

Dell Technologies Advances Enterprise AI Innovation With NVIDIA

November 17, 2025

Dell AI Factory with NVIDIA advancements accelerate deployment of AI applications, from traditional to agentic, creating integrated, efficient IT environments

- Dell ObjectScale and PowerScale, the Dell AI Data Platform's storage engines, accelerate workflows for AI applications and high performance compute with NVIDIA Dynamo integration
- Dell Automation Platform expansion offers automated deployment of full-stack AI workloads, accelerating time to value
- Server and networking advancements help future-proof data centers for next-generation workloads
- Dell Professional Services provide expert guidance and tailored solutions to drive seamless operations

ST. LOUIS--(BUSINESS WIRE)--Nov. 17, 2025-- **SC25** – Dell Technologies (NYSE: DELL), the world's top provider of AI infrastructure,¹ and NVIDIA are delivering the future of enterprise AI with advancements to the Dell AI Factory with NVIDIA. These advancements provide simplicity, performance and flexibility for advanced environments and support enterprises across AI applications, from traditional to agentic.

Why it matters

A new era of high performance computing (HPC) and enterprise AI is here. As enterprises scale, many face challenges such as managing complex ecosystems of hardware and software or gaining control over their data. According to recent research, 95% of enterprises believe working with a trusted technology partner like Dell Technologies reduces the risks associated with adopting new technologies.² Additionally, 90% agree that bringing AI to their data creates greater value through enhanced control, fresh insights and secure access.³

Enter the Dell AI Factory with NVIDIA—a game-changer for enterprises looking to accelerate outcomes, reduce complexity and maximize ROI. By integrating Dell's robust end-to-end infrastructure with NVIDIA AI technology, backed by expert guidance from Dell Professional Services, organizations can transform ideas into tangible results and stay ahead of evolving technologies and scaling needs.

Accelerate deployment with integrated, automated platforms

Dell's storage and AI solutions offerings help enterprises automate deployments, optimize performance and deliver real-time AI applications with greater efficiency and reliability.

- **Dell ObjectScale and PowerScale**, the Dell AI Data Platform's storage engines for unstructured data, are now integrated with the NVIDIA NIXL library, part of [NVIDIA Dynamo](#). This integration enables scalable KV Cache storage, reuse and sharing, achieving a 1-second Time to First Token (TTFT) at a full context window of 131K tokens – 19X faster than standard vLLM – while reducing infrastructure costs and overcoming GPU memory capacity bottlenecks.⁴
- The Dell AI Factory with NVIDIA now includes **solutions with Dell PowerEdge XE7740/XE7745 servers featuring NVIDIA RTX PRO 6000 Blackwell Server Edition GPUs** and [NVIDIA Hopper GPUs](#). These proven and validated offers feature next-level AI acceleration and computing power to execute advanced use cases—from large-scale multimodal models to emerging agentic AI applications and from enterprise-grade inferencing to training workloads.
- The **Dell Automation Platform**, now expanded to the Dell AI Factory with NVIDIA, will deliver smarter, more automated experiences by deploying validated, optimized solutions with a secure framework. This approach will produce repeatable outcomes, eliminates guesswork, and helps unlock the full potential of AI-driven use cases powered by NVIDIA accelerated computing.
- Software-driven tools like the AI code assistant with Tabnine and agentic AI platform with Cohere North are now automated, getting AI workloads into production faster, streamlining operations and enhancing scalability.
- Beyond the traditional data center, **Dell's ecosystem enablers for AI PCs** offer organizations expanded silicon options, now supporting [NVIDIA RTX Blackwell GPUs](#) and [NVIDIA RTX Ada GPUs](#), ensuring compatibility across a broader range of Dell devices.
- **Dell Professional Services** provide turnkey interactive AI use case pilots using real customer data to validate business value ahead of scaled investments. These expert-led pilots offer a hands-on preview for experimentation with clear success metrics and KPIs, delivering tangible ROI.

Uplevel AI performance with next-generation infrastructure

Dell infrastructure updates accelerate HPC and AI innovation by delivering platforms with powerful performance, scalability and streamlined management. These solutions will help organizations build efficient systems for modern workloads.

- The **Dell PowerEdge XE8712** server, available next month, sets a new standard by delivering the industry's highest GPU density in a standard rack with up to 144 [NVIDIA Blackwell GPUs](#) per Dell IR7000 rack.⁵ This breakthrough enables next-generation AI and HPC at rack scale, complemented by unified rack-level automation and monitoring via Integrated Dell Remote Access Controller (iDRAC), OpenManage Enterprise (OME) and the Integrated Rack Controller (IRC) for

advanced thermal controls.

- **Enterprise SONiC Distribution by Dell Technologies now supports NVIDIA Spectrum-X platforms**, in addition to NVIDIA's Cumulus OS. Designed to bring hyperscale networking capabilities to a unified, multi-vendor environment, SONiC combines the flexibility of open-source technology with Dell's enterprise-grade features and trusted support. This empowers customers to confidently build open, standards-based AI infrastructure, delivering powerful performance and streamlined management for modern workloads.
- The latest release of **SmartFabric Manager** will now be extended to Dell's Enterprise SONiC on NVIDIA's Spectrum-X platforms as well, helping customers with faster, error-free setups that reduce deployment times significantly, with minimal manual intervention.

Expanded AI ecosystem offers enterprises choice

As enterprises right-size their AI investments, picking the right tools to help get the job done is a top priority. **Red Hat OpenShift for the Dell AI Factory with NVIDIA** is now validated on more Dell PowerEdge systems, helping enterprises operationalize AI at scale to transform business operations.

In addition to the Dell PowerEdge R760xa, the Dell PowerEdge XE9680 featuring NVIDIA H100 and H200 Tensor Core GPUs, is now supported, offering more choice for enterprises looking to accelerate their AI adoption at scale. This combination of Red Hat OpenShift tools, controls and governance with Dell's secure, trusted infrastructure lets organizations scale AI with confidence.

Perspectives

Jeff Clarke, vice chairman and chief operating officer, Dell Technologies: "The Dell AI Factory with NVIDIA solves the problem every enterprise is facing: how to move from AI pilots to production without rebuilding their infrastructure. We've done the integration work so customers don't have to, which means they can deploy faster and scale with confidence."

Justin Boitano, vice president, Enterprise AI products, NVIDIA: "Enterprise AI is shifting from experimentation to transformation—advancing at unprecedented speed and redefining how businesses operate. Together, Dell and NVIDIA are driving this evolution with a fully integrated platform that unites advanced infrastructure, intelligent automation, and powerful data engines to help organizations deploy AI at scale and realize measurable impact."

Ashish Nadkarni, group vice president and general manager, Infrastructure Systems, Platforms and Technologies, IDC: "As enterprises shift from AI experimentation to scaled deployment, they need infrastructure that is not only powerful but also integrated and easy to operationalize. The combination of Dell's end-to-end AI portfolio with NVIDIA's advanced technology represents a significant step forward in delivering enterprise-ready AI."

Availability

- Dell AI workload blueprints are available in tech preview now.
- Dell Services AI use case pilots are globally available now.
- Dell AI Factory with NVIDIA with Dell PowerEdge XE7740/XE7745 updates are globally available now.
- Dell ObjectScale and PowerScale with NVIDIA Dynamo is globally available now.
- Updates to Dell's ecosystem enablers for AI PCs are globally available now.
- The Dell PowerEdge XE8712 server is globally available in December.
- Enterprise SONiC Distribution by Dell Technologies support for NVIDIA Spectrum-X will be globally available in 1H26.
- Dell SmartFabric Manager with NVIDIA Spectrum-X support will be globally available 1H26.
- Updates for Red Hat OpenShift for the Dell AI Factory with NVIDIA are globally available now.

Additional resources

- [Find out more](#) about new Dell AI Factory innovations also announced today.
- [SC25: Accelerating Enterprise AI with the Dell AI Factory](#)
- [Dell AI Data Platform Innovations Announced at Supercomputing 2025](#)
- [Dell at SC25: Unified compute and networking solutions for AI and HPC with NVIDIA](#)
- [Open Ethernet for AI: NVIDIA Spectrum-X with Dell SONiC](#)
- [Scale and Automate AI success with Dell and NVIDIA](#)
- [Prove AI Value with Pilots that Work](#)
- [Powering progress with a new era of AI PC tools](#)
- 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 IDC Worldwide Quarterly Artificial Intelligence Infrastructure Tracker, 2025Q2

² Vanson Bourne, Dell Technologies survey of 750 IT and business decision makers across US, UK, DE, FR and JP, all segments, November 2025.

³ Vanson Bourne, Dell Technologies survey of 750 IT and business decision makers across US, UK, DE, FR and JP, all segments, November 2025.

⁴ Performance results are based on internal Dell testing using the vLLM inference engine, LMCache with Dell connector, and NVIDIA NIXL library. Tests were conducted with a 100% KV Cache hit rate on 4x NVIDIA H100 GPUs, using the LLaMA-3.3-70B Instruct model with Tensor Parallelism=4. Actual performance may vary based on system configuration, workload, and other factors. Oct, 2025

⁵ Based on Dell Technologies analysis of publicly available information, November 2025.

View source version on [businesswire.com](https://www.businesswire.com/news/home/20251117790858/en/): <https://www.businesswire.com/news/home/20251117790858/en/>

Media Relations: Media_Relations@Dell.com

Source: Dell Technologies