FOR IMMEDIATE RELEASE

Wedge Networks Launches WedgeSecure Agent: A Sovereign Workflow-Governance Platform for Securing Agentic AI at the Edge

First agentic AI security platform purpose-built for sovereign edge deployment, protecting intelligence at the point of contact, powered by Intel Xeon 6 processors and packaged on Advantech Edge AI hardware

---

CALGARY, Canada – May 30, 2026 – Wedge Networks Inc. today announced the early release of WedgeSecure Agent, a workflow-governance platform that secures autonomous AI systems against a new class of cybersecurity threat: attacks that target the reasoning and judgment of AI agents rather than the software infrastructure they run on.

*"Enterprise AI has crossed a line. We are no longer building systems that answer questions; we are deploying systems that plan, delegate, and act – and the security boundary has moved with them. The integrity of an agent's judgment is now the asset that has to be defended, and conventional cybersecurity tools were never built to see inside that judgment," said **Dr. Hongwen Zhang, CEO of Wedge Networks**. "WedgeSecure Agent is our answer. It is a sovereign, CPU-native platform that governs agentic AI workflows where they actually run – at the edge, on customer-controlled infrastructure, in the regulated environments that cannot send their data or their decisions to someone else's cloud."*

Sovereign Security for a New Threat Class

Traditional cybersecurity products protect networks, endpoints, and applications against software-level exploits. With the proliferation of Agentic AI, there emerges a new class of threats that attacks the reasoning of AI agents. The guardrails shipped by AI framework and LLM vendors are best-effort based demo-grade solutions. They do not provide the holistic security and trust that govern what an autonomous agent intends to do, what context it reasons over, which model it consults, what output it produces, or what real-world actions it takes.

WedgeSecure Agent is purpose-built to protect against this new class of threat. It operates as a set of containerized Policy Decision Point microservices that evaluate trust at consequential checkpoints across all trust domains, Intent, Context, Model, Content, and Action, and returns structured Decision Bundles that existing applications, agent frameworks, network controls, or cyber-physical systems can enforce. The platform does

not replace existing security infrastructure or agentic frameworks; it plugs into them at the points they already expose.

At the core of the platform is the WedgeSecure Unified Semantic Core (WedgeSecure-USC), a fine-tuned open-weight language model substrate that gives all trust services a shared semantic basis. This unified foundation enables WedgeSecure Agent to correlate weak signals across workflow stages, maintain judgment consistency across an agentic session, and produce audit-grade evidence for enterprise governance, capabilities that stitched-together point classifiers cannot deliver.

The core capabilities provided in the early release platform include:

- Prompt Injection Classification (PIC) detects adversarial instructions, jailbreaks, role-hijack attempts, and indirect injection embedded in tool outputs or retrieved content.
- Content Safety Classification (CSC) evaluates inputs and outputs against a granular multi-category taxonomy spanning harmful content, regulated topics, and operational risk categories aligned with emerging AI safety standards.
- One of the industry's first security-rating databases scoring leading LLMs across multiple security and trust attributes – not just leaderboard quality, but adversarial robustness, judgment-manipulation resistance, content-safety behavior, and stability under attack.
- Native OpenTelemetry instrumentation: every decision, classification verdict, and policy evaluation is emitted as standards-compliant traces, metrics, and logs – plugs directly into existing SIEM, SOC, and compliance pipelines – no proprietary telemetry silo.

WedgeSecure Agent is packaged as self-contained, easy-to-install, fully-supported microservices in Docker containers. The platform serves as a PDP layer for three deployment patterns:

- Embedded alongside agentic applications
- Inline at network gateways protecting fleets of agentic workloads, and
- Integrated with cyber-physical systems running autonomous agentic workflows.

Agentic AI Security with Trust Inferencing Powered by Intel Xeon 6 Processors

WedgeSecure Agent runs its trust-classification stack on Intel Xeon 6 processors, using a trust-fine-tuned open-weight LLM accelerated through Wedge Networks' Standing-Wave edge AI inference optimization technology.

Utilizing Intel Advanced Matrix Extensions (Intel AMX) present in Intel Xeon 6 processors, this CPU-only capability is foundational to WedgeSecure Agent's sovereign deployment posture. Organizations in defense, government, healthcare, finance, telecommunications, and industrial automation can run the full WedgeSecure Agent stack on their own infrastructure, in their own jurisdiction, with no dependency on hyperscaler model services and no requirement for customer data or policy decisions to leave their control plane. For these sectors, sovereignty is not a feature – it is a deployment prerequisite.

“Intel Xeon 6 with Intel AMX is enabling a new class of agentic AI security workload to run at the edge without discrete accelerators specific to AI,” commented ***Srini Krishna, Intel Fellow, Data Center Products, Intel Corporation****.”This work with Wedge Networks demonstrates how CPU-native inference and hardware-rooted security can meet the sovereignty and compliance requirements of the world's most regulated industries."*

Turn Key Agentic AI Security solution with Advantech Edge AI

WedgeSecure Agent ships as a containerized stack that runs on Advantech's industrial-grade Edge AI platforms. This places agentic-AI policy decisioning inside the operational environments Advantech already anchors – telco aggregation sites, energy substations, ITS roadside cabinets, defense field nodes, smart-factory control rooms – within their existing power, thermal, lifecycle, and sovereignty constraints. Together, Intel’s CPU-native inference and Advantech’s Edge AI hardware form a validated, turnkey route from architecture to sovereign-edge deployment.

"Advantech's edge computing platforms are now the physical foundation for a new category of AI workload – agentic AI governance,” commented ***Senior Director Sven Freudenfeld at Advantech****. “The Advantech FWA and AES product family provides the foundation for an economically valuable solution with agentic AI in a single platform without the need for costly, purpose-built processing of agentic AI. With this technology enhancement of edge appliances, Wedge Networks brings security intelligence to the edge locations where Advantech hardware is already trusted by enterprises in telecommunications, energy, defense, and industrial automation.”*

Availability and Pilot Programs

WedgeSecure Agent v1 is available in early release to pilot partners across the telecommunications, intelligent transportation, smart grid and healthcare industry sectors.

WedgeSecure Agent v1 with today’s announced capabilities is targeted for production release in Q3–Q4 2026. Subsequent releases will deliver the additional capabilities that implement the full vision set out in the WedgeSecure Agent white paper, “Trustworthy Autonomy at the Agentic Edge.”

Organizations building or deploying agentic AI where security, compliance, or audit-readiness is a requirement are invited to contact Wedge Networks to discuss pilot engagement.

About Wedge Networks

Wedge Networks Inc. is a Calgary-based cybersecurity company that develops network security and AI security products for enterprise and government customers. The WedgeSecure product family includes WedgeSecure Cloud, Edge, Guard and Wi-Fi for network-layer threat prevention, and WedgeSecure Agent for agentic AI workflow governance.

For more information, please visit https://www.WedgeNetworks.com or contact Wedge Networks Media Relations at Info@WedgeNetworks.com

About Advantech

Advantech's corporate vision is to enable an intelligent planet. The company is a global leader in the fields of IoT intelligent systems and embedded platforms. To embrace the trends of IoT, big data, and artificial intelligence, Advantech promotes IoT hardware and software solutions with the Edge Intelligence WISE-PaaS core to assist business partners and clients in connecting their industrial chains. Advantech is also working with business partners to co-create business ecosystems that accelerate the goal of industrial intelligence. (www.advantech.com)

Advantech Media Contacts

- Sven Freudenfeld - Senior Director Business Development Edge and Edge AI
  - Sven.Freudenfeld@advantech.com or +1 514-912-2903
- Matt Tsai - Product Manager of Network Security
  - Matt.Tsai@advantech.com.tw or +886-2-7732-3399 Ext. 1613
- Iris Lee - Senior Marketing Specialist of Network Security
  - Iris1.Lee@advantech.com.tw or +886-2-2792-7818, Ext. 1329

###