

## Cognizant Expands AI Infrastructure Capabilities with Launch of Cognizant AI Factory

- *Multi-tenant cloud offering, powered by Dell Technologies and NVIDIA, enables more secure, scalable AI deployment across hybrid and multi-cloud environments.*
- *Proprietary Cognizant fractional GPU technology built on NVIDIA Multi-Instance GPU helps enterprises operationalize AI faster, reduce costs and support the governance of AI workloads with confidence.*

TEANECK, N.J., March 16, 2026 /PRNewswire/ -- Cognizant ([NASDAQ: CTSH](#)) announced the launch of Cognizant AI Factory, a multi-tenant, enterprise-grade offering powered by Dell Technologies and NVIDIA AI infrastructure and software platform. Designed to help organizations scale artificial intelligence more securely, efficiently and responsibly, Cognizant AI Factory aims to unify the management of the AI lifecycle in a single environment – from ideation and experimentation to deployment, orchestration and day-to-day operations.

In line with Cognizant's strategy to be an AI Builder for the enterprise, the offering aims to create a path for enterprises to move from proof-of-concept to operationalized AI at scale, accelerating time to market and the ability to drive better outcomes. Powered by the Dell AI Factory with NVIDIA, the solution aims to bring together high-performance, enterprise-grade infrastructure and NVIDIA's advanced AI software stack to unify management of the AI lifecycle across private, public and hybrid cloud environments.<sup>1</sup>

A standout innovation is Cognizant's proprietary Fractional GPU technology, interoperable with [NVIDIA Multi-Instance GPU \(MIG\)](#) enabling secure, isolated GPU "slices" that allow multiple business units – or even multiple clients – to run AI workloads concurrently in a unified environment. This approach is designed to preserve data integrity, maximize utilization and help significantly lower the barriers to adopting advanced AI capabilities across an enterprise. Dell Technologies' proven expertise in delivering secure, scalable and high-performing AI infrastructure aim to help Cognizant AI Factory meet the rigorous demands of enterprise AI workloads, enabling organizations to innovate with confidence.

"Enterprises everywhere are racing to operationalize AI, but too often run into barriers around scale, cost and governance," said Sriram Kumaresan, Global Head of Cloud and Infrastructure Services, Cognizant. "Cognizant AI Factory changes the equation. By pairing best-in-class infrastructure with intelligent orchestration and enterprise-grade guardrails, we're supporting clients to turn AI into a durable engine for business value. The initiative also reflects Cognizant's evolution into a frontier organization, applying AI systematically within its own operations to continuously refine its platforms, governance models and delivery approaches before scaling them for clients."

Cognizant AI Factory is intended to give clients full control over where their AI runs, how it scales and how it's governed, while supporting clients in meeting emerging standards such as ISO/IEC 42001:2023. Provisioning, governance, and cost management are consolidated into a streamlined experience designed specifically for the complexity of enterprise AI workloads. Internal benchmarking and controlled testing of Cognizant AI Factory demonstrate the platform's potential to help achieve 50-60% lower total cost of ownership and up to 30% faster AI processing, compressing time-to-value from months to weeks.

Running on Dell AI Factory with NVIDIA, Cognizant offers a full-stack managed service for the complete AI lifecycle.

"Cognizant AI Factory helps enterprises turn AI's potential into real value and progress," said Denise Millard, Chief Partner Officer, Dell Technologies. "Powered by the Dell AI Factory with NVIDIA, a full stack solution, including Dell's high-performance PowerEdge Servers, PowerSwitch Networking, PowerScale and Services, we're giving businesses a clear, fast path from pilots to full-scale deployment."

Included among the features of Cognizant AI Factory are:

- Ready-to-use sandbox environments for experimentation and rapid internal pilots;
- Pre-built blueprints and MLOps pipelines to shorten time-to-value;
- An AI resiliency layer for monitoring, supporting governance and lifecycle management;
- Consumption-based pricing model to help enable predictable spend and optimized utilization.

"Across industries, enterprises need solutions to quickly move AI from proof-of-concept to a secure, governed and cost-efficient operational engine," said John Fanelli, Vice President, Enterprise Software at NVIDIA. "Together, NVIDIA, Cognizant and Dell Technologies are delivering the AI infrastructure and software foundation that equips organizations with the performance and flexibility needed to scale AI deployments with confidence."

Cognizant AI Factory is a key element of Cognizant's broader AI Builder strategy, providing clients a foundation to deploy intelligent agents, integrate advanced models and modernize legacy platforms.

### About Cognizant

Cognizant (NASDAQ: CTSH) is an AI builder and technology services provider, building the bridge between AI investment and

enterprise value by building full-stack AI solutions for our clients. Our deep industry, process and engineering expertise enables us to build an organization's unique context into technology systems that amplify human potential, realize tangible returns and keep global enterprises ahead in a fast-changing world. See how at [www.cognizant.com](http://www.cognizant.com) or @cognizant.

**For more information, contact:**

U.S.

Name Ben Gorelick

Email [benjamin.gorelick@cognizant.com](mailto:benjamin.gorelick@cognizant.com)

Europe / APAC

Name Sarah Douglas

Email [sarah.douglas@cognizant.com](mailto:sarah.douglas@cognizant.com)

India

Name Vipin Nair

Email [Vipin.nair@cognizant.com](mailto:Vipin.nair@cognizant.com)

<sup>1</sup> This includes Dell PowerEdge servers, PowerSwitch networking, PowerScale storage, and Dell Services, together with NVIDIA AI Enterprise, NVIDIA NIM microservices, NVIDIA NeMo for LLM lifecycle management, RAPIDS for accelerated data pipelines, and CUDA-X libraries, as well as Cognizant's proprietary Fractional GPU technology built on NVIDIA Multi-Instance GPU (MIG).

SOURCE Cognizant

---

<https://news.cognizant.com/2026-03-16-Cognizant-Expands-AI-Infrastructure-Capabilities-with-Launch-of-Cognizant-AI-Factory>