy to Start
- Easy to Price
- Easy to Ship

The portfolio includes solutions that scale from focused AI inferencing environments to larger AI Factory deployments.

Rather than spending months assembling infrastructure from individual components, organizations can begin with validated architectures that reduce deployment risk and accelerate time to value.

The image below illustrates the progression from compact inferencing platforms through larger-scale AI deployments.

The AI Express small, medium, and large configurations are designed to align with different stages of enterprise AI adoption, from inferencing and Agentic AI to large-scale generative AI environments.

Lenovo AI Express: Fast Lane to the Hybrid AI Factory

<i>Small</i>	<i>Medium</i>	<i>Large</i>
		
ThinkSystem SR650a V4 Flexible 2U AI Inferencing with 2× NVIDIA RTX PRO 6000	ThinkSystem SR675 V3 Versatile 3U Large AI Inferencing With 4/6/8× RTX PRO 6000	ThinkSystem SR680a V4 Maximum-density Generative AI Base 1× NVIDIA HGX B300 NVL8

A Practical Foundation for Enterprise AI

The Lenovo Hybrid AI 221 and 241 platforms provide a compact, enterprise-ready foundation for inferencing, RAG, and Agentic AI workloads.


Unlike massive AI training environments that operate across thousands of GPUs and petabytes of data, these platforms are designed to bring AI closer to enterprise information and business processes.

This is where the **Lenovo ThinkSystem DM and DG portfolio** becomes a critical part of the solution.

While compute powers AI reasoning, the ThinkSystem DM and DG platforms provide access to the business knowledge that makes AI useful.

Organizations can start with entry-level systems and scale across the DM and DG portfolio as performance requirements, data volumes, and AI adoption grow, while continuing to use the same ONTAP data management framework.

This creates a consistent foundation for both today's RAG deployments and tomorrow's Agentic AI environments.



Why Unified Storage Matters for RAG and Agentic AI

Enterprise knowledge rarely resides in a single location.

Engineering documentation may live on NFS shares. Business documents may be stored on SMB file systems. Applications may use block storage. Archives and unstructured content increasingly reside in object storage.

A RAG system often needs information from all of these sources.

An AI agent may need not only to find information but also interact with business systems, validate records, update workflows, and execute tasks across them.

Managing separate storage platforms for file, block, and object workloads introduces unnecessary complexity and creates barriers between AI and the information it needs.

The ThinkSystem DM and DG platforms address this challenge through ONTAP's **unified architecture**.

Organizations can access:

- File data through NFS and SMB
- Block data through iSCSI, NVMe/TCP, NVMe/FC, and Fibre Channel
- Object data through ONTAP S3

From a business perspective, **unified storage** means less time moving data between silos and more time turning enterprise knowledge into business outcomes.

It allows organizations to bring AI closer to their data instead of constantly moving data closer to AI.

Turning Data Into an AI Advantage


The real role of storage in AI is not simply storing information.

Its role is enabling organizations to convert information into an operational advantage.

Enable Faster AI Innovation

AI teams continuously experiment with embeddings, vector databases, retrieval methods, prompts, and agent workflows.

ONTAP helps organizations create rapid, space-efficient copies of data sets, allowing teams to innovate without disrupting production environments or duplicating massive amounts of data.




A Snapshot copy records only changes to the active file system since the last Snapshot copy.

Protect Business Knowledge

As organizations become increasingly dependent on AI-driven processes, protecting the information behind those processes becomes critical.

Replication, backup, recovery, and data protection capabilities help ensure that the knowledge powering AI remains available even during outages, failures, or cyberattacks.



Manage AI Data Growth

AI environments create more than source documents.

They also generate embeddings, indexes, logs, audit records, and agent-generated outputs.

Intelligent tiering, storage efficiencies, and lifecycle management help organizations scale sustainably without continuously expanding premium storage capacity.

Taken together, these capabilities help transform enterprise data from a static repository into a dynamic foundation for AI innovation.

Protecting and Governing AI Data

As organizations move AI from experimentation into production, security and governance become just as important as performance.


The data used to ground AI responses often contains intellectual property, customer records, engineering know-how, financial information, operational procedures, and business-critical processes.

Protecting that information is essential to maintaining trust in AI outcomes.

Protect AI Data

As AI becomes integrated into business operations, compromised data can lead to inaccurate decisions, failed processes, and business disruption.

ONTAP Autonomous Ransomware Protection helps detect suspicious behavior and create protective recovery points when attacks are suspected.



Capabilities such as multi-admin verification help reduce operational risk by requiring additional authorization before critical administrative actions can be performed.

Secure AI Data

AI systems frequently interact with sensitive information.

The ThinkSystem DM and DG family supports:

- Role-based access control
- Directory integration
- Encryption at rest
- Encryption in transit
- Secure administration
- Tenant-level isolation

These capabilities help ensure that only authorized users and applications can access sensitive information.

Govern AI Data

As Agentic AI becomes more common, governance becomes increasingly important.

Organizations need confidence in how information is retained, protected, replicated, audited, and recovered.

ONTAP provides capabilities that support retention policies, immutability, backup, disaster recovery, and data management requirements, helping organizations maintain control of the information that powers AI.

Looking Beyond Today's AI Models

Five years from now, the winners will not be the companies with the largest GPU clusters.

They will not be the companies using the latest AI model.

Those advantages are temporary.

New models will emerge.

Existing models will improve.

Costs will fall.

Capabilities will become more widely available.

What will remain valuable is the ability to connect AI to proprietary data, business processes, and institutional knowledge.

Organizations that establish a trusted, governed foundation for enterprise data today will be able to adopt future models, future AI platforms, and future agentic systems without rebuilding their strategy every time technology changes.

The future advantage is not owning the model.

The future advantage is owning the data that makes the model useful.

Conclusion

Successful AI projects are not built on models alone.

They are built on trusted business knowledge.

The Lenovo Hybrid AI 221 and 241 platforms provide the compute foundation for enterprise inferencing, RAG, and Agentic AI workloads. Lenovo AI Express extends this vision by delivering validated AI Factory configurations that help organizations move from concept to production more quickly.

The Lenovo ThinkSystem DM and DG product family provides the unified data foundation that makes these AI platforms useful. By bringing file, block, and object data together under a single ONTAP architecture, organizations can connect AI directly to the information that powers their business while maintaining the security, governance, and operational simplicity that enterprise environments require.

Together, these solutions help organizations prepare not just for today's AI use cases, but also for the next generation of intelligent agents and business automation.

AI models will come and go.

Business knowledge compounds over time.

Own Your Data. Rent The AI.

