

Cognizant Neuro AI Trust delivers real-time assurance for enterprises scaling AI at speed

New command center helps enterprises trust and scale AI with confidence, delivering real-time visibility and supporting continuous governance across every model, agent and application

TEANECK, N.J., July 1, 2026 /PRNewswire/ -- Cognizant (Nasdaq: CTSH) today announced [Cognizant Neuro[®] AI Trust](#), a new platform designed to provide enterprises with continuous governance and real-time assurance across all AI systems. As AI environments grow more autonomous and complex, Neuro AI Trust empowers enterprises to monitor, manage and help control AI behavior and performance in real time, aiming to enable organizations to scale AI with confidence.

With enterprises deploying multiple AI models, multi-agent networks, and applications, managing visibility and risk is becoming more difficult as systems continuously evolve and interact with one another with increasing levels of human-defined autonomy. Governance approaches built for static systems cannot keep pace with the dynamic nature of AI. [According to Gartner[®]*](#), "organizations that deployed AI governance platforms are 3.4 times more likely to achieve effectiveness in AI governance than those that do not." Cognizant believes this reinforces the need for centralized platforms that enable continuous, real-time oversight across AI systems.

Neuro AI Trust addresses these challenges by introducing an interoperable control and intelligence layer for enterprise AI, purpose-built to give organizations a centralized way to oversee and manage increasingly complex AI environments across a wide range of models and agents. The control layer provides real-time observability across AI systems, using Guardian Agents to continuously monitor behavior, interactions and outcomes, aiming to deliver clear visibility into system health, performance, security and risk.

In parallel, the intelligence layer governs how these systems operate, evaluating interactions in real time and applying configured policies through centralized decisioning, guardrails and automated controls designed to align to business objectives and regulatory requirements. Insights and enforcement actions from both layers are brought together in a comprehensive dashboard, enabling organizations to identify issues early, take action with confidence and help reduce operational, regulatory, and reputational risk. Together, these capabilities aim to enable adaptive oversight as AI systems evolve and interact.

"As agentic AI moves into enterprise operations, the constraint is no longer capability but trust. Technology leaders expect governance, accountability and transparency to be addressed by AI platforms," said Jennifer Hamel, Research Vice President, Enterprise Data and AI Services at IDC. "Increasingly, organizations look to service providers for agentic AI platforms, such as Cognizant Neuro AI Trust, that combine technical integration, governed deployment and auditability as a strategic operating layer, not isolated tooling."

The Neuro AI Trust platform has already been deployed internally across Cognizant's [agentified intranet](#), serving its 350,000 employees.

"Neuro[®] AI Trust was built to govern AI as it actually behaves: autonomously, continuously, and across systems that interact in ways no single policy check can anticipate. We know it is effective because we have applied it to our own AI systems," said Amir Banifatemi, Chief Responsible AI Officer at Cognizant.

Neuro AI Trust leverages specialized multi-agent networks embedded across both the intelligence and control layers to continuously evaluate AI systems, interactions and workflows in real time. These agents operate across distinct domains such as policy enforcement, risk management and governance, enabling system-wide visibility and coordinated control across complex AI environments.

Neuro AI Trust is designed to enable enterprises to:

- **Gain end-to-end observability into every AI system:** A comprehensive trust score and full lifecycle observability give operators clear visibility into model behavior, agent interactions, and outcomes across the entire AI stack, including early detection of model drift and coordination risks that span multiple agents.
- **Deploy Guardian Agents for system-wide oversight** A dedicated multi-agent system continuously monitors agent interactions across steps, tools and turns, catching coordination failures such as escalation loops, circular disputes, risky tool use and emergent patterns that single-message checks would never surface.
- **Enforce policy across AI interactions:** The platform evaluates all AI interactions at runtime, returning permissive, warning or blocking outcomes based on configurations aligned with frameworks including NIST AI RMF, EU AI Act, OECD Principles and ISO/IEC 42001, as well as any internal custom policies.
- **Predict and surface risks before they escalate:** Neuro AI Trust is designed to move governance upstream, using signals from AI traces to anticipate potential policy violations earlier in the workflow lifecycle.
- **Update governance rules without code changes:** Policies, policy packs and risk thresholds are dynamically loaded at runtime, so compliance, legal and risk teams can update controls as requirements evolve, without waiting on a code

release.

- **Escalate to a human when necessary:** Higher-risk or ambiguous decisions can be paused and routed to a human reviewer with the full context needed to approve, reject, or request more information before any action is taken.
- **Build trust with audit-ready records:** Audit-ready records and replay views allow operators and auditors to reconstruct captured AI interactions in detail, understanding what happened, why it happened, which policy applied, and how the governance layer responded at every step.

Neuro AI Trust integrates with Cognizant's broader AI portfolio, including offerings such as the [Neuro[®] AI Multi-Agent Accelerator](#), as well as any other agentic application. Built on the [Cognizant Trust[™] framework](#), Neuro AI Trust helps AI systems operate in a transparent, fair, safe, accountable and reliable manner, advancing the responsible adoption of AI at scale. This reflects Cognizant's broader strategy as an AI Builder: helping enterprises maintain accountability for AI in production by providing centralized oversight, trust and governance.

For more information on Cognizant Neuro AI Trust, please [visit this page](#).

*Gartner Press Release, [Global AI Regulations Fuel Billion-Dollar Market for AI Governance Platforms](#), February 17, 2026

GARTNER is a trademark of Gartner, Inc. and/or its affiliates.

About Cognizant

Cognizant (Nasdaq: CTSH) is an AI Builder and technology services provider, bridging the gap 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, drive tangible outcomes and keep global enterprises ahead in a fast-changing world. See how at www.cognizant.ai or @cognizant.

For more information, contact:

U.S.	Europe / APAC	India
Name Paul Jarratt	Name Sarah Douglas	Name Vipin Nair
Email paul.jarratt@cognizant.com	Email sarah.douglas@cognizant.com	Email Vipin.nair@cognizant.com

SOURCE Cognizant Technology Solutions

<https://news.cognizant.com/2026-07-01-Cognizant-Neuro-AI-Trust-delivers-real-time-assurance-for-enterprises-scaling-AI-at-speed>