Dell Technologies Closes the Gap Between AI Ambition and AI Outcomes

May 18, 2026

Dell AI Factory with NVIDIA advancements span agentic AI, data orchestration, next-generation infrastructure and an expanding open ecosystem built to advance enterprise AI from experiment to reality

- **Agentic AI at scale:** Deskside solution with NVIDIA NemoClaw support lets enterprises build and run secure, autonomous agents locally with data that never leaves their environment, plus OpenShell integration across the Dell AI Factory provides seamless scaling from deskside to data center
- **AI-ready data:** Faster indexing of billions of files of all types, up to 6x faster SQL query performance and new NVIDIA Omniverse integration turn enterprise data into AI fuel
- **Next-generation infrastructure:** Turnkey rack deployment, intent-based networking and the industry's most efficient rack-mount CDU advance the AI foundation
- **An open ecosystem that delivers outcomes:** New solutions with Google, Hugging Face, OpenAI, Palantir, Reflection, ServiceNow, SpaceXAI and others, plus security solutions and services, and a new Dell AI Ecosystem Program give enterprises more ways to deploy AI on infrastructure they control

LAS VEGAS--(BUSINESS WIRE)--May 18, 2026-- Dell Technologies World – Dell Technologies (NYSE: DELL) announces a broad set of advancements to the Dell AI Factory with NVIDIA, delivering the foundation enterprises need to move from AI ambition to realized outcomes. With more than 5,000[1] customers already deploying the Dell AI Factory, these portfolio additions are designed to help organizations adopt AI with confidence, scale with purpose, and achieve results on infrastructure they control, with data they trust.

This press release features multimedia. View the full release here: https://www.businesswire.com/news/home/20260518066830/en/

Dell PowerRack in data center

Why it matters

Most enterprises don't have an AI ambition problem. They have an AI execution problem. Data availability and quality remain the top implementation challenges across organizations at every stage of AI maturity.[2] Without a trusted, AI-ready data foundation, even the best infrastructure falls short. Pilots stall before they reach production, and the promise of agentic AI remains out of reach. Dell and NVIDIA address this with a simplified, integrated approach that can accelerate time-to-value by up to 84%[3] and gives enterprises the confidence to scale.

Agentic AI at for every workload

As agentic AI workloads grow in complexity, cloud costs are becoming increasingly unpredictable. Organizations are seeking a more controlled approach to deploying autonomous AI where performance, data sovereignty and cost efficiency are paramount.

Dell Deskside Agentic AI, a new solution powered by Dell's high-performance workstations and NVIDIA NemoClaw, allows enterprises to more securely build and run autonomous agents locally with data that never leaves the device. Supported by end-to-end Dell services, the solution is designed for specialized groups in software engineering, academic research and regulated industries, converting variable cloud token costs into a controlled infrastructure investment. With Dell Deskside Agentic AI, organizations can break even versus public cloud API costs in as little as three months.[4]

NVIDIA OpenShell, the secure runtime for autonomous agents, is now supported across the entire Dell AI Factory with NVIDIA. This allows organizations to build, deploy and govern agents with privacy controls, from Dell Pro Precision towers and Dell Pro Max with GB10 and GB300 through to Dell PowerEdge XE servers. The Dell-NVIDIA AI-Q 2.0 Reference Architecture, powered by the Dell AI Data Platform with NVIDIA, extends this foundation with a production-ready multi-agent research workflow for regulated industries. Read more here.

Turning enterprise data into AI fuel

AI is only as good as the data it can find, trust and act on. Dell is announcing significant advancements to the **Dell AI Data Platform** that make enterprise data AI-ready at scale across the full lifecycle, from discovery and preparation to analytics and AI-driven experiences.

- **Unify and orchestrate AI data pipelines at scale:** Enhancements to the Dell AI Data Platform's orchestration and search capabilities index billions of unstructured files and connect them into governed pipelines, accelerating data discovery and dataset creation for AI. Integrated services for Dell AI Data Platform help customers tackle challenges like data preparation, skills gaps and operational complexity so they can move from pilots to production faster.
- **Accelerate SQL analytics for NVIDIA Blackwell and future NVIDIA Vera CPU platforms:** Within Dell AI Data Platform, the Dell Data Analytics Engine, powered by Starburst, brings GPU-accelerated SQL analytics to enterprise AI, delivering up to 6x faster query performance on NVIDIA Blackwell GPUs today[5] with support designed for future platforms including Vera. This accelerates insights for both traditional data analytics and data-intensive agentic AI applications.
- **Higher density, lower TCO:** The new Dell ObjectScale X7700 ultra-dense appliance delivers up to 45%[6] more HDD capacity than the previous generation, with flexible compute-to-storage scaling and improved TCO. Forthcoming 245 TB

all-flash drive support will more than triple[7] ObjectScale flash density.
- **Power digital twins and AI-driven experiences with unified data:** Within Dell AI Data Platform with NVIDIA, Dell storage and search engines integrate with NVIDIA Omniverse libraries to combine scalable object storage with semantic, vector-based asset search. This helps connect PLM systems and repositories directly into Omniverse, feeding digital twins and physical AI training and validation workflows with trusted, well-organized data.

Next-generation infrastructure built for the demands of modern AI

Dell is expanding its AI infrastructure portfolio with new systems built for modern enterprise AI workloads. As the top rack-scale infrastructure provider,[8] shipping more than twice the number of rack-scale servers compared to the closest competitor,[9] Dell is adding PowerRack to the industry's broadest AI infrastructure portfolio.[10]

Dell PowerRack is a fully integrated system – compute, networking and storage engineered as one – with thermal design, power management and software optimization built to work together from the ground up. The result is accelerated AI and HPC workloads at enterprise scale, without the integration overhead of component assembly. Dell PowerRack for storage and networking are simplified, rack-scale platforms delivering factory-integrated dedicated Dell Exascale storage and Dell PowerSwitch networking with a system-level approach to performance, power and cooling, managed consistently through the Dell Integrated Rack Controller.

Additional infrastructure updates include:

- **The industry's only 4-in-1 storage built for extreme-scale[11]:** Dell is adding PowerFlex to Dell Exascale Storage, completing a unified rack architecture for Dell PowerRack that supports block (PowerFlex), file (PowerScale, Lightning File System), and object (ObjectScale) for AI, HPC and demanding enterprise workloads.
- **Compact, mountable rack workstation:** The Dell Pro Precision 7 R1 brings high-performance computing to space-constrained environments in a 1U form factor with NVIDIA RTX PRO Blackwell Max-Q Workstation Edition GPUs and up to 64TB of storage.
- **Unified rack management**: New releases of the Dell Integrated Rack Controller and Dell OpenManage Enterprise deliver a unified control plane for integrated compute, with expanded remote device connectivity and orchestration across the entire rack.
- **Next-generation cooling:** The Dell PowerCool CDU C7000 is the first rack-mount cooling distribution unit to meet the cooling needs for the NVIDIA Vera Rubin NVL72 platform in a compact 4U, 19" form factor, and extends Dell's cooling capacity and support for up to 40°C facility water.

Scalable solutions with an expanding open ecosystem

The new **Dell AI Ecosystem Program** gives AI software providers a structured path to validate solutions on Dell AI Factory infrastructure, turning fragmented innovation into proven, deployable outcomes. For enterprises, this means lower-risk paths to production-scale AI, faster POC-to-production and the ability to run AI solutions where data lives.

Bringing AI leaders and frontier models to the enterprise helps organizations maintain control over their data, models and operations within their trusted environments.

- **Google** and Dell are collaborating to bring Gemini 3 Flash models on Google Distributed Cloud on Dell PowerEdge XE9780 servers. This fully integrated, on-premises solution allows enterprises to run advanced generative AI workloads within a private, confidential computing environment. By leveraging a secure BIOS and robust security attestation, organizations can more seamlessly meet strict data protection, residency, and sovereignty requirements. The collaboration supports the latest Gemini models—featuring expanded 1M+ context windows and advanced AI tools like Gemini CLI—delivering the security and control modern enterprises demand.
- **Dell Enterprise Hub on Hugging Face** gives enterprises on-premises access to a curated collection of the latest open-weight models, including MiniMax-M2.7, DeepSeek Pro, DeepSeek-V4, GLM 5.1 and Kimi K2.6, optimized for Dell AI Factory infrastructure. As the industry moves toward highly efficient architectures delivering frontier-level reasoning at long context lengths, this collaboration shifts the tokenomics of enterprise AI, giving organizations a trusted, more secure path to deploy the most capable open models where their data lives, at a fraction of the cost.
- **OpenAI** and Dell Technologies are collaborating to help more enterprises deploy Codex in the environments where their most important data, systems, and workflows already live. Through this collaboration, Codex will connect with the Dell AI Data Platform, which many businesses already use to store, organize, and govern enterprise data on-premises. The collaboration will help customers bring Codex closer to the internal context that makes agents useful: codebases, documentation, business systems, operational knowledge, and team workflows. Dell and OpenAI will also explore how Codex can connect with the Dell AI Factory, which businesses use to power their AI workloads.
- **Palantir**'s Foundry and AIP platform is coming on-premises to the Dell AI Factory, where Palantir's Ontology layer will be deployed on Dell ObjectScale and PowerFlex to ingest data from enterprise sources and automate business workflows using AI models deployed on the Dell AI Factory. This will allow enterprises and sovereign entities to connect all their data sources across their enterprise, define and dynamically manage relationships between those data sources and optimize their business operations with the full weight of AI, all within their organization boundaries.
- **Reflection's** open-source frontier AI models are coming on-premises on the Dell AI Factory. Open models help enterprises

in regulated industries including governments and sovereign entities to deploy AI in fully controlled environments. Reflection's frontier-level quality models deployed on the Dell AI Factory, integrated with the Dell AI Data Platform, will help customers securely extract knowledge from on-premises data sources.
- **SpaceXAI** and Dell deliver Grok's advanced reasoning and multimodal capabilities as more secure, enterprise-grade AI assistants, deployable fully on-premises or in a hybrid approach.

In addition, **ServiceNow** customers will be able to leverage the Dell AI Factory to bring together infrastructure and enterprise workflow automation, enabling organizations to discover, govern, and operationalize AI focused on business outcomes

New validated AI solutions for common enterprise outcomes spanning **agentic AI** with Mistral, **computer vision** with Fogsphere and Ipsotek, an Eviden business, **immersive AI** with UneeQ Digital Humans, and **code assistants** with Poolside — are deployable directly from the Dell Automation Platform catalog. New **security** solutions and services using CrowdStrike, Fortanix and F5 provide full-stack, 24/7 protection and confidential AI across AI infrastructure, data, models and applications for more resilient AI foundations. JFrog and Dell deliver a central hub for securely managing AI models, MCPs, Agent Skills and software artifacts at scale.

Perspectives:

Michael Dell, chairman and chief executive, Dell Technologies:

“With the advent of Agentic AI, every organization now faces the same challenge to turn intelligence into impact at speed or become obsolete. At Dell Technologies, we’re helping customers turn their data into AI fuel on infrastructure they control with security, governance and cost efficiency.”

Jensen Huang, founder and CEO, NVIDIA:

"Agentic AI has arrived - enterprise AI adoption is going parabolic. Dell and NVIDIA are building the full-stack AI factory for this moment, with accelerated computing, networking, storage, software and services that scale from the desktop to the data center — turning AI’s potential into unprecedented productivity for enterprises everywhere.”

Availability

- Dell Deskside Agentic AI Solutions are now available.
- Dell AI Data Platform orchestration and search advancements will be available in Q2 2026.
- Dell Data Analytics Engine accelerated by NVIDIA Blackwell and NVIDIA Vera will be available in Q1 2027.
- Dell ObjectScale NVIDIA Certified Storage validation will be available in Q2 2026.
- Dell ObjectScale and Dell Data Search Engine integration with NVIDIA Omniverse is available now.
- Dell Exascale with the Dell PowerFlex addition will be available in 1H 2027.
- Dell Pro Precision 7 R1 will be available in July 2026.
- Dell PowerRack for compute is available now.
- Dell PowerRack for Dell PowerSwitch networking will be available in September 2026
- Dell PowerRack for Dell Exascale storage will be available in 2H 2026.
- Dell Integrated Rack Controller and Dell OpenManage Enterprise releases will be available in May 2026.
- Dell PowerCool CDU C7000 will be available in Q3 2026.
- Dell AI solutions will be available throughout 2026.

Additional resources

- Dell Technologies World 2026 Press Kit
- Learn more about Dell AI Solutions.
- Blog: Securing Your AI Workloads with Dell AI Factory with NVIDIA
- Blog: The Agentic AI Continuum: When Deskside Makes Sense
- Blog: Density by Design: Inside Dell Pro Precision 7 R1
- Blog: Dell AI Data Platform with NVIDIA: Full Throttle AI
- Blog: Securing Storage: Dell’s Layered Approach to PQC Readiness - is the PQC Blog2 one
- Blog: Dell PowerRack Transforms AI Infrastructure with Scalable Compute, Networking and Storage
- Blog: Simplifying Enterprise AI: Introducing the Dell AI Ecosystem Program
- Blog: Choice Without Compromise: Inside Dell's Expanding AI Ecosystem
- Blog: Dell Enterprise Hub Brings Frontier Open Models to PowerEdge XE9780
- Blog: Dell and Palantir Introduce an On-Premises AI Operating System
- Blog: Dell Technologies and SpaceXAI Collaborate to Bring Grok On-Premises to Enterprise Environments
- 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 April 2026 Dell analysis of customer order data.

[2] Gartner, “AI Maturity Matters: Increased Trust, Improved Effectiveness, Optimized Operations,” June 30, 2025.

[3] Based on a Principled Technologies report commissioned by Dell, Accelerate AI time to value with Dell Services, April 2026.

[4] Based on validated analysis by Signal 65 and Futurum Group: "The Economics of Agentic AI: On-premises Deployments with Dell AI Factory vs. Cloud", May 2026. Based on publicly available API pricing and Dell solution pricing and performance data provided by Dell. Savings assume a multi-year deployment and a range of Dell Pro Max Workstations and PowerEdge Servers being used for general knowledge, sales, and software development workloads all supported by agentic AI over a 5-day work week. Analysis is inclusive of estimated cloud discounts, and infrastructure hosting, energy, infrastructure management, and Dell support services costs. Individual results may vary.

[5] Based on Dell internal analysis, May 2026.

[6] Based on Dell comparison of maximum available storage capacity on planned ObjectScale X7700 versus previous-generation ECS 5000, May 2026. Actual results may vary.

[7] Support for 245 TB drives and an expanded drive count per node are targeted for an ObjectScale release in 2H 2026. Actual usable capacity and density may vary.

[8] Based on IDC data on CY25 rack-scale units sold worldwide.

[9] Based on IDC data on CY25 rack-scale units sold worldwide.

[10] Based upon internal analysis, March 2026

[11] Based on internal analysis, May 2026. Comparison refers to distinct file, object, block, and parallel-file engines on one reusable hardware platform, excluding single-engine multi-protocol designs. Block storage is expected to be available in 1H CY2027.

View source version on businesswire.com: https://www.businesswire.com/news/home/20260518066830/en/

Media Relations: Media.Relations@Dell.com

Source: Dell Technologies