



CONFERENCE DISPATCH

DAY 1

CrowdStrike Stops Renting Its Intelligence

Day Zero widened the pipes. Day 1 filled them with something CrowdStrike built itself: **two frontier security models, a research lab to keep making them, and a runtime control plane for AI agents.** The company that runs its own discovery pipeline on OpenAI and Anthropic is now shipping models of its own.

6

ANNOUNCEMENTS

2

FRONTIER MODELS

1

RESEARCH LAB

\$2B

OPTIV TCV

FEATURING

CrowdStrike

Fal.Con 2026 · Mandalay Bay, Las Vegas · 1 September 2026

AUTHOR

Vatsal Sharma

PEER REVIEWED BY

Pinki Sharma, Devesh Taneja

PUBLISHED

September 2, 2026

CONTACT

vatsal@fieldciso.com

The Shape of Day 1

Six announcements, and three of them are the same announcement told three ways.

Day Zero was about distribution: where Falcon could be bought, who was paid to sell it, and whose data flowed in. Day 1 is about what CrowdStrike puts inside that expanded surface. Six announcements landed, and they split cleanly. Three are the company building its own AI capability. Three extend it through partners.

The three that matter are one story. **CrowdStrike announced that it now builds frontier models rather than only consuming them.** SafeMind is the model family, the Cyber Superintelligence Lab is the organisation that produced it and will produce the next one, and Falcon Guardian is the enforcement layer those models are meant to feed. Announced separately, they describe a single vertical: research, model, control point.

That is a meaningfully different posture from Day Zero. Project QuiltWorks, the initiative CrowdStrike built its Day Zero product around, runs its discovery step on models from OpenAI and Anthropic. Twenty-four hours later the company introduced models of its own, built with NVIDIA, trained on CoreWeave. It has not stopped being a customer of the frontier labs. It has started being a producer alongside them, in one narrow domain.

ANNOUNCEMENT	TYPE	WHAT IT ACTUALLY MOVES
CrowdStrike SafeMind	Models	Two purpose-built security models, trained adversarially
Cyber Superintelligence Lab	Organisation	The research group that produced SafeMind and succeeds it
Falcon Guardian	Product	Discovers and controls AI agents at the endpoint
Google Cloud enterprise AI	Integration	Puts Falcon controls inside Gemini Enterprise workflows
EY.ai Trust Layer	Route to market	Falcon becomes the security layer under EY's AI framework
Optiv milestone	Commercial	Evidence that Falcon Flex is driving platform consolidation

SafeMind and the Lab

Two models that attack and defend each other, and the research organisation built to keep them improving.

Cyber Superintelligence Lab ORGANISATION

1 September 2026 · CrowdStrike press release, via BusinessWire

CrowdStrike describes the Lab as the first frontier AI research organisation built for cyber defence and AI safety. It puts the company's AI researchers, offensive operators and incident responders under a single charter, and is led by Dr. Bartley Richardson, chief AI and autonomous systems officer.³

The stated input is CrowdStrike's data. The Falcon sensor reports from endpoints, identity systems, cloud workloads, data stores and Falcon Next-Gen SIEM into one architecture, producing trillions of events daily, and inside that sit 15 years of threat intelligence and incident response records where analysts have attached verified verdicts to what an attack was, how it worked and what stopped it.³

THE ARGUMENT UNDERNEATH

CrowdStrike's claim is that volume alone does not teach a defensive model to be correct, and that confirmed outcomes do. That is a real distinction and the strongest version of its data moat argument. It is also unfalsifiable from outside, since nobody else can inspect the corpus.

CrowdStrike SafeMind **LAUNCH**

1 September 2026 · CrowdStrike press release, via BusinessWire

SafeMind is a family of purpose-built security models and harnesses produced by the Lab. It launches with two models. Red Tempest is offensive, built to emulate AI-driven adversaries and run attack scenarios against an environment. Blue Solano is defensive, applying the containment measures CrowdStrike responders use on live incidents. The harnesses run both in a closed loop that pits one against the other so each improves, and the same harnesses can drive frontier and open-source models from other providers.⁷

SafeMind operates natively in the Falcon platform. Standalone access to the models and harnesses is handled through Project QuiltWorks.^{2,7} NVIDIA is the AI design partner, with Jensen Huang framing cyber defence as among the most compute-intensive applications of AI, and CoreWeave provides the cloud for training and inference.⁷

WHAT IS ACTUALLY NEW

Adversarial self-play is not a new idea, but applying it to a paired attack and containment model, seeded with real incident response tradecraft, is a specific and checkable architecture rather than a generic AI claim. That specificity is the most credible thing about the announcement.

THE BENCHMARK PROBLEM

CrowdStrike published performance claims for SafeMind: a 29 percent higher detection rate, six times faster end-to-end remediation, and 99 percent cost savings on detection and remediation compared with leading frontier models and open-source baselines. Trade coverage noted that the announcement did not provide further methodological detail.⁸

Treat all three as vendor-run evaluations against unnamed comparators. The 99 percent cost figure in particular is the kind of number that depends entirely on what was measured and against what. Several secondary outlets have also circulated different figures for the same product. Until a methodology or an independent evaluation appears, these are marketing claims, not results.

03

Falcon Guardian

The shipping product of Day 1, and the one with the clearest problem behind it.

Falcon Guardian **LAUNCH**

1 September 2026 · CrowdStrike press release, via BusinessWire

Falcon Guardian extends runtime endpoint security and visibility across AI agents, covering data, models, prompts and identities across cloud, SaaS and browser environments under one architecture.⁹

CrowdStrike's framing is that agents should be secured where they execute, which is the endpoint, rather than only at the prompt layer where its earlier AI detection and response work operated.¹

Four capabilities are named. Discovery and inventory identifies AI agents across Windows and macOS, including agents security never approved, and records who deployed each one. Runtime visibility traces a causal chain from prompt, identity, tool call and skill use through to the downstream system action. Access controls determine which agents are permitted to run and block the rest. Runtime detection and response reconstructs the execution chain during an attack and determines blast radius.¹

WHY THE ENDPOINT FRAMING MATTERS

The problem Guardian describes is arrival, not misuse. Agents reach corporate machines through abandoned pilots, developer-installed coding assistants and vendor desktop applications, and none of those routes involve a security approval. Because Guardian runs on the Falcon sensor already deployed, there is no second agent to install, which removes the usual reason a discovery tool never gets rolled out. CrowdStrike states that sensor runs on hundreds of millions of devices, which is the company's own figure.¹

04

The Ecosystem Layer

Three announcements that extend the above through other people's platforms and practices.

Falcon across Google Cloud's enterprise AI ecosystem

INTEGRATION

1 September 2026 · CrowdStrike and Google Cloud

Distinct from the Day Zero announcement that put Falcon on Google Cloud infrastructure. This one places Falcon controls inside Google's AI stack. Falcon Guardian extends through Google Agent Gateway to bring runtime protection to enterprise AI applications built on Google Cloud. Falcon MCP and Charlotte AI extend into Gemini Enterprise, and Falcon Shield covers Agent Registry.⁴

WHY IT MATTERS

Guardian's value depends on being present wherever agents run. Shipping it into a hyperscaler's agent gateway on the same day it launches is the fastest available route to that coverage.

EY.ai Value Blueprints Trust Layer

ROUTE TO MARKET

1 September 2026 · CrowdStrike and Ernst & Young LLP

EY US selected CrowdStrike as an accelerator to the Trust Layer of EY.ai Value Blueprints, integrating Falcon into EY's AI transformation framework. The stated components are a unified security foundation across identities, cloud, data and AI workloads, and protecting AI from development through runtime using Falcon Guardian. It builds on the companies' existing QuiltWorks collaboration.⁵

WHY IT MATTERS

This is the Day Zero partner strategy paying out one day later. EY was already a QuiltWorks remediation partner, and the relationship has now moved from delivering CrowdStrike's work to embedding Falcon inside EY's own framework, which is a better position for CrowdStrike to hold.

Optiv passes \$2 billion in total contract value

COMMERCIAL

1 September 2026 · CrowdStrike and Optiv

The companies announced they have surpassed \$2 billion in lifetime total contract value, a milestone reached in less than half the time it took them to reach the first \$1 billion. Both sides attribute it to customers standardising on the platform, with Optiv CEO Kevin Lynch naming Falcon Flex as the mechanism.⁶

WHY IT MATTERS

The acceleration is the number worth noting, not the total. A second billion arriving in under half the time of the first is the clearest external evidence offered this week that Falcon Flex is genuinely changing how quickly customers consolidate.

05

What Holds Up, What Does Not Yet

An assessment of the Day 1 announcements as published. FieldCISO Advisory has not had hands-on access to any of them.

HOLDS UP

- Guardian runs on the installed Falcon sensor, so the deployment friction argument is real rather than rhetorical.
- SafeMind's architecture is specific and named: two models, a closed loop, a stated training and inference provider.
- The harnesses are documented as working with third-party frontier and open-source models, which is a more honest position than full replacement.
- The Lab has a named leader, a defined charter, and a stated data input.
- Optiv's milestone includes a comparison point, which is rarer than the number itself.
- EY builds on an existing QuiltWorks relationship rather than starting from a press release.

NOT YET PROVEN

- SafeMind's three headline performance claims were published without methodology or named comparators.
- No independent evaluation of Red Tempest or Blue Solano exists.
- No pricing for standalone SafeMind access, which is gated behind Project QuiltWorks.
- The sensor deployment figure is CrowdStrike's own and not independently verified.
- The Lab has published no research output yet, so its claim to be a frontier research organisation rests on intent.
- Secondary outlets circulated AI application and endpoint counts that appear nowhere in the releases.

06

The Question Day 1 Raises

A security vendor becoming a model producer is a strategic bet, not a product decision.

CrowdStrike is now simultaneously a customer, a partner and a competitor of the frontier labs. Its own QuiltWorks discovery pipeline runs on OpenAI and Anthropic models. Its new SafeMind harnesses are documented as driving frontier and open-source models from other providers. And its own two models are

benchmarked against, in its words, leading frontier models and open-source baselines. All three relationships are live at once.

Read generously, that is a hedge rather than a contradiction. CrowdStrike is not claiming general models are useless. It is claiming that on a narrow set of security tasks, models trained on verified incident outcomes beat general models, and it has kept the option to run either. That is a defensible position and probably the right one.

The harder question is which of the two assets announced on Day 1 actually compounds. **Models are perishable.** A model family launched in September 2026 will be unremarkable within eighteen months, and CrowdStrike does not control the pace of the underlying capability curve. **The Lab and the data are not perishable.** Verified incident outcomes accumulate, and a research organisation that keeps producing successors is a more durable advantage than any particular model it ships.

Our read is that the Cyber Superintelligence Lab is the more significant of the two announcements, and it received a fraction of the coverage that SafeMind did. If CrowdStrike's bet is correct, the Lab is what makes it repeatable. If the bet is wrong, the Lab is what lets the company find out quickly and change course.

THE TEST TO APPLY

Watch for two things. First, whether CrowdStrike publishes a methodology or submits SafeMind to independent evaluation, which would convert the benchmark claims into evidence. Second, whether the Lab publishes research that is useful to defenders who are not CrowdStrike customers. A frontier research organisation that only produces product features is a product team with a better name.

07

What Day 1 Set Up

Day Zero widened distribution. Day 1 supplied the thing to distribute and the control point to enforce it. Read together, the two days describe a company trying to hold the layer where AI agents are discovered, governed and stopped, and to reach that layer through every hyperscaler, integrator and reseller it can sign.

Day 2 brings Mike Sentonas on platform and Adam Meyers on threat intelligence, plus an investor briefing. Meyers is the one to watch for anyone assessing the claims above, since adversary research is where the case for domain-specific security models either gains evidence or does not. Day 3 closes with Bartley Richardson, who now also runs the Lab, and Alex Ionescu.

Coverage note. This report covers announcements dated 1 September 2026 only. Sources swept: CrowdStrike newsroom, both investor relations feeds, BusinessWire, PR Newswire, partner newsrooms, and security trade press. Partner-led releases are distributed on different wires and do not all reach the investor relations feeds, so the newsroom index was treated as the authoritative list. Not reachable at time of writing: session content not published as text, and independent evaluation of any product announced. Performance figures attributed to CrowdStrike have not been verified by FieldCISO Advisory and are identified as vendor claims throughout.

References

1. CrowdStrike. "CrowdStrike Unveils Falcon Guardian to Secure AI Agents Where They Execute: On the Endpoint at Runtime." Press release, September 1, 2026.
<https://www.crowdstrike.com/en-us/press-releases/crowdstrike-unveils-falcon-guardian-ai-agent-security/>
2. CrowdStrike. "CrowdStrike Launches Frontier Models for Cybersecurity, Created with NVIDIA." Press release, September 1, 2026.
<https://www.crowdstrike.com/en-us/press-releases/crowdstrike-launches-frontier-models-for-cybersecurity-with-nvidia/>
3. CrowdStrike. "CrowdStrike Establishes Cyber Superintelligence Lab, Bringing Frontier AI Research to Cyber Defense for the First Time." Press release, September 1, 2026.
<https://www.crowdstrike.com/en-us/press-releases/crowdstrike-establishes-cyber-superintelligence-lab/>
4. CrowdStrike. "CrowdStrike Extends Falcon Platform Capabilities Across Google Cloud's Enterprise AI Ecosystem." Press release, September 1, 2026.
<https://www.crowdstrike.com/en-us/press-releases/crowdstrike-extends-falcon-platform-across-google-cloud-enterprise-ai-ecosystem/>
5. EY. "EY selects CrowdStrike to power the trust layer of EY.ai Value Blueprints." Newsroom, September 1, 2026.
https://www.ey.com/en_us/newsroom/2026/09/ey-selects-crowdstrike-to-power-the-trust-layer-of-ey-ai-value-blueprints
6. CrowdStrike and Optiv. "CrowdStrike and Optiv Surpass \$2 Billion in Total Contract Value." Press release, September 1, 2026.
<https://www.crowdstrike.com/en-us/press-releases/crowdstrike-optiv-surpass-2-billion-total-contract-value/>
7. SiliconANGLE. "CrowdStrike builds security frontier models with Nvidia and opens an AI lab." September 1, 2026.
<https://siliconangle.com/2026/09/01/crowdstrike-builds-security-frontier-models-with-nvidia-and-opens-an-ai-lab/>
8. IT Brief Asia. "CrowdStrike launches SafeMind cyber defence AI lab." September 1, 2026.
<https://itbrief.asia/story/crowdstrike-launches-safemind-cyber-defence-ai-lab>
9. Seeking Alpha. "CrowdStrike expands AI security suite with Falcon Guardian, SafeMind." September 1, 2026.
<https://seekingalpha.com/news/4639398-crowdstrike-expands-ai-security-suite-with-falcon-guardian-safemind>

Author: Vatsal Sharma | Peer Reviewed by: Pinki Sharma, Devesh Taneja | FieldCISO Advisory | vatsal@fieldciso.com | fieldciso.com

This report summarises and analyses third-party vendor announcements for informational purposes. Claims attributed to CrowdStrike, EY, Google Cloud, NVIDIA, CoreWeave and Optiv reflect those companies' own press releases and have not been independently verified by FieldCISO Advisory. All SafeMind performance figures, including detection rate, remediation speed and cost savings, are CrowdStrike's own evaluations published without methodology or named comparators, and should not be read as independently established results. Analysis and opinion in the "What Holds Up," "The Question Day 1 Raises," and "What Day 1 Set Up" sections represent FieldCISO Advisory's own editorial assessment and are not a hands-on product evaluation; FieldCISO Advisory has not had test access to Falcon Guardian, SafeMind, or any other product named here. Product capabilities, pricing and availability are subject to change. Readers should confirm current details directly with the relevant vendor before making procurement decisions. FieldCISO Advisory is not affiliated with CrowdStrike, EY, Google, NVIDIA, CoreWeave, Optiv or any other company named in this report. All product names and trademarks belong to their respective owners.