

## Dell Technologies Turns Enterprise Data Into Trusted Context for AI Agents

October 6, 2026

*Dell AI Data Platform, the data foundation of the Dell AI Factory, expands with new data orchestration and storage capabilities*

- Unified Semantic Layer and Enterprise Knowledge Graph give every application and agent the same trusted view of enterprise data
- Dell AI Data Platform shortens the wait between question and insight with accelerated data processing that processes data nearly 4 times faster on average than CPUs alone
- Dell AI Data Platform storage enhancements improve security and benchmark performance
- New AI-ready data services move organizations from deployment to production quickly and reliably

ROUND ROCK, Texas--(BUSINESS WIRE)--Oct. 6, 2026-- Dell Technologies (NYSE: DELL) expands the [Dell AI Data Platform](#) so AI systems deliver better answers, cost less to run and reach production faster, giving applications and agents trusted, current context from a company's own data.

### Why it matters

In most cases the model isn't the bottleneck to achieving enterprise scale for AI. The data is. A company's most useful information sits in files, databases, cloud services and systems across multiple sites. Most of it was created long before anyone expected an AI agent to read it. It was never labeled, connected or organized with a machine reader in mind. As a result, agents burn tokens reconstructing answers that should already exist. Every query means more steps, more compute, more cost and slower, less trustworthy results.

### Giving data meaning and context

The Dell AI Data Platform is built end to end, from where data is stored to where agents use it. Storage moves critical data to accelerated compute without bottlenecks; orchestration connects and prepares it; analytics and search transform and index it. NVIDIA Nemotron Retriever models enable document parsing, embedding and reranking. NVIDIA cuVS accelerates vector indexing and search. Because every layer runs on one platform, everything it produces traces back to governed data the customer already owns. Agents return to the same documents, history and decisions repeatedly, and most systems rebuild their understanding from scratch every time. That means answering the same two questions on every request: what does this term mean and how does it relate to everything else? Dell is extending the Dell AI Data Platform with three new features to answer both at once.

- The **Unified Semantic Layer** gives structured and unstructured information consistent business meaning with rules, definitions and a searchable glossary, so a term means the same thing everywhere it appears. If one system calls it a client and another calls it an account, AI now knows they're the same thing. Users can also import existing ontologies and classification taxonomies, reusing industry-standard or enterprise-specific taxonomy assets. Dell is also enabling NVIDIA Auto-Ontology, an open-source library that builds knowledge graphs from enterprise data, to extend the Unified Semantic Layer.
- The **Enterprise Knowledge Graph** maps how structured and unstructured data is related, so an application or an agent can find the right context. It uses metadata, lineage and query history to keep tuning the graph as activity changes. When an agent asks a question, the platform pulls in every related piece it's allowed to see: the right tables, data products, multimodal data and vector indexes, wherever they live. For instance, a manufacturer chasing a production line problem can trace one odd sensor reading to the machine, its repair history, the supplier batch and the orders at risk.
- **Knowledge Agents** put that context to work. Each one acts as a trusted advisor on a single topic, grounded in a defined slice of the Enterprise Knowledge Graph. Customers set the rules: what guidance it follows, what data it can see, what quality bar it must clear and how much it's allowed to spend. NVIDIA Nemotron Retriever models provide reasoning and visual understanding to Knowledge Agents.

The Unified Semantic Layer, Knowledge Graph and Knowledge Agents hold a customer's most valuable and sensitive information, so they stay inside the platform in the customer's data center. The platform builds context from that data, keeps it current as the business changes and shares it across every application and agent, without locking customers into any single model, data or storage provider.

### Acceleration and security built into the data path

Business context is only reliable if data is accurate, current and governed from source to application. Customers should not have to adopt another proprietary layer to accelerate different parts of their data environment. Dell brings acceleration into the data engines themselves, working with NVIDIA across the data path. Data must be transformed and analyzed, then reach accelerated compute fast enough to keep GPUs saturated and stay protected along the way. Dell data and storage advancements include:

- **cuDF accelerated data processing with Apache Arrow:** New testing shows that the [Dell Data Processing Engine](#), powered by NVIDIA cuDF on NVIDIA RTX PRO 4500 Blackwell Server Edition GPUs, processes data nearly 4 times faster on average than CPUs alone across a range of workloads, and up to 20 times faster on batch processing workloads.<sup>1</sup> Apache Arrow moves data efficiently between Dell storage and processing, so jobs can query data in place. This can

reduce data-preparation time and shorten the path from raw enterprise information to trusted context.

- **Expanded multitenancy and security:** [Dell PowerScale](#) now supports up to 500 tenants in a single cluster, with mTLS over NFS to encrypt and authenticate file traffic and more granular role-based access control for managing each tenant. AI service providers and enterprises running shared AI platforms can serve more teams and customers from one cluster while keeping each tenant's data isolated and secure.
- **Benchmark workload performance:** The new [Dell Storage Performance Tool](#) helps customers confidently size AI infrastructure, compare vendor solutions and keep GPUs fed by testing the performance of their S3-compatible object storage across training, inference and checkpointing.

### Expertise to get to production

Dell is expanding Dell AI Data Platform implementation services across its data and storage engines to move organizations from deployment to production results quickly and reliably. Dell Services establish a production-ready foundation, activating the platform's analytics, processing, search and orchestration capabilities while keeping the platform tuned to the demands of AI workloads.

### Perspectives

#### Arthur Lewis, president, Infrastructure Solutions Group, Dell Technologies:

"Data without context is just noise. Most enterprises have spent years making their data accessible. That's not the same as making it usable. An agent that can find a customer record but doesn't know what it means, how it connects to everything else, or whether it can be trusted isn't intelligent. It's just fast. The companies that solve it won't just deploy AI faster or run more agents. They'll get more out of the data they already have."

#### Jason Hardy, vice president of storage technology, NVIDIA:

"AI agents are only as effective as the data they can access, understand, and trust, and NVIDIA accelerated computing and AI software help turn governed data into AI-ready context. By bringing NVIDIA cuDF acceleration directly into the Dell Data Processing Engine, Dell helps shorten the path from stored data to GPU-accelerated AI-ready data their agents can use to drive real world innovation."

### Availability

- The Unified Semantic Layer, Enterprise Knowledge Graph and Knowledge Agents will be released in 1H 2027.
- Dell Data Processing Engine enhancements with the NVIDIA acceleration stack will be available in December 2026, with further acceleration using Apache Arrow expected 1H 2027.
- Dell PowerScale security and multitenancy enhancements will be available in November 2026.
- The Dell Storage Performance Tool is available now.
- AI-ready data services are available now.

### Additional resources

- **Blog:** [Getting Enterprise AI into Production Starts with Context](#)
- Connect with Dell on [X](#) and [LinkedIn](#)

1. Based on Dell internal testing in September 2026 comparing GPU-accelerated and CPU-only Apache Spark runs across diverse data processing workloads on a Dell PowerEdge R770 with NVIDIA RTX PRO 4500 Blackwell Server Edition GPUs. Results showed a 3.9x average speedup and a 20.4x peak speedup on a batch processing (data mining) workload. The numbers are from default configurations without any performance tuning to the engine or infrastructure. Actual results may vary.

### About Dell Technologies

Dell Technologies (NYSE: DELL) helps organizations and individuals build their digital future and transform how they work, live and play. The company provides customers with the industry's broadest and most innovative technology and services portfolio for the AI era.

View source version on businesswire.com: <https://www.businesswire.com/news/home/20261006142491/en/>

### Media Contacts

Media Relations, [Media.Relations@Dell.com](mailto:Media.Relations@Dell.com)

Source: Dell Technologies