Trend Micro Launches End-to-End Protection for Agentic AI Systems with NVIDIA

Extends Agentic AI Safety from infrastructure to application with agentless EDR and integrated guardrails to secure next-generation AI factories

DALLAS, Oct. 28, 2025 /PRNewswire/ -- Trend Micro Incorporated (TYO: 4704; TSE: 4704), a global cybersecurity leader, today announced a new integration with NVIDIA BlueField, which embeds protection directly into the data center layer for more secure and scalable AI factories. This new offering enables organizations to deploy AI faster and reduce risks across multi-tenant AI clouds, all while meeting stringent compliance requirements.

To learn more about the combination of Trend Vision One™ Endpoint Security and NVIDIA BlueField DPUs, please visit: https://www.trendmicro.com/en_us/business/ai/factory.html

NVIDIA BlueField data processing unit (DPU) is a processor designed to offload, accelerate, and isolate infrastructure and cybersecurity tasks from the CPU. Trend Vision One™ Endpoint Security (AI Factory Endpoint Detection and Response) is deployed on NVIDIA BlueField to collect and monitor host and network information, and correlate with Trend threat intelligence to detect suspicious behavior. Complementing this BlueField integration, Trend is also among the first cybersecurity vendors to validate its solution on NVIDIA RTX PRO Servers, bringing purpose-built, enterprise-class security to AI factories.

This will also now include AI factories in federal agencies and high-assurance organizations, as Trend is included in the NVIDIA AI Factory for Government reference design, which provides full-stack, end-to-end guidance for deploying AI workloads while meeting the compliance needs of regulated industries.

Rachel Jin, Chief Enterprise Platform Officer at Trend Micro:"Agentic AI promises to unleash a new era of productivity, efficiency, and business agility, but only if we build it on secure foundations. That's why Trend is committed to advancing AI safety through innovation in zero-trust enforcement and AI-native threat detection. Our combined offering with NVIDIA will establish a new market standard for peak performance deployments."

Ofir Arkin, Sr. Distinguished Architect for Cybersecurity at NVIDIA: "As enterprises deploy AI factories, they need to secure large-scale, high-speed infrastructures without slowing innovation. By integrating with NVIDIA BlueField, Trend Vision One establishes a new class of endpoint detection and response for AI factories, combining hardware-enforced isolation with real-time threat insights to safeguard critical AI assets at the data center layer."

According to Gartner®, "AI infrastructure security includes the built-in security features of the underlying technology stack, such as vector and graph databases, and third-party security controls that could be easily expanded by incumbent vendors to cover AI security use cases."*

Building on this foundation, Trend is also extending protection to the application layer of Agentic AI.Trend Vision One™ AI Application Security (AI Guard) integrates natively with NVIDIA NeMo Guardrails, part of the NVIDIA NeMo framework, a scalable rail orchestration solution for ensuring the security, safety, accuracy, and topical relevance of LLM interactions. This integration streamlines how teams define, test, and orchestrate AI guardrails, including multimodal rails through a microservice and simple APIs. This joint capability enables security teams to align guardrails with enterprise policy, map them to key risks, such as prompt injection, data leakage, tool or agent abuse, jailbreaks, and hallucinations, and enforce them consistently from development to runtime.

The platform ingests guardrail telemetry for observability, risk scoring, and incident response, then automates policy-as-code updates and playbook-driven remediation across the AI stack and cloud. This combined approach detects credential dumps, reverse shells, and other advanced threats while strengthening Agentic AI safety in three critical areas:

Content Moderation: Filters toxic or biased AI outputs without slowing inference.

Security: NVIDIA BlueField-accelerated, hardware-enforced isolation prevents prompt injection and jailbreak attempts.

Privacy: Built-in encryption, GDPR, HIPAA, and CCPA compliance, and zero-trust segmentation.

** Gartner, Use an AI Security Platform to Launch Your AI Security Strategy,Dennis Xu, Kevin Schmidt, Jeremy D'Hoinne, 26 February 2025*

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally, Magic Quadrant is a registered trademark of Gartner, Inc. and/or its affiliates and is used herein with permission. All rights reserved.

About Trend Micro
Trend Micro, a global cybersecurity leader, helps make the world safe for exchanging digital information. Fueled by decades of

security expertise, global threat research, and continuous innovation, Trend Micro's AI-powered cybersecurity platform protects hundreds of thousands of organizations and millions of individuals across clouds, networks, devices, and endpoints. As a leader in cloud and enterprise cybersecurity, Trend's platform delivers a powerful range of advanced threat defense techniques optimized for environments like AWS, Microsoft, and Google, and central visibility for better, faster detection and response. With 7,000 employees across 70 countries, Trend Micro enables organizations to simplify and secure their connected world. www.TrendMicro.com.

SOURCE Trend Micro Incorporated

For further information: Media Contact: Trend Micro Communications, 817-522-7911, media_relations@trendmicro.com

---

https://newsroom.trendmicro.ca/2025-10-28-Trend-Micro-Launches-End-to-End-Protection-for-Agentic-AI-Systems-with-NVIDIA