Dell AI Data Platform with NVIDIA Supercharges Enterprise AI with Breakthrough Data Orchestration and Storage Innovations

March 16, 2026

- Dell AI Data Platform with NVIDIA advancements automate the complete AI data lifecycle and deliver extreme AI storage performance for demanding agentic AI workloads
- Dell Technologies will support all of NVIDIA's latest AI storage and data management innovations

SAN JOSE, Calif., March 16, 2026 /PRNewswire/ -- Dell Technologies (NYSE: DELL) announces Dell AI Data Platform with NVIDIA advancements that help enterprises discover and activate enterprise data while delivering extreme storage performance to power AI applications and autonomous AI agents.

Why it matters
AI is rapidly shifting from assistive tools to autonomous, agentic systems, but its effectiveness is constrained by the data it can access, trust and act upon. Many enterprises hit a wall because much of their data remains trapped in silos, lacking structure, business context, and governance. The result: AI initiatives stall, investments underdeliver and competitive advantages slip away.

Dell and NVIDIA are removing one of the biggest blockers to enterprise AI: data that's too slow, too siloed, or too messy to use. As a core component of the Dell AI Factory with NVIDIA, **the Dell AI Data Platform with NVIDIA** activates enterprise data for AI while maintaining security, governance, and best-in-class performance at scale. Customers see up to **12X faster vector indexing[1], 3X faster data processing,[2] and 19X faster time-to-first-token[3]** than traditional computing approaches.

Automating the entire AI data lifecycle
Dell data engines, accelerated by NVIDIA AI infrastructure, automate the complete AI data lifecycle and dramatically reduce data preparation time while maintaining enterprise governance.

- The **Dell Data Orchestration Engine**, powered by technology from Dell's recent Dataloop acquisition, redefines how enterprises operationalize data for AI. The no-code, low-code engine orchestrates the AI data lifecycle—automatically discovering, labeling, enriching, and transforming structured, unstructured, and multimodal data into governed, AI-ready datasets at scale. By combining automated pipelines with active learning and human-in-the-loop workflows, organizations can continuously improve dataset quality and model accuracy while maintaining governance and control. The **Data Orchestration Engine Marketplace** lets organizations deploy production-ready data workflows without having to build them from scratch with a curated library of **NVIDIA NIM microservices**, **NVIDIA AI Blueprints** and more than 200 other models, applications and templates.

- Dell Technologies supports the latest **NVIDIA AI-Q** blueprint, helping enterprises build customizable AI agents that deliver actionable insights for smarter decision-making. NVIDIA-accelerated data engine integrations in the Dell AI Data Platform enable high-performance data preparation, retrieval, and reasoning pipelines across structured and unstructured data. Customers also gain access to a growing library of **pre-built NVIDIA blueprints** and **NIM microservices**, along with the **NVIDIA Nemotron 3 Super** model on **Dell Enterprise Hub on Hugging Face.**

- Dell Technologies will also support **NVIDIA STX**, a new modular reference design powered by next-generation **NVIDIA Vera Rubin NVL72**, **NVIDIA BlueField-4 DPUs**, and **NVIDIA Spectrum-X™ Ethernet** networking that accelerates how organizations manage, process, and retrieve data for AI.

- The new **AI Assistant** within the **Dell Data Analytics Engine** brings conversational natural language interface directly into SQL analytics. Business users can query, visualize and collaborate on governed data products with a common semantic understanding of key metrics intuitively without specialized SQL knowledge. This democratizes data access, streamlines decision-making and unlocks deeper insights faster, which is particularly critical for organizations deploying AI agents that need to access structured data.

- Within the Dell AI Data Platform with NVIDIA, the introduction of **NVIDIA RTX PRO™ Blackwell Server Edition GPUs** will bring acceleration directly into the data platform layer. Accelerated **NVIDIA CUDA-X libraries** including **NVIDIA cuDF** for structured data processing, and **NVIDIA cuVS** for vector indexing and search applied to unstructured data, work alongside Dell's data engines and optimized infrastructure to deliver up to **3x faster SQL queries[4]** and **12x faster vector indexing**.[5] These technologies help organizations develop more responsive AI applications and improved infrastructure efficiency

when processing and preparing data at scale.

Extreme-scale storage software innovations keep GPUs running at full speed
As enterprises move from AI experimentation to production deployment, storage becomes the critical constraint. Traditional storage architecture slows down as it scales, creating bottlenecks that leave GPUs idle and waste infrastructure investments. Dell's AI-optimized storage engines solve this problem with purpose-built architectures that maintain performance at massive scale.

- **Dell Lightning File System**, the world's fastest parallel file system[6], delivers extreme performance density for AI training and inferencing environments with up to 150 GB/second per rack[7], up to 20X greater performance versus traditional flash-only scale out file competitors[8] and up to 2X greater throughput per rack unit than competing parallel file systems.[9] Purpose-built fabric architecture with direct storage access prevents slowdowns, keeping GPUs fully utilized at massive scale. Lightning FS integrates seamlessly into NVIDIA-based AI infrastructures, keeping training and inference workloads running at full speed.

- **Dell Exascale Storage**, the only 3-in-1 storage built for extreme-scale AI and HPC[10], gives IT teams the flexibility to deploy Dell's best-of-breed file, object, and parallel file system storage software on the latest Dell PowerEdge servers. Customers can allocate Dell PowerScale, Dell ObjectScale, and/or Dell Lightning File System storage resources on a common hardware platform to support the most demanding AI and HPC environments like high-frequency trading and neoclouds. With support for NVIDIA CX-8 and CX-9 SuperNICs and planned network connectivity up to 800GbE, **Exascale delivers read performance up to 6TB/second per rack[11]**, providing the high throughput required by multimodal AI workloads.

- **NVIDIA CMX context memory storage platform support** and **inference acceleration with KV Cache** on shared storage across Dell PowerScale, Dell ObjectScale and Dell Lightning File System allows organizations to offload KV cache from GPU memory to Dell CMX Storage and high-speed shared network storage based on performance needs. This dramatically improves GPU utilization for long-context and agentic AI workloads, allowing AI systems to maintain context across extended interactions without exhausting GPU memory. This capability is essential for enterprises deploying AI agents that need to reference extensive historical data or maintain long conversation threads.

- **PowerScale performance testing:** New testing demonstrates that **Dell PowerScale's** software-driven Parallel Network File System (pNFS) architecture delivers up to **6X faster performance with large files** in enterprise AI environments compared to NFSv3.[12] This keeps GPU-intensive AI workloads continuously fed with data, reducing bottlenecks across the entire pipeline and ensuring expensive GPU resources don't sit idle waiting for data.

Dell AI Factory with NVIDIA delivers proven path to enterprise AI ROI
Dell Technologies today marks the two-year anniversary of the Dell AI Factory with NVIDIA with advancements spanning its end-to-end AI infrastructure, software, solutions, and services portfolio that help enterprises move AI from pilot to production at scale. With over 4,000 customers deploying the Dell AI Factory, and early adopters seeing up to 2.6x ROI within the first year[13], Dell proves that an end-to-end approach delivers measurable business results.

Perspectives

Travis Vigil, senior vice president, ISG Product Management, Dell Technologies:
"The number one problem enterprises face when moving AI pilots to production is curating the data they already have and putting it to work. The Dell AI Data Platform with NVIDIA automates the entire data lifecycle and delivers the speed and scale AI workloads demand. We've done the integration work, so customers deploy faster, scale with confidence and see real returns. Together with NVIDIA, we're defining what enterprise AI infrastructure needs to be."

Jason Hardy, vice president, Storage Technologies, NVIDIA:
"The shift to autonomous agents requires a fundamentally different approach to data infrastructure, with automated orchestration, AI-native storage and GPU-optimized performance architected to work together. Dell's enterprise expertise, combined with full-stack NVIDIA AI infrastructure, creates the foundation organizations need to deploy AI at scale."

Availability

- Dell Data Orchestration Engine and Marketplace are available in Q1 CY26.
- Dell and NVIDIA Blueprints are available now.
- Dell support for NVIDIA AI-Q blueprint is available now.
- AI Assistant for the Dell Analytics Engine will be available in 1H CY26.
- NVIDIA GPU-accelerated data processing and data indexing in the Dell AI Data Platform will be available in 2H CY26.
- Dell Lightning File System will be available in April 2026.
- Dell Exascale Storage is targeted for availability in early 2H CY26.
- Dell support for NVIDIA's latest innovations will roll out throughout the year.

Additional resources

- Blog: Introducing the Supercharged AI Data Platform with NVIDIA
- Blog: Scale Your AI Ambitions with Dell Storage and NVIDIA
- Connect with Dell on X and LinkedIn

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.

1. Dell results are based on internal testing compared to Elasticsearch's results, Dec. 2025.
2. Disclaimer: Based on Dell internal analysis, Sept. 2025.
3. Dell results are based on internal testing using the Qwen3-32B model, Oct. 2025.
3. Disclaimer: Based on Dell internal analysis, Sept. 2025
4. Disclaimer: Based on Dell internal analysis, Sept. 2025
5. Dell results are based on internal testing compared to Elasticsearch's results, Dec. 2025.
6. Based on Dell preliminary testing comparing random and sequential throughput per rack unit, May 2025. Actual performance may vary.
7. Based on internal analysis of sequential and random read I/O. Actual results may vary. Feb. 2026
8. Based on Dell internal testing comparing IOPs performance per node, Mar. 2026. IOPs rates based on FIO over a remote file system. Actual performance may vary.
9. Based on Dell preliminary testing comparing random and sequential throughput per rack unit, May 2025. Actual performance may vary.
10. Based on publicly available documentation from leading enterprise storage vendors as of March 2026. Comparison refers to distinct file, object and parallel-file engines on one reusable hardware platform, excluding single-engine multi-protocol designs.
11. Based on internal analysis of sequential and random read I/O for Lightning File System, Feb. 2026. Actual results may vary.
12. Based on preliminary internal testing on single-client random-I/O performance of large files. Results will vary by workload and configuration. March 2026
13. Based on Enterprise Strategy Group paper commissioned by Dell, "Analyzing the Economic Benefits of the Dell AI Factory with NVIDIA," comparing the ROI of on-premises Dell and NVIDIA solution, August 2025. Estimated costs were modeled utilizing Llama 3 70B LLM for inferencing and model fine-tuning workloads by organizations over a 4-year period. Server models used were XE9680s with 8 x H100 GPUs. Actual results may vary.

View original content to download multimedia:https://www.prnewswire.com/news-releases/dell-ai-data-platform-with-nvidia-supercharges-enterprise-ai-with-breakthrough-data-orchestration-and-storage-innovations-302715096.html

SOURCE Dell Technologies

Media Relations: Media.Relations@Dell.com