



Dell AI Factory with NVIDIA Delivers Proven Path to Enterprise AI ROI

March 16, 2026

Dell AI Data Platform, end-to-end infrastructure and services advancements remove barriers to enterprise AI deployment

- 4,000+ Dell AI Factory customers, up to 2.6x ROI, clear path from experiment to production
- Broad portfolio updates streamline enterprise AI deployments, turn siloed data into AI fuel and compress pilot-to-production timeline

SAN JOSE, Calif., March 16, 2026 /PRNewswire/ -- Today, Dell Technologies (NYSE: DELL) marks the two-year anniversary of the Dell AI Factory with NVIDIA by announcing advancements across its AI data platform, end-to-end AI infrastructure, and AI 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,¹ Dell proves that an end-to-end approach delivers measurable business results.



Why This Matters

The enterprise AI landscape is undergoing a fundamental shift. As AI code assistants and agentic workflows drastically lower the cost and time to build custom applications, CIOs are increasingly choosing to develop AI capabilities in-house, on-premises—driving the need for owned infrastructure.

Yet unclear ROI remains the top obstacle preventing AI deployments at scale. Two years of the Dell AI Factory with NVIDIA has revealed three critical requirements for achieving measurable returns: data platforms that make enterprise information AI-ready, infrastructure that seamlessly scales the latest innovations efficiently from pilot to production, and solutions and services that compress time to value by simplifying deployments and accelerating ROI. Dell is the premier provider delivering all three with NVIDIA technology at the core, creating a proven path from AI investment to business outcome.

Three Capabilities That Define Enterprise AI Leadership

As the top AI infrastructure provider², Dell's AI infrastructure portfolio—the industry's broadest³—delivers integrated capabilities across data, infrastructure, solutions and services.

Data platforms that turn institutional knowledge into AI fuel

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. The [Dell AI Data Platform with NVIDIA](#) addresses this challenge with a unified platform for AI that combines Dell's high-performance storage, modular data engines, and NVIDIA accelerated computing, networking, software and CUDA-X libraries. As the data foundation of the Dell AI Factory with NVIDIA, it handles workloads from retrieval-augmented generation (RAG) and multimodal search to agentic workflows and large-scale data processing. Advancements [announced](#) today make it faster and easier for companies to turn data into real AI results.

Infrastructure that enables AI workflows from desktop to data center

Dell's next-generation infrastructure supports AI workflows at every stage, from rapid prototyping to production deployment at scale.

For desktop AI development and autonomous agents:

- [Dell Pro Max with GB10 and Dell Pro Max with GB300](#) – Purpose-built desktop AI supercomputers that build and run AI, delivering the compute power, memory capacity and always-on reliability that agentic AI workflows and autonomous agent development require. With NVIDIA NemoClaw and NVIDIA OpenShell, enterprises can build safe, autonomous, long-running agents at the desk with local AI that keeps data secure and private. Dell is the first OEM to ship a desktop with [NVIDIA GB300 Grace Blackwell Ultra Desktop Superchip](#), delivering up to 20 petaFLOPS of FP4 performance and 748GB of coherent memory for developing and deploying autonomous AI agents at trillion-parameter scale.
- [Dell Pro Precision workstations](#) deliver the power and expandability AI developers and data scientists need. Tower configurations feature up to five [NVIDIA RTX PRO Blackwell Desktop Generation GPUs](#), while mobile workstations with [NVIDIA RTX PRO Blackwell Generation Laptop GPUs](#) deliver more performance in thinner, sleeker designs.

For production AI at scale:

- **PowerEdge XE9812** is Dell's flagship liquid-cooled server leverages the [NVIDIA Vera Rubin NVL72](#) platform for massive real-time training and inference.
- **PowerEdge XE9880L, XE9882L, and XE9885L** are liquid-cooled servers featuring [NVIDIA HGX™ Rubin NVL8](#) designed to accelerate validated AI performance within existing data center footprints and power constraints.

For enterprise workloads in the data center:

- **PowerEdge R770, R7715 and R7725** configured with the new [NVIDIA RTX PRO™ 4500 Blackwell Server Edition GPU](#) allows teams to add AI acceleration to general-purpose infrastructure.
- **PowerEdge R9822 and M9822** debut [NVIDIA Vera CPU](#), giving organizations more choice for demanding enterprise workloads.

For high-performance networking and emerging technologies:

- **Dell PowerSwitch SN6000-series** are [NVIDIA Spectrum-6](#) Ethernet switches with 1.6Tbs, liquid cooling and co-packaged optics options for Vera Rubin-based Dell platforms.
- **PowerSwitch SN5610 and SN2201** now offer expanded network OS choices including Cumulus Linux and Enterprise SONiC Distribution by Dell Technologies.
- **NVIDIA Quantum-X800 InfiniBand Q3300-LD** liquid-cooled switches deliver high-bandwidth networking for AI and cloud-native workloads.
- **Dell Integrated Rack Scalable Systems (IRSS)** expands to include Dell PowerSwitch and NVIDIA liquid-cooled switching, providing unified, rack-level power and cooling management for AI infrastructure.

[NVIDIA NVQLink](#) and [NVIDIA CUDA-Q support](#) – Dell is the first OEM to integrate NVIDIA NVQLink with CUDA-Q across PowerEdge servers featuring NVIDIA AI infrastructure, allowing enterprises and research institutions to explore emerging quantum-classical computing use cases. These capabilities accelerate discoveries in advanced drug development and materials science simulations by combining the processing power of Quantum Processing Units with NVIDIA accelerated computing for quantum systems control and error correction on a trusted foundation of Dell PowerEdge servers.

Solutions and services that accelerate deployment and prove ROI

Updated Dell AI Solutions combine new modular architecture with Dell Automation Platform blueprints and NVIDIA AI Enterprise software to deliver enterprise outcomes while simplifying operations and reducing deployment complexity. New services bridge skill gaps and scale deployments from experimentation to production.

Accelerating enterprise AI workloads:

- **Knowledge assistant** provides the foundation for designing, deploying and managing intelligent assistants, working with industry leaders like Aible, Cohere's North and NVIDIA.
- **ClearML blueprint** improves agentic AI environments for enterprises with secure, efficient GPU cluster management and workload scheduling.
- **Agentic AI platform**, in collaboration with Cohere's North, DataRobot and NVIDIA allows enterprises to securely deploy and manage AI agents with orchestration, governance and observability.
- **Dell Accelerator Services for Agentic AI** provide packaged capabilities to support businesses at any stage, from experimentation and validation to enterprise-wide integration, closing skill gaps and reducing technical complexity.

Simplifying AI infrastructure deployment:

- **Dell AI Factory with NVIDIA modular architecture** offers a clear, simplified path to enterprise AI by addressing deployment complexity, managing rapid technology change and supporting continuous adoption. Integrated automation gives organizations the flexibility to start at the right size and scale as needs evolve.

Perspectives:

Michael Dell, chairman and chief executive officer, Dell Technologies:

"Two years ago, enterprises were asking how to access AI technology. Today, they're asking how to make their data AI-ready, how to operationalize AI at scale and how to prove ROI. The Dell AI Factory with NVIDIA answers all three questions. We're brought in from the start as a trusted advisor, helping customers navigate their entire AI journey—from turning raw data into AI fuel, through deployment and to measurable business outcomes."

Jensen Huang, founder and chief executive officer, NVIDIA:

"AI infrastructure is being built everywhere — every company will be powered by it, every country will build it— and it demands integrated data platforms, scalable infrastructure and deployment expertise. Dell Technologies delivers all three, with NVIDIA at the core. The Dell AI Factory with NVIDIA is a proven infrastructure blueprint for every phase of AI powering the next industrial era."

Availability

- Dell Pro Precision 5 and 7 Series mobile workstations with NVIDIA RTX PRO Blackwell GPUs will be available in May.
- Dell Pro Precision 9 T2/T4/T6 will be available in May.
- Dell has shipped Dell Pro Max with GB300 to select customers in March 2026, with plans to ship more broadly in the coming months.
- Dell PowerEdge XE9812 will be globally available 2H 2026.
- Dell PowerEdge XE9880L, XE9885L will be globally available Q3 2026.
- Dell PowerEdge R770, R7715 and R7725 with NVIDIA RTX PRO 4500 Blackwell Server Edition GPUs are globally available now.
- Dell PowerEdge M9822 AND R9822 will be globally available in September.

- Dell PowerSwitch SN6000-Series will be globally available starting in July.
- Dell SONiC with Spectrum-based PowerSwitch SN5610 and S2201 will be globally available in March.
- NVIDIA Quantum-X800 Q3300-LD will be globally available by Dell Technologies in Q4 2026.
- Dell PowerEdge NVIDIA NVQLink and CUDA-Q integration is available now.
- Knowledge assistant is globally available now.
- Agentic AI platform with Cohere's North and DataRobot are available now, agentic AI platform with ClearML will be available in March.
- Dell AI Factory with NVIDIA modular architecture will be globally available in April.
- Dell Accelerator Services for Agentic AI are available now.

Additional resources

- [Check out](#) Dell ISG President, Arthur Lewis' take on how Dell is ushering in the era of enterprise AI.
- [Find out more](#) about the Dell AI Data Platform with NVIDIA.
- [Learn more](#) about how AI at scale stares with your data.
- [See here](#) for more information about bringing the AI lab to your desktop.
- [Get more details](#) about Dell compute and networking innovations.
- [Learn how](#) Dell AI Solutions and Services will accelerate your AI journey.
- 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 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.

² IDC Worldwide Quarterly Trackers for Servers, External Storage Systems, PBBA, and AI Infrastructure Q32025: showing Dell #1 in AI infrastructure, servers, external storage systems, and purpose built backup appliances worldwide

³ Based on Prowess Consulting's paper commissioned by Dell, "Which Vendor Offers the Broadest AI Portfolio for Scalable Innovation?" October 2025

 View original content to download multimedia: <https://www.prnewswire.com/news-releases/dell-ai-factory-with-nvidia-delivers-proven-path-to-enterprise-ai-roi-302715157.html>

SOURCE Dell Technologies

Media Relations: Media.Relations@Dell.com