



CrowdStrike Unveils Secure-by-Design AI Blueprint for AI Agents Built with NVIDIA

Architecture will integrate protection from the Falcon platform with NVIDIA OpenShell to run safer, autonomous AI agents both locally on DGX Spark and in the cloud

AUSTIN, Texas & SAN JOSE, Calif.--(BUSINESS WIRE)--Mar. 16, 2026-- GTC – [CrowdStrike](#) (NASDAQ: CRWD) today unveiled a Secure-by-Design AI Blueprint built with [NVIDIA](#) that integrates protection from the [CrowdStrike Falcon® platform](#) directly into NVIDIA OpenShell, an open-source runtime that enforces policy-based guardrails to make autonomous agents safer to deploy.

The architecture integrates security natively into the AI agent stack, enabling organizations to operationalize autonomous systems with governance, visibility, and control from development through runtime, wherever agents run.

As organizations shift from copilots to AI agents that think, reason, and act autonomously at machine speed, security models must evolve. AI agents introduce a fundamentally different security challenge as privileged identities with direct access to data, applications, compute resources, and other agents. Traditional static controls were not designed to govern systems that move at the speed of AI. Securing AI agents requires continuous enforcement across the AI stack, not point in time controls – delivered at machine speed.

By integrating the Falcon platform directly into the NVIDIA OpenShell runtime, the Secure-by-Design AI Blueprint can embed security at the foundation of autonomous systems. Part of the NVIDIA Agent Toolkit, the open-source OpenShell runtime provides isolated sandboxes with private inference and built-in policy enforcement. The Falcon platform extends protection to local agents running on NVIDIA DGX Spark or NVIDIA DGX Station, and can also extend security to agents in the cloud that are built on the open-source NVIDIA AI-Q Blueprint for deep research.

Organizations will gain unified visibility and continuous runtime monitoring and enforcement to constrain unsafe behavior, prevent prompt manipulation, and enforce policy across the full AI lifecycle.

Key capabilities of the Secure-by-Design AI Blueprint include:

- **AI Policy Enforcement Across the Agent Stack:** [Falcon® AI Detection and Response \(AIDR\)](#) will integrate with the OpenShell runtime to secure every prompt, response, and agent action in real time.
- **Endpoint Protection for Local AI Agents:** [Falcon® Endpoint Security](#) will secure local agents on NVIDIA DGX Spark or DGX Station running OpenShell, enforcing host-level controls and continuous behavioral monitoring across system activity and agent execution.
- **Cloud Runtime Protection for AI Agent Deployments:** [Falcon® Cloud Security](#) will protect agents built based on the NVIDIA AI-Q Blueprint in cloud and data center environments, delivering unified visibility and runtime controls across infrastructure and AI workloads.
- **Identity-Based Governance for Agent Access:** [Falcon® Next-Gen Identity Security](#) will deliver dynamic identity management for local agents, enforcing access controls across data, APIs, and services so agents operate within defined privilege boundaries.

CrowdStrike and NVIDIA are also advancing intent-aware controls that govern how agents plan and execute tasks, enabling flexible autonomy while limiting the blast radius of unintended or malicious behavior.

“As we enter the agentic era, agents no longer simply assist – they act,” said Daniel Bernard, Chief Business Officer, CrowdStrike. “This shift fundamentally changes the security equation, and security must be embedded into the AI stack itself. Together with NVIDIA, we are delivering a Secure-by-Design architecture that enables organizations to operationalize agents with confidence and control.”

“Autonomous agents will fundamentally reshape how we work,” said Justin Boitano, Vice President, Enterprise Platforms, NVIDIA. “By integrating CrowdStrike’s security platform with the NVIDIA Agent Toolkit, we’re enabling enterprises to build and scale safer, autonomous AI agents to help transform their operations, empower every employee, and securely generate intelligence at the speed of business.”

“AI infrastructure is moving from experimentation to mission-critical production,” said James Higgins, Chief Information Security Officer, CoreWeave. “As we scale GPU-accelerated environments, AI agents must be observable, governed, and resilient by design. The collaboration between CrowdStrike and NVIDIA secures AI systems at the foundation – enabling high-performance AI environments without compromising control.”

The Secure-by-Design AI Blueprint reinforces CrowdStrike’s position as *cybersecurity for enterprise AI* – embedding security directly into the AI stack, wherever AI lives.

About CrowdStrike

[CrowdStrike](#) (NASDAQ: CRWD), a global cybersecurity leader, has redefined modern security with the world’s most advanced cloud-native platform for protecting critical areas of enterprise risk – endpoints and cloud workloads, identity and data.

Powered by the CrowdStrike Security Cloud and world-class AI, the CrowdStrike Falcon® platform leverages real-time indicators of attack, threat intelligence, evolving adversary tradecraft and enriched telemetry from across the enterprise to deliver hyper-accurate detections, automated protection and remediation, elite threat hunting and prioritized observability of vulnerabilities.

Purpose-built in the cloud with a single lightweight-agent architecture, the Falcon platform delivers rapid and scalable deployment, superior protection and performance, reduced complexity and immediate time-to-value.

CrowdStrike: We stop breaches.

Learn more: <https://www.crowdstrike.com/>

Follow us: [Blog](#) | [X](#) | [LinkedIn](#) | [Instagram](#)

Start a free trial today: <https://www.crowdstrike.com/trial>

© 2026 CrowdStrike, Inc. All rights reserved. CrowdStrike and CrowdStrike Falcon are marks owned by CrowdStrike, Inc. and are registered in the United States and other countries. CrowdStrike owns other trademarks and service marks and may use the brands of third parties to identify their products and services.

Forward-Looking Statements

This press release contains forward-looking statements that involve risks and uncertainties, including statements regarding a Secure-by-Design AI Blueprint for AI Agents and the benefits of such deployments to CrowdStrike and its customers. You should not place undue reliance on these forward-looking statements, as actual outcomes and results may differ materially from those contemplated as a result of risks and uncertainties. There are a number of factors that could cause actual results to differ materially from statements made in this press release, including the risks and uncertainties described in the filings CrowdStrike makes with the Securities and Exchange Commission from time to time, including CrowdStrike's most recently filed Annual Report on Form 10-K, most recently filed Quarterly Report on Form 10-Q, and subsequent filings. All forward-looking statements in this press release are based on information available to CrowdStrike as of the date hereof, and CrowdStrike does not assume any obligation to update any of these forward-looking statements to reflect events that occur or circumstances that exist after the date on which they were made.

View source version on [businesswire.com](https://www.businesswire.com/news/home/20260316048931/en/): <https://www.businesswire.com/news/home/20260316048931/en/>

Media Contact

Jake Schuster

CrowdStrike Corporate Communications

press@crowdstrike.com

Source: CrowdStrike