CROWDSTRIKE
Investor Relations

# CrowdStrike Launches Frontier Models for Cybersecurity, Created with NVIDIA

September 1, 2026

*CrowdStrike SafeMind, the first agentic system for defenders, built with NVIDIA Nemotron, protects more and costs less than generic frontier models*

AUSTIN, Texas & LAS VEGAS--(BUSINESS WIRE)--Sep. 1, 2026-- **Fal.Con 2026 –** CrowdStrike (NASDAQ: CRWD) today introduced CrowdStrike SafeMind, a family of purpose-built security models and harnesses from the CrowdStrike Cyber Superintelligence Lab. The SafeMind agentic system will operate natively in the CrowdStrike Falcon® platform. Trusted access for standalone models and harnesses will be part of the Project QuiltWorks program.

**SafeMind – Delivering Cybersecurity's First Agentic System Built for Defenders**

Unlike frontier AI models, SafeMind runs as one system designed to deliver AI safety: an offensive model that finds the attack path, a defensive model that closes it, and the harnesses that operate both in the same loop.

SafeMind model training data comes from CrowdStrike Falcon sensor telemetry, the world's largest pureplay cyber dataset and edge install base. Model training also includes CrowdStrike's threat intelligence, Falcon Complete MDR event annotations, and fifteen years of incident response fieldwork, where human responders stopped breaches on the front lines. SafeMind launches with two distinct models, delivering unique security competencies:

- **Red Tempest:** Offensive red team model release, built for advanced attack scenarios, emulating AI adversaries.
- **Blue Solano:** Defensive blue team model release, built for protecting enterprise assets by deploying battle-tested measures that defenders use in real-life.

CrowdStrike builds the models using NVIDIA Nemotron open models in collaboration with NVIDIA, its AI design partner. The program also includes CoreWeave's AI Cloud for training and inference.

SafeMind's harnesses operationalize both the red and blue CrowdStrike models in a closed-loop system that continuously pits the models against each other to improve. The harnesses also work with frontier and open-source models, maximizing user choice and model preference while maintaining cost control. Frontier labs can tell a defender a risk exists. CrowdStrike goes beyond with the harnesses that can autonomously act on risk.

"The future of cybersecurity won't be defined by AI that simply identifies threats, it will be defined by AI that defeats them," said George Kurtz, CEO and founder of CrowdStrike. "SafeMind brings offensive and defensive models together in a system trained on CrowdStrike's unique cyber data. It finds weaknesses, strengthens protection, and gets smarter with every cycle, advancing our mission to stop breaches at machine speed."

"Cybersecurity in the age of AI will be a continuous contest between adversaries using AI to scale attacks and defenders using AI to expand detection and response. Cyber defense will be among the most compute-intensive applications of AI," said Jensen Huang, founder and CEO of NVIDIA. "SafeMind combines NVIDIA Nemotron open models with CrowdStrike's deep cybersecurity expertise, trusted security data, purpose-built agent harnesses, rigorous evaluations, and safeguards – creating a frontier agentic cybersecurity stack designed to operate at machine speed. Powered by NVIDIA accelerated computing, SafeMind makes AI a force multiplier for defenders."

"The real test of AI is what it can do in production, at scale, when the stakes are highest. Few environments put that to the test more than cybersecurity," said Michael Intrator, co-founder and chief executive officer, CoreWeave. "We're proud to power SafeMind across training and inference as CrowdStrike puts specialized AI to work against real-world threats."

**What Evaluations Show**

Compared to leading frontier models and open-source baselines, SafeMind delivers:

- 29% higher detection rate
- 6x faster end-to-end remediation
- 99% cost savings on detection and remediation

"Our models are the start of a new chapter for cyberdefense," said Dr. Bartley Richardson, chief AI and autonomous systems officer at CrowdStrike. "With the models and harnesses together in a co-evolving agentic system, defenders can now act at machine speed. This is the foundation for the next decade of AI security, and CrowdStrike is the only company that owns the entire stack, from sensor to harness to model."

**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



View source version on businesswire.com: https://www.businesswire.com/news/home/20260901951891/en/

**Media Contact**
Jake Schuster
CrowdStrike Corporate Communications
press@crowdstrike.com

Source: CrowdStrike