PRODUCTBEACON

ProductBeacon — State of Cyber Security Markets 2026, Front 7: The Emerging AI Theater

Yohay Etsion

Managing Director, ProductBeacon

2026-06-21

Disclosure: *This is ProductBeacon's independent, impartial market research, based solely on publicly verifiable signals, published on the open web. No vendor sponsors, no paywalled data, no analyst-firm reuse. It reflects ProductBeacon's independent view and is not investment advice. Full independence and sourcing methodology: Methodology →.*

By Yohay Etsion · Head of Product (Fractional), AXIA · Creator of Product Org OS · Author of Leading the Charge (2023) and Vision to Value (coming 2026)

AXIA is an AI-powered data-security (Insider Risk Management / DLP) company.

7.1 The Playing Ground

The AI Security front is the youngest battleground in this report, and the only one where the category itself is still being assembled in front of the buyer. It is not one market. It is three, bound together by a single demand driver — agentic AI — but sold to different buyers, on different timetables, against different threats. Naming them cleanly is the first job of this front, because the marketing collateral does not.

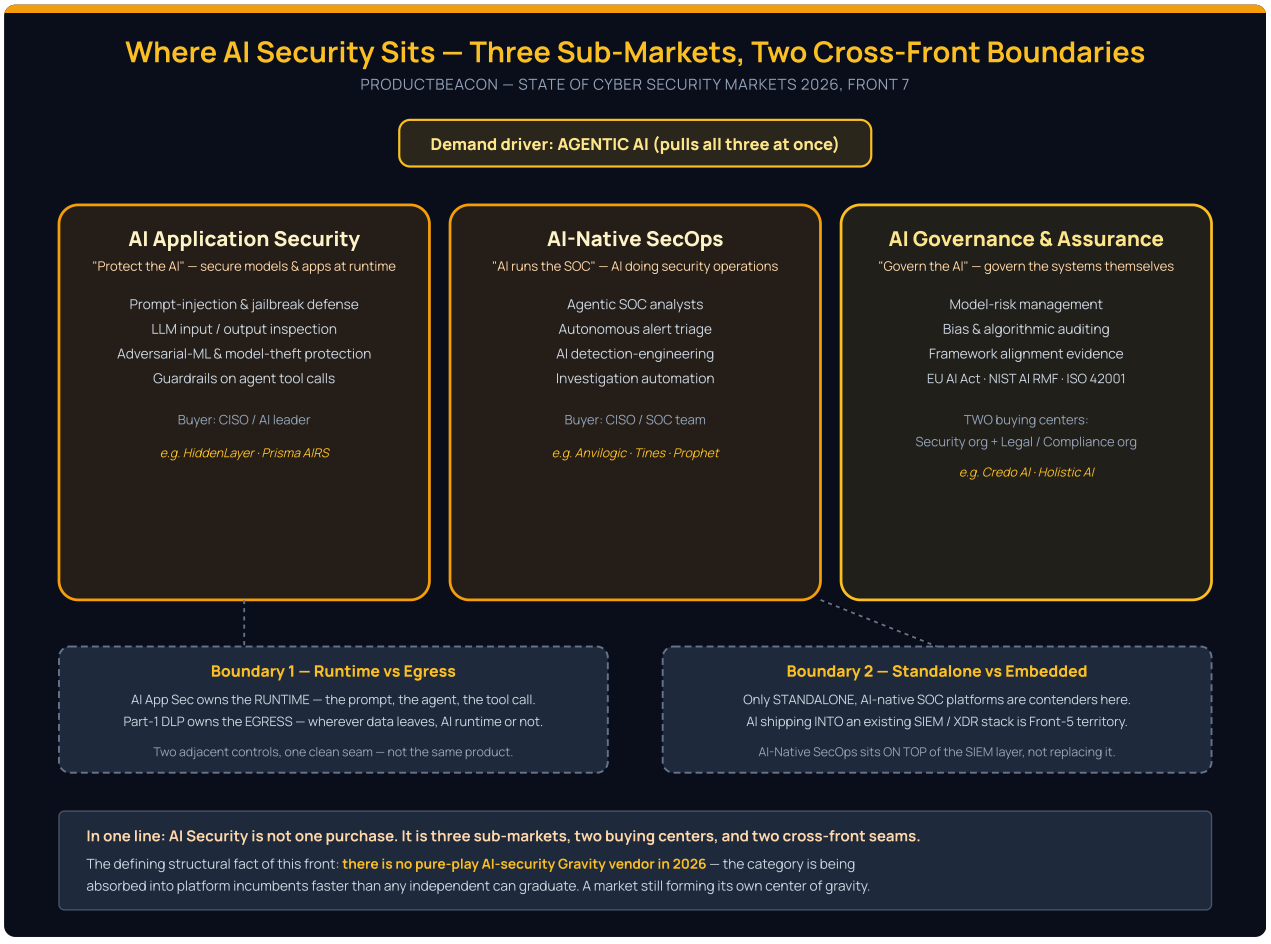

Three sub-markets — runtime defense, AI-run security operations, and AI governance — share an agentic-AI demand driver but split across two buying centers and two cross-front boundaries (DLP egress; the SOC stack).

AI Application Security (AI App Sec) secures AI applications and models at *runtime*: prompt-injection and jailbreak defense, LLM input/output inspection, adversarial-ML and model-theft protection, and — increasingly the center of gravity — guardrails on agent tool calls. Palo Alto's Prisma AIRS, for example, markets itself as blocking 30-plus prompt-injection and jailbreak techniques and discovering AI agents

across cloud, SaaS, and endpoints with per-agent identity and audit trails.[1] **AI-Native SecOps** is the inverse: AI *doing* security operations — agentic SOC analysts, autonomous alert triage, AI detection-engineering. Microsoft's phishing-triage agent claims to surface 6.5x more malicious alerts than analysts working alone, and its Security Analyst Agent runs multi-step investigations across Defender and Sentinel telemetry.[2] **AI Governance & Assurance** governs AI systems themselves: model-risk management, bias and algorithmic auditing, and alignment to frameworks like the EU AI Act, NIST AI RMF, and ISO 42001.

Two cross-front boundaries are decisive, and both are drawn in this report's shared taxonomy. The first separates AI App Sec from Part 1's DLP front. The boundary: **AI App Sec owns the *runtime* — the prompt, the agent, the tool call — while DLP owns the *egress*, regardless of whether the data leaves through an AI runtime or any other channel.**[3] The second separates AI-Native SecOps from Front 5's SOC platform war. AI-augmented detection that ships *into* an existing SIEM/XDR stack is characterized in Front 5; only *standalone, AI-native* SOC platforms are contenders here. As the taxonomy frames it, AI-Native SecOps sits as an analytic layer *on top of* the SIEM/XDR stack — it does not replace it.[3]

The defining structural feature of this front, and its sharpest finding, is what is *missing*: **as of mid-2026 there is no pure-play AI-security Gravity vendor** — none public, none past the roughly $100M-private mark that the report's vendor-tier rubric uses for Gravity. The category is being absorbed into platform incumbents — CrowdStrike's Falcon AIDR and Continuous Identity for AI Agents, Microsoft's Defender for Cloud AI, Palo Alto's Prisma AIRS — faster than any pure-play can graduate. The clearest evidence is acquisition. Palo Alto bought Protect AI for a reported figure above $500M (the deal value was not officially disclosed), and Cisco acquired Robust Intelligence — both pure-play AI-security startups, folded into platforms before they could reach escape velocity alone.[4] I read this as the market's defining condition for 2026: the incumbents are buying the category faster than it can self-organize. The test is public and falsifiable. If a pure-play crosses into Gravity (an IPO, or a disclosed $100M-plus round at a standalone valuation) within the next four quarters, the absorption thesis weakens. If the next two notable AI-security exits are again acquisitions into platforms, it holds.

A second structural note: **AI Governance & Assurance has two distinct buying centers** — the security org and the legal/compliance org — and they do not buy the same way. The CISO buys runtime controls and SOC tooling; the GC or Chief Compliance Officer buys model auditing and regulatory-alignment evidence. A vendor selling "AI governance" to both is selling two products under one word, which is one reason the analyst sizings for this slice diverge so widely.

Three buyer misconceptions are worth retiring.

- **AI security is one purchase.** It is three sub-markets with different buyers, and a runtime-defense tool does not govern a model or run a SOC.
- **A standalone AI-security pure-play is the safe bet.** The absorption pattern above says the durable control is more likely to arrive bundled into a platform the buyer already owns.
- **"AI governance" means the security team's problem.** For the compliance-driven slice, the buying center sits outside the SOC entirely.

This front shares its layer map with the rest of the report. The runtime/egress and embedded/standalone boundaries follow the report's shared layer map, which sets the control-plane boundary policy and cross-Part vendor rules this front inherits.

7.2 The Terrain

Market sizing — three lenses, deliberately not averaged. Because AI Security is three sub-markets, there is no single number. The three analyst lenses below measure genuinely different things, so read them side by side, not blended.

The first lens is **agentic-AI-security** — the slice closest to AI App Sec plus agent-runtime defense. MarketsandMarkets sizes it at USD 1.65 billion in 2026, growing to USD 13.52 billion by 2032 at a 42.0% CAGR. Threat-detection-and-response is the largest segment at 23.1% of 2026 spend, and North America is 41.9%.[5]

The second lens is much broader: **AI-in-cybersecurity**, which folds in every AI-augmented security tool — most of which this report characterizes inside Fronts 5 and 6 rather than here. The same firm puts 2026 at USD 25.53 billion, reaching USD 50.83 billion by 2031 at a 14.8% CAGR. Grand View Research, on a similar definition, estimates USD 25.35 billion in 2024 reaching USD 93.75 billion by 2030 at a 24.4% CAGR.[6] The gap between MarketsandMarkets' 14.8% and Grand View's 24.4% on nominally the same category is a definitional artifact: it turns on how much embedded-AI tooling each firm counts. That is exactly why averaging them would destroy information.

The third and narrowest lens is **AI governance**. MarketsandMarkets sizes it at USD 0.89 billion in 2024 growing to USD 5.78 billion by 2029 at a 45.3% CAGR. On its own cut, Research and Markets puts 2026 at roughly USD 0.61 billion reaching USD 2.63 billion by 2030 at a 44.3% CAGR.[7]

The honest summary across all three: the narrower and newer the slice, the steeper the headline growth rate. Agentic security and governance both clear 40% CAGR off small bases, while the broad embedded-AI category grows slower off a base ten-to-fifteen times larger.

Buyer and user trends. The buyer for this front is not one person. The runtime-defense and AI-SOC slices sell to the CISO and the SOC; the governance slice splits between security and the legal/compliance org, per the two-buying-centers note in §7.1. The user-side signal worth watching is the AI-SOC analyst. KuppingerCole stood up a dedicated *Emerging AI Security Operations Center* report category in 2026 and named Microsoft an overall leader. The creation of an analyst category is itself a public signal that AI-Native SecOps has crossed from feature to market.[8] Buyers are also consolidating their entry point through platforms they already own. Microsoft's Security Copilot now hosts partner agents (OneTrust, Tanium, BlueVoyant, Aviatrix, Fletch), turning the incumbent platform into the distribution channel for third-party AI-security capability rather than a competitor to it.[9]

Technology trend anchor — agentic AI across all three sub-markets. The single demand driver tying this front together is agentic AI, and it pulls on all three slices at once. AI applications and agents need runtime defense (App Sec); AI is increasingly running the SOC itself (AI-Native SecOps); and the agents being deployed must be governed and audited (Governance). The clearest production signals are converging on agent *identity and authorization* as the control point. CrowdStrike's June 2026 Continuous Identity for AI Agents (built on its SGNL acquisition) grants, denies, and revokes agent access in real time with no standing privileges, and Palo Alto's Prisma AIRS 3.0 assigns per-agent identity with RBAC and audit trails.[10] I read agent identity as the contested center of this front for the next several quarters. The

falsifiable test is whether the platform incumbents report named agent-security or AI-runtime SKUs as disclosed growth drivers in upcoming earnings, rather than burying them as undifferentiated platform line-items.

Regulatory trend (context only). Three frameworks form the regulatory backdrop for the Governance slice: the EU AI Act, the NIST AI Risk Management Framework, and ISO 42001, the AI-management-system standard.[11] These are real and they shape what the governance buyer must eventually demonstrate. But they are *context*, not a market-timing mechanism: this report does not predict buying on the strength of a regulatory effective date or a policy-adoption forecast. The forward-looking claims in this front ground only on observable public signals. Those signals are the agentic-AI-security sizing clearing 40% CAGR off a real revenue base, the platform incumbents' acquisitions of Protect AI and Robust Intelligence, KuppingerCole's creation of an AI-SOC analyst category, and CrowdStrike's and Palo Alto's shipped agent-identity products. If the absorption-into-platforms pattern continues to show up in dated acquisitions and earnings disclosures, the thesis holds; if a pure-play graduates to Gravity on its own, it weakens. The regulation is the weather, not the trade.

Sourcing note. AI governance is the thinnest-sourced slice here. The named-analyst sizings diverge by base and definition more than in any other front in this report. Several firms publish governance numbers an order of magnitude apart (sub-$1B versus low-hundreds-of-millions in 2026), depending on whether tooling, services, or model-audit consulting is counted. The two figures cited above are the most clearly defined and consistently scoped; the wider divergence is flagged rather than resolved, and no governance number carries the weight of a forward claim.

7.3 The Contenders

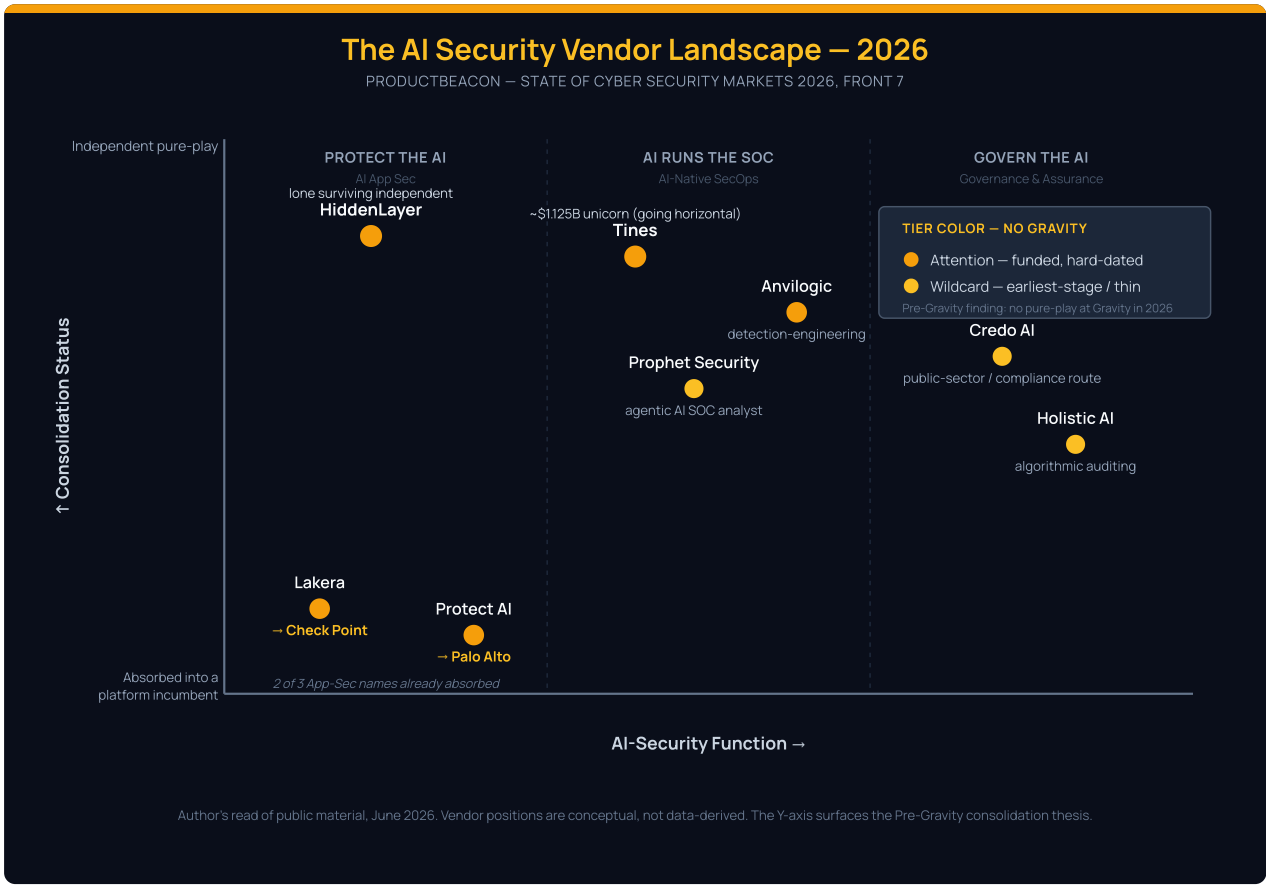

Author's read of public material, June 2026. Vendor positions are conceptual, not data-derived.

The first thing to say about this front is what it does not have. Every other chapter in this report opens its contender section with a Gravity tier — the public companies and the late-stage private platforms whose pricing and packaging set the reference point the rest of the field prices against. The emerging AI-security front has none. No pure-play AI-security vendor is public, and none has crossed the roughly $100-million-private threshold that the Gravity label requires. The largest independents in this set have raised tens of millions, not hundreds. The vendors that would otherwise be the heavyweights have either been absorbed into platform owners or are still living on Series A and Series C rounds. I call this the Pre-Gravity finding, and it is the single most important structural fact about the front: this is a market still forming its own center of gravity, and the gravity it has is borrowed from the larger platforms buying into it.

That borrowing is the second structural fact. The clearest example sits in AI Application Security, where two of the three notable vendors are no longer independent companies — Lakera is now inside Check Point, and Protect AI is now inside Palo Alto Networks. Both moves landed inside a single twelve-month window in 2025, and they leave one surviving independent, HiddenLayer, carrying the pure-play App-Sec story alone. That pattern — a category being bought into existence by the platform players before any independent reaches escape velocity — is the through-line of this section.

The eight vendors below are grouped by sub-market rather than by tier, because the sub-market is where the structure lives. Within each group I note the tier each vendor would carry on the report's standard scale: Attention for the established, funded, hard-dated players; Wildcard for the earliest-stage or thinnest-sourced. But the reader should hold the Pre-Gravity finding throughout. With no Gravity anchor and most vendors raised below sixty million dollars, every card here is descriptive positioning extraction only. No vendor in this section is framed as a winner or a loser. The two acquired App-Sec vendors are described through their public, dated acquisitions — that is factual, not a verdict. Each card is anchored to a single dated vendor-controlled surface and at least one verbatim messaging pillar lifted from that surface on 2026-06-16.

AI Application Security

This is the sub-market where the buy-in story is loudest, so it leads. Of the three vendors that defined model-and-application-layer AI defense, two are now units inside larger platforms and one remains independent. The independent — HiddenLayer — gets the elevated, central treatment here, because it is the only place left to read what a pure-play AI Application Security company sounds like when it speaks for itself.

HiddenLayer

> *"The most comprehensive security platform for AI. Backed by patented technology and industry-leading adversarial AI."* — [hiddenlayer.com, accessed 2026-06-16][12]

HiddenLayer is the lone surviving independent in AI Application Security, and the elevated card of this front. Its stated USP is model-specific runtime and adversarial-ML defense delivered as a platform rather than a point tool. The homepage positions it as "the most comprehensive security platform for AI," built on "patented technology and industry-leading adversarial AI," and frames the coverage as securing "agentic, generative, and predictive AI applications" through AI Discovery, AI Supply Chain Security, and AI Attack Simulation.[13] The product story is the security of the model itself — the weights, the supply chain, and the runtime behavior — rather than the network or the data around it. That is what distinguishes the App-Sec category from the SecOps and governance sub-markets below. The target buyer is the CISO or AI leader inside organizations deploying models into production, with a stated lean toward financial services, technology, and US federal; the surfaces carry no public pricing signal, routing instead through a book-a-demo motion. Architecturally HiddenLayer classifies as a model-layer AI-security platform, vendor-controlled and pure-play.

The capital picture is modest by the standards of every other front in this report, but central to the Pre-Gravity reading. HiddenLayer has raised roughly $56 million in total, anchored by a roughly $50-million Series A in September 2023 led by M12 (Microsoft's venture arm) and Moore Strategic Ventures. It remains independent, with a headcount around 169 as of March 2026, and was founded in Austin in 2022 by Chris Sestito, Tanner Burns, and Jim Ballard.[14] Published-material tier is vendor-controlled plus named-outlet funding coverage; as an early-stage private company there is no public reporter. My read: HiddenLayer is the cleanest articulation of what a stand-alone AI Application Security company is — a

model-defense platform speaking the model's language rather than the network's. The fact that it is the last independent of its original three is itself the descriptive center of gravity for this sub-market, which is precisely why the front has no Gravity tier of its own.

Check Point / Lakera (acquired unit)

> *"AI Is Moving From Language To Action. Are You Ready? Secure AI usage across browsers, SaaS, and copilots. Protect model inputs and outputs at the application edge. Control tool calls, file access, and autonomous runtime behavior."* — [Check Point AI Security, checkpoint.com/ai-security, accessed 2026-06-16][15]

Lakera is no longer an independent contender, and that is the story. Check Point acquired the LLM-security company for a figure reported at roughly $300 million, announced on 2025-09-16 and closed on October 22, 2025, and made it the foundation of a Global Center of Excellence for AI Security.[16] On Check Point's live AI-security surface the Lakera heritage now reads as a three-part frame: securing AI usage at the workforce layer ("across browsers, SaaS, and copilots"), protecting "model inputs and outputs at the application edge," and controlling agent runtime behavior ("tool calls, file access, and autonomous runtime behavior"). That language carries Lakera's original LLM input/output-security DNA up into a network-security platform.[15] Because the acquisition is the descriptive fact, this card extracts positioning from the acquirer's page, which is the live surface; Lakera's pre-acquisition prompt-injection and Gandalf-education heritage survives mainly as a brand reference on that page. The target buyer is Check Point's installed base extending AI controls across an existing security relationship; the architecture classifies as application-edge AI security now embedded inside a larger platform. Published-material tier is named-outlet acquisition coverage plus vendor-controlled. My read: the Lakera acquisition is the clearest single data point behind the App-Sec buy-in finding — an independent LLM-security pure-play folded into a network-security platform as a center-of-excellence anchor, which describes the direction of the sub-market more than any independent's roadmap does.

Palo Alto Networks / Protect AI (acquired unit)

> *"Secure the entire enterprise AI lifecycle and get the visibility, assurance and runtime governance to deploy autonomous AI agents safely."* — [Prisma AIRS, paloaltonetworks.com, accessed 2026-06-16][17]

Protect AI is the second of the three App-Sec vendors to leave independence. Palo Alto Networks announced its intent to acquire the ML/AI model-security company on 2025-04-28 and completed the deal on July 22, 2025 (its fiscal Q4 FY2025), folding Protect AI's heritage into Prisma AIRS — Palo Alto's AI runtime security platform.[18] On the live Prisma AIRS surface the inherited capability reads as a lifecycle frame: "secure the entire enterprise AI lifecycle" across Discover, Assess, and Protect. The platform statement is that "the Prisma AIRS platform secures all AI agents, apps, models and data from development to deployment," with a capability set spanning agent security, AI red teaming, AI runtime security, AI model security, and AI posture management.[17] As with Lakera, the descriptive fact is the

acquisition, so positioning is extracted from the acquirer's page. The target buyer is Palo Alto's enterprise installed base extending into AI-lifecycle security; the architecture classifies as full-lifecycle AI security embedded inside a larger platform. Published-material tier is named-outlet acquisition coverage plus vendor-controlled (PANW, NASDAQ). My read: Protect AI's absorption into Prisma AIRS completes the App-Sec buy-in pattern — two of the three notable model-security independents are now platform modules, leaving HiddenLayer as the descriptive anchor of an independent category that the larger platforms are otherwise building out by acquisition.

AI-Native SecOps

The second sub-market is the security operations center rebuilt around AI — vendors applying detection engineering, triage, and investigation automation to the SOC workflow. The three here range from an established detection-engineering platform to an agentic automation layer that is consciously growing beyond security, to an early-stage agentic SOC analyst. The agentic-SOC thesis itself is examined in Front 5; these cards extract positioning only.

Anvilogic (Attention)

> *"Scale your SOC, Keep your stack. Agentic SecOps platform that onboards, searches, detects, and investigates across every data source — without moving a single byte."* — [anvilogic.com, accessed 2026-06-16][19]

Anvilogic is the detection-engineering play in this sub-market — the vendor applying AI across the SOC lifecycle without forcing data migration. Its stated USP is captured in the homepage promise to "scale your SOC, keep your stack": an "agentic SecOps platform that onboards, searches, detects, and investigates across every data source" while leaving data in place — "without moving a single byte" — and running "standalone on your security data lakes or alongside your SIEM."[20] The product applies AI throughout the detection lifecycle: AI-enhanced detection and investigation, automated detection tuning with ML recommendations, and agentic triage via an agentic workbench. It supports Snowflake, Databricks, and Azure data lakes directly as the underlying store. The target buyer is the SOC team, detection engineer, or security architect modernizing detection without ripping out the existing SIEM. No public pricing signal surfaces; the motion routes through a demo. Architecturally Anvilogic classifies as a data-lake-native agentic SecOps platform, founded in 2019 and independent. The capital picture sits squarely in the Pre-Gravity band: Anvilogic raised a $45-million Series C led by Evolution Equity Partners, bringing total funding to roughly $85 million.[21] Published-material tier is vendor-controlled plus named-outlet funding coverage. My read: Anvilogic's distinctive descriptive claim is the no-data-movement architecture. It positions as the SOC AI layer that sits on top of the customer's existing data estate rather than demanding a new one — a different bet than the rip-and-replace SOC platforms, and a clean fit with the data-lake-centric direction of modern security operations.

Tines (Attention)

> *"The intelligent workflow platform. Securely scale AI and automation. Integrate agents, teams, and tools with speed and control."* — [tines.com, accessed 2026-06-16][22]

Tines is the automation-layer incumbent going horizontal, and that repositioning is the central descriptive fact of this card. Tines came up as a security-automation platform — a SOAR alternative built for analysts rather than coders. Its current surfaces consciously frame it instead as "the intelligent workflow platform," positioned as a foundational layer "no matter how many tools you change," explicitly spanning "Security, IT Ops, Infrastructure, Engineering, and Product teams" rather than security alone.[23] The product organizes around three workflow types — human-led, deterministic, and agentic — with the agentic layer carrying the AI message ("flexibility, ambiguity, safe actions"), and the pitch leans hard on time-to-value and non-developer accessibility. The target buyer has broadened accordingly: the surfaces describe multi-functional enterprise teams, not just the SOC. That is why Tines reads as the automation layer widening beyond its security origin rather than a pure AI-SOC analyst. Architecturally Tines classifies as a horizontal agentic-workflow-automation platform with security heritage; it is based in Dublin. The capital picture is the largest independent valuation in this front: Tines raised a $125-million Series C that lifted its valuation to roughly $1.125 billion — unicorn status — on roughly $272 million raised in total.[24] Published-material tier is vendor-controlled plus named-outlet funding coverage. My read: Tines is best described not as an AI-SOC vendor but as the automation layer that started in security and is deliberately becoming horizontal. It is the most-capitalized name in this section, but its center of gravity is moving toward IT and ops workflow generally, which makes "AI security vendor" an increasingly partial label for it.

Prophet Security (Wildcard)

> *"AI SOC Agents that Accelerate Detection and Response. Investigate every alert in minutes. Hunt threats before they surface. Close detection gaps before attackers find them."* — [prophetsecurity.ai, accessed 2026-06-16][25]

Prophet Security is the agentic AI SOC analyst of the trio — the earliest-stage of the three SecOps names and a Wildcard on the report's scale. Descriptively, its surfaces lead with "AI SOC Agents that Accelerate Detection and Response." They frame the product as agents that "investigate every alert in minutes, hunt threats before they surface, and close detection gaps before attackers find them," working "every alert end-to-end."[26] The described capability is first-pass alert triage and investigation handled by AI agents alongside the human SOC. That places Prophet directly in the agentic-SOC thesis that Front 5 examines in depth, where Dropzone is the Wildcard analog — a deliberate cross-front link rather than a separate finding. The described buyer is the SOC team seeking to absorb alert volume through autonomous investigation. Architecturally Prophet presents as an agentic AI SOC analyst, cloud-delivered, founded by Kamal Shah and Arun Sreeshankar. On capital, Prophet has raised roughly $41 million across a Bain Capital Ventures seed, a $30-million Accel-led Series A, and a strategic round joined by American Express Ventures and Citi Ventures dated 2026-02-25. That is the freshest hard-dated and

strongest-sourced funding signal of the SecOps trio.[27] Published-material tier is vendor-controlled plus named-outlet funding coverage. This is a Wildcard-tier card under the front-wide descriptive discipline: the treatment is positioning extraction only, and Prophet is not carried into any directional pattern claim. My read: Prophet is the purest expression of the "AI SOC analyst" descriptive category in this set — an agentic-triage layer with strategic-investor validation from two financial incumbents. Its cleanest read is as the independent counterpart to the agentic-SOC dynamic that Front 5 treats as a thesis in its own right.

AI Governance & Assurance

The third sub-market is the assurance layer — vendors that do not defend the model at runtime but govern it: cataloguing the AI an organization runs, testing it for bias and safety, and mapping it to regulatory frameworks. Both names here carry the Wildcard tier on the report's scale. A note on framing discipline applies to this pair specifically: any forward-looking observation about governance vendors is grounded only on a vendor, analyst, or funding proxy, never on a regulatory-effective-date or policy-adoption forecast. The regulation is the customers' problem, and treating its timeline as a vendor's tailwind would be exactly the kind of unverifiable claim the report's method excludes.

Credo AI (Wildcard)

> *"Govern AI Everywhere. One platform to discover, assess, and govern every AI agent, model, and application — continuously and in context."* — [credo.ai, accessed 2026-06-16][28]

Credo AI is the governance-platform Wildcard with the freshest hard-dated 2026 signal of the pair. Descriptively, its surfaces lead with "Govern AI Everywhere" and frame the product as "one platform to discover, assess, and govern every AI agent, model, and application — continuously and in context." A continuous risk-assessment pillar covers "bias, security, privacy, and compliance," with explicit mapping to the EU AI Act, NIST AI RMF, ISO 42001, SOC 2, GDPR, and HITRUST.[29] The described buyer skews toward Fortune 500 and regulated-industry enterprises scaling AI beyond pilots. Architecturally Credo AI presents as an AI governance and assurance platform; it is based in Palo Alto and was founded in 2020. The verifiable proxies are unusually dense for a Wildcard: a January 2026 Carahsoft partnership opening public-sector procurement, a #6 placement on Fast Company's Most Innovative Companies 2026 (Applied AI), and a Gartner "Cool Vendor in AI Cybersecurity Governance" recognition. Total funding sits at roughly $41.3 million, anchored by a $21 million Series B in July 2024.[30] Published-material tier is vendor-controlled plus named-outlet and analyst recognition. This is a Wildcard-tier card under the descriptive discipline: positioning extraction only, no directional verdict. My read: Credo AI's descriptive identity is the model-and-agent governance catalog mapped to named frameworks. Its messaging is squarely about governing AI systems for bias, safety, and regulatory alignment, not about monitoring people or data flows, and the 2026 public-sector and analyst proxies are what make it the better-evidenced of the governance pair.

Holistic AI (Wildcard)

> *"The end-to-end AI governance platform trusted by global enterprises running AI at scale."* — [holisticai.com, accessed 2026-06-16][31]

Holistic AI is the algorithmic-auditing Wildcard — the governance vendor whose heritage is bias testing and algorithm assessment, framed today around end-to-end AI governance. Descriptively, its surfaces carry the verbatim positioning line "the end-to-end AI governance platform trusted by global enterprises running AI at scale," organized around the pillars "Identify," "Protect," and "Enforce." The coverage spans AI model discovery and inventory, bias and safety testing (a stated 40-plus specialized tests), shadow-AI discovery across cloud and code, and regulatory mapping to the EU AI Act, NIST AI RMF, and ISO 42001.[32] The EU AI Act framing is the deepest of the pair, consistent with the company's UK origin (London, founded 2020). The described buyer is the global enterprise running AI at scale that needs auditing and compliance evidence. Architecturally Holistic AI presents as an algorithmic-auditing and AI governance platform. On capital, the honest currency note is that the widely circulated $200-million figure does not belong to this London AI-governance vendor. It is the funding of a similarly-named but distinct company, and is therefore not stated here as Holistic AI's funding. The London governance vendor's own disclosed funding is in the low tens of millions, and no fresh post-baseline round was confirmed in the bounded scan. The position rests instead on the dated vendor surface (with blog activity into May 2026) and named-outlet and buyer-guide corroboration.[33] Published-material tier is vendor-controlled plus named-outlet corroboration. This is a Wildcard-tier card under the descriptive discipline: positioning extraction only. My read: Holistic AI's descriptive identity is auditing the model — bias, safety, and EU-AI-Act conformity of the AI system itself — with no messaging touching employee monitoring, insider risk, or data-loss prevention. Like Credo AI, it governs models, not people, and its sharpest differentiator is the depth of its algorithmic-audit and EU-AI-Act feature framing.

7.4 Their Plays

The AI-security front is the youngest theater in this report, and its strategic moves read differently from the mature fronts. There are no decade-old incumbents defending installed bases here; there are venture-funded specialists racing a clock, and platform incumbents deciding whether to buy the specialists or build past them. The four plays below are not four answers to one question the way the edge front's were. They are four distinct motions happening at once: the platforms absorbing the independent "Security-for-AI" layer; the cross-front incumbents embedding AI-security into tools the buyer already owns; AI agents being pointed at the SOC's own first-pass work; and a governance category testing whether it is a business or a feature. Each play below is stated as a public observation, its named participants, and a conditional outcome that resolves against a signal an outsider can watch.

Play 1: The Platform-Absorbs-AI-App-Sec Consolidation Play

- **Observation.** The independent "Security-for-AI" layer — the pure-play vendors building runtime protection, red-teaming, and guardrails for AI applications — is being absorbed into platform incumbents faster than it is going public. Check Point acquired Lakera, the AI-application-security

and red-teaming specialist, for a reported ~$300 million, announced September 16, 2025.[34] Palo Alto Networks announced its intent to acquire Protect AI on April 28, 2025, folding it into the Prisma AIRS AI-security runtime offering.[35] Cisco acquired Robust Intelligence in 2024, building it into the company's AI-security narrative.[36] Three of the most visible independent AI-application-security vendors have now exited into platforms inside roughly eighteen months — before any reached a public listing or a disclosed nine-figure private valuation. HiddenLayer is the most prominent independent not yet absorbed.[37]

- **Participants.** Check Point (acquirer of Lakera), Palo Alto Networks (acquirer of Protect AI → Prisma AIRS), and Cisco (acquirer of Robust Intelligence) as the absorbing platforms; Lakera, Protect AI, and Robust Intelligence as the absorbed specialists; HiddenLayer as the lone visible independent still standing.
- **Conditional outcome.** *Two signals would confirm the pattern over the next two-to-four quarters: a fourth named AI-application-security independent acquired by a platform incumbent in named-outlet coverage, or HiddenLayer itself acquired rather than raising an independent growth round. Either one and we expect the absorb-the-AI-App-Sec-layer pattern to validate as the default exit for this category, with the independent layer continuing to thin. The inverse signal points the other way: if HiddenLayer instead closes a disclosed independent growth round at an up valuation and no further pure-play absorption surfaces, the independent Security-for-AI layer holds as a survivable standalone category rather than a way station to a platform acquisition.*
- **Evidence.** Check Point, "Check Point to Acquire Lakera," September 16, 2025 (checkpoint.com, accessed 2026-06-16); Palo Alto Networks, "Palo Alto Networks to Acquire Protect AI," April 28, 2025 (paloaltonetworks.com, accessed 2026-06-16); Cisco, Robust Intelligence acquisition, 2024 (cisco.com, accessed 2026-06-16).

Play 2: The Incumbent-Embeds-AI-Security Play

- **Observation.** The largest cross-front platform incumbents are not (mostly) buying AI-security independents at scale; they are embedding AI-security capabilities into the detection-and-response platforms enterprises already run. Microsoft ships Defender for Cloud AI security-posture management as part of its cloud-security suite; CrowdStrike ships Falcon for AI; Palo Alto Networks ships XSIAM AI alongside its Prisma AIRS runtime offering.[38] All three are characterized as primary contenders in Front 5 (Detection & Response), where their EDR, SIEM, and SOC platforms are graded. Here they appear only as AI-security moves, named by the AI-security product line each is attaching to its existing platform.[39] The competitive vector is the same one the edge front showed: the AI-security line item competes against a capability the buyer can switch on inside a platform they already license, rather than procure from an independent vendor.
- **Participants.** Microsoft (Defender for Cloud AI), CrowdStrike (Falcon for AI), and Palo Alto Networks (XSIAM AI / Prisma AIRS) as the embedding incumbents — each holding its primary contender slot in Front 5; the independent AI-security specialists of §7.3 as the standalone vendors whose offerings the embedded capability compresses.
- **Conditional outcome.** *If any of these three names embedded AI-security as a disclosed product-adoption or attach metric — in an earnings call or a named-outlet enterprise win over the next two-to-four quarters — we expect the embed-rather-than-acquire motion to validate as real pressure on the independent AI-security layer. If the embedded AI-security products stay absent from these*

incumbents' own earnings narratives, and the independents keep closing disclosed funding rounds and named enterprise logos, the embedded capability reads as a platform checkbox rather than a displacement of the specialist layer.

- **Evidence.** Microsoft, Defender for Cloud AI security posture management documentation, 2026 (learn.microsoft.com, accessed 2026-06-16); CrowdStrike, Falcon for AI product page, 2026 (crowdstrike.com, accessed 2026-06-16); Palo Alto Networks, Prisma AIRS / XSIAM AI product pages, 2026 (paloaltonetworks.com, accessed 2026-06-16).

Play 3: The Agentic-SOC Play

- **Observation.** A cohort of venture-funded vendors is pointing AI agents at the security operations center's own first-pass work — alert triage, investigation, and enrichment that a Tier-1 analyst would otherwise do by hand. Prophet Security raised roughly $41 million across seed and Series A and disclosed strategic participation from American Express Ventures and Citi Ventures in a round dated February 2026, positioning an "agentic SOC" that automates first-pass investigation.[40] Anvilogic raised a $45 million Series C led by Evolution Equity Partners, building a detection-engineering and AI-triage layer that sits across existing SIEM and data-lake tooling.[41] The thesis is that the SOC's most repetitive labor — the first read of every alert — is the work AI agents can do soonest, and that an agent doing it well is a wedge into the SOC budget. This thread is examined as a primary matter in Front 5, where Dropzone AI is the detection-and-response front's Wildcard analog of the same agentic-SOC bet; the cross-front link is noted here rather than re-graded.[42]
- **Participants.** Prophet Security and Anvilogic as the AI-security-front agentic-SOC challengers; the detection-and-response incumbents and the F5 Wildcard (Dropzone AI) whose first-pass-analyst territory the agentic-SOC vendors contest, characterized primarily in Front 5.
- **Conditional outcome.** *If an agentic-SOC vendor in this cohort discloses a named enterprise deployment or a follow-on growth round tied to production SOC adoption over the next two-to-four quarters, we expect agentic first-pass SOC work to validate as a fundable wedge into the operations budget rather than a pilot-stage curiosity. If these vendors' disclosed traction stays in pilots and design-partner counts, with no production-deployment or growth-round signal, the agentic-SOC pitch remains an unproven layer that the detection-and-response platforms absorb as a feature rather than cede as a category.*
- **Evidence.** Prophet Security, Series A / strategic investment coverage, February 2026 (prophetsecurity.ai, accessed 2026-06-16); Anvilogic, "Anvilogic Raises $45M Series C," Evolution Equity Partners (anvilogic.com, accessed 2026-06-16).

Play 4: The AI-Governance-as-a-Commercial-Category Play

- **Observation.** A set of AI-governance specialists is testing whether AI governance — model inventory, risk assessment, policy mapping, and audit evidence — becomes a standalone commercial software category or settles in as a compliance bolt-on inside larger platforms. Credo AI, the AI-governance specialist, has raised roughly $41.3 million in total, anchored by a $21 million Series B in July 2024, and built its go-to-market partly through public-sector distribution, including a Carahsoft partnership that puts its governance platform onto government procurement vehicles.[43] Holistic AI occupies the same governance-and-assurance wedge, positioning AI audit and risk tooling for enterprise and regulated buyers.[44] The thesis is that AI governance is large and recurring enough to

support independent vendors with their own buyers, rather than being subsumed as a feature of the security platforms that already sit in the enterprise.

- **Participants.** Credo AI (governance platform, public-sector distribution via Carahsoft) and Holistic AI (AI audit and risk assurance) as the standalone-governance challengers; the platform incumbents and GRC suites whose compliance modules could absorb governance as a feature.
- **Conditional outcome.** *Suppose a standalone AI-governance vendor in this cohort converts a distribution signal into disclosed traction over the next two-to-four quarters — a named-outlet enterprise or public-sector win through the Carahsoft vehicle, a follow-on funding round, or recognition in a Gartner or analyst category for AI governance. We then expect AI governance to validate as its own commercial software category with independent buyers. If instead these vendors show no fresh funding, no named enterprise win, and no analyst-category recognition while platform GRC and security suites add governance modules, AI governance reads as a compliance feature absorbed into existing platforms rather than a standalone category.*
- **Evidence.** Credo AI, Carahsoft partnership announcement (credo.ai / carahsoft.com, accessed 2026-06-16); Credo AI funding history (crunchbase.com, accessed 2026-06-16); Holistic AI company and product pages (holisticai.com, accessed 2026-06-16).

7.5 War Chests & Casualties

The capital picture for the AI-security front is unlike any other front in this report, and the difference is the whole point. The snapshot below mixes the independent specialists with the two largest exits, which were not IPOs but acquisitions into platforms.

Vendor	Most Recent Round	Valuation (if public)	Strategic Investor	Distress Signal
HiddenLayer (private)	$50M Series A, Sep 2023; ~$56M total[45]	Not disclosed	M12 (Microsoft's venture fund)	(empty — no public distress event)
Tines (private)	$125M Series C, ~$1.125B valuation[46]	~$1.125B (private)	Goldman Sachs Alternatives (lead)	(empty — no public distress event)
Anvilogic (private)	$45M Series C[47]	Not disclosed	Evolution Equity Partners (lead)	(empty — no public distress event)
Prophet Security (private)	~$41M total; strategic round Feb 2026[48]	Not disclosed	American Express Ventures, Citi Ventures	(empty — no public distress event)
Credo AI (private)	~$41.3M total; $21M Series B Jul 2024[49]	Not disclosed	Sands Capital, Decibel	(empty — no public distress event)
Lakera → Check Point (exit)	Acquired ~$300M, Sep 16 2025[50]	Consolidation exit	Check Point	(empty — consolidation exit, not distress)
Protect AI → Palo Alto (exit)	Acquisition intent, Apr 28 2025[51]	Consolidation exit	Palo Alto Networks	(empty — consolidation exit, not distress)

The defining feature of this front is the absence of a war chest at scale: no pure-play AI-security vendor has crossed a public listing or a disclosed $100M-plus private valuation as a standalone company, and the largest "exits" are acquisitions into platform incumbents rather than independent milestones. Lakera went to Check Point for a reported ~$300 million and Protect AI to Palo Alto Networks under its April 2025 intent — both public, dated, positive-or-neutral consolidation events, not distress. The independent capital here is early-to-mid venture. HiddenLayer's $50 million 2023 Series A with Microsoft's M12 on the cap table is the most prominent surviving independent. Anvilogic's $45 million Series C and Prophet Security's ~$41 million with American Express and Citi strategic backing fund the agentic-SOC bet. Credo AI's ~$41.3 million in total (anchored by a $21 million July-2024 Series B) funds the governance bet. Tines, the workflow-automation and agentic-security-operations vendor, is the rare unicorn-mark in the adjacency at a ~$1.125 billion valuation off its $125 million Series C, though its center of gravity is automation rather than AI-application security proper. No vendor in this front carries a citable public distress event; the front's defining motion is consolidation into platforms, not casualty.

7.6 Winning & Losing

This front is not won or lost the way the mature theaters are. There is no installed base to defend, no decade-old incumbent to dislodge, no public-market scoreboard to read. As Part 1 established, AI Security has *zero* pure-play Gravity in 2026 — not one independent vendor has reached the public market or the roughly $100M-private mark the report's tier rubric uses for Gravity. That single structural fact governs everything I can responsibly say here. With no Gravity vendor to anchor a verdict, and with most contenders sub-$60M-raised privates whose individual survival I cannot call from public material, **I am deliberately not labeling any individual early-stage vendor a winner or a casualty**. Every claim below is about the *category's motion* — how the sub-markets are forming, where the spend is consolidating, and which structural patterns the public evidence supports — not a directional bet on any one company's fate.

That discipline is not a hedge; it is the honest read of a pre-Gravity market. When the high ground itself is still being assembled in front of the buyer, the contest that matters is not vendor-versus-vendor but *category-versus-absorption*. The question is whether any of these three sub-markets can organize into a durable standalone business before a platform incumbent buys the capability and folds it into a box the buyer already owns. Three patterns explain the motion, each stated as a public observation, a labeled read, and a falsifiable conditional that resolves against a signal an outsider can watch. The first is the spine: the window in which a pure-play could still graduate to Gravity is closing because the platforms are buying the category faster than it can self-organize. The second is a bifurcation: AI governance is splitting across two buying centers, and the vendor that bridges them defines the category. The third is a layering claim: AI-native security operations is forming as a layer *on top of* the SOC stack, not a replacement for it.

Pattern Claim 1 — The Pre-Gravity Window Thesis

Observation. - As of mid-2026 there is no pure-play AI-security vendor at Gravity — none public, none past the roughly $100M-private threshold the report uses — and the most visible independents are exiting into platform incumbents before they can graduate. - Check Point acquired AI-application-security and red-teaming specialist Lakera for a reported ~$300M, announced September 16, 2025, making Lakera the foundation of Check Point's Center of Excellence for AI Security [52]. - Palo Alto Networks announced its intent to acquire Protect AI on April 28, 2025, folding it into the Prisma AIRS AI-security runtime line [53]. - Cisco acquired Robust Intelligence in 2024 and built it into the company's AI-security narrative [54]. - Across the platform field the same motion shows up as embedding rather than acquisition: Microsoft's Defender for Cloud AI security posture, CrowdStrike's Falcon protections for AI, and Palo Alto's Prisma AIRS all ship AI-security capability inside platforms the enterprise already buys [55]. - Three of the most prominent independent AI-application-security vendors have exited into platforms inside roughly eighteen months, before any reached a public listing or a disclosed nine-figure standalone valuation. HiddenLayer is the most prominent independent that has not been absorbed [56].

My read. - I read this as a pre-Gravity *window* that is closing — and closing on the absorption side, not the graduation side. - The structural reason no pure-play has reached Gravity is not that the category lacks demand. It is that the platform incumbents can buy a runtime-defense or model-security capability and cross-sell it into an installed base faster than a venture-funded specialist can compound to escape velocity alone. - When the durable control most likely arrives bundled into a platform the buyer already owns, the pure-play's path to Gravity narrows to a footrace it is structurally disadvantaged to win. - I want to be precise about what I am *not* saying: I am not calling any individual independent a casualty, and I am explicitly not reading the two acquisitions above as failures. They are consolidation exits — the category's most-funded specialists being bought *because* they were good, not because they were dying. - The claim is about the category's shape: the window in which an AI-security pure-play could still graduate to standalone Gravity is the thing under pressure, and the public M&A record is the evidence that it is narrowing.

Conditional prediction. - Watch the next two-to-four quarters for a fourth named AI-application-security or AI-SecOps independent absorbed into a platform incumbent in named-outlet M&A coverage. If the next notable AI-security exits are again acquisitions of that kind, the pre-Gravity window stays open only as a way station, and absorption is confirmed as the category's default exit. - The window's *close* would be marked instead by the inverse signal: the first pure-play AI-security acquisition above ~$200M as a standalone valuation, or the first AI-security pure-play to cross to a public listing or a disclosed $100M-plus standalone private round. Either of those, in named-outlet M&A coverage, SEC filings, or a dated funding announcement, would mark a pure-play graduating to Gravity rather than being absorbed below it. - I am grounding this on the public M&A and funding record, not on any regulatory date.

Sources. [52] [53] [54] [55] [56]

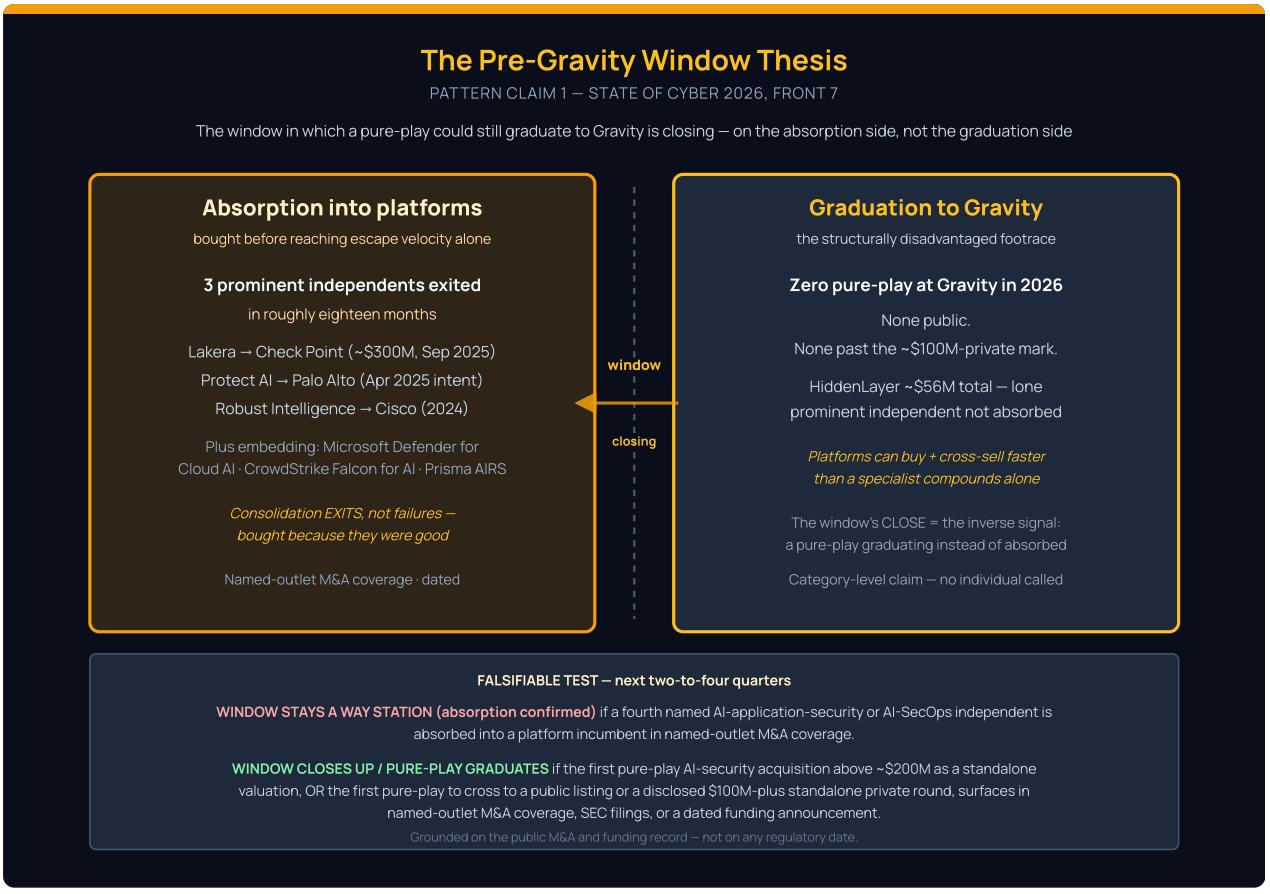
The Pre-Gravity Window Thesis
PATTERN CLAIM 1 — STATE OF CYBER 2026, FRONT 7
The window in which a pure-play could still graduate to Gravity is closing — on the absorption side, not the graduation side
Absorption into platforms
bought before reaching escape velocity alone
3 prominent independents exited
in roughly eighteen months
Lakera → Check Point (~$300M, Sep 2025)
Protect AI → Palo Alto (Apr 2025 intent)
Robust Intelligence → Cisco (2024)
Plus embedding: Microsoft Defender for Cloud AI · CrowdStrike Falcon for AI · Prisma AIRS
Consolidation EXITS, not failures — bought because they were good
Named-outlet M&A coverage · dated
window
closing
Graduation to Gravity
the structurally disadvantaged footrace
Zero pure-play at Gravity in 2026
None public.
None past the ~$100M-private mark.
HiddenLayer ~$56M total — lone prominent independent not absorbed
Platforms can buy + cross-sell faster than a specialist compounds alone
The window's CLOSE = the inverse signal: a pure-play graduating instead of absorbed
Category-level claim — no individual called
FALSIFIABLE TEST — next two-to-four quarters
WINDOW STAYS A WAY STATION (absorption confirmed) if a fourth named AI-application-security or AI-SecOps independent is absorbed into a platform incumbent in named-outlet M&A coverage.
WINDOW CLOSES UP / PURE-PLAY GRADUATES if the first pure-play AI-security acquisition above ~$200M as a standalone valuation, OR the first pure-play to cross to a public listing or a disclosed $100M-plus standalone private round, surfaces in named-outlet M&A coverage, SEC filings, or a dated funding announcement.
Grounded on the public M&A and funding record — not on any regulatory date.

Pattern Claim 2 — The Governance-vs-Compliance Bifurcation Thesis

Observation. - AI Governance & Assurance is being sold to two different buyers under one word. The same "AI governance" label is pitched to the security organization (model-risk, runtime inspection, the CISO's stack) and to the legal/compliance organization (algorithmic auditing, regulatory-alignment evidence, the GC's or Chief Compliance Officer's stack), and the public go-to-market signals split along that line. - Credo AI took an explicitly public-sector, compliance-led route. On January 7, 2026 it named Carahsoft its Master Government Aggregator, putting its governance platform onto NASA SEWP V, ITES-SW2, NASPO ValuePoint and other government procurement vehicles, positioned around NIST AI RMF and OSTP alignment for agencies that must audit AI systems and maintain regulatory compliance [57]. - Credo was also recognized as a Gartner Cool Vendor in AI cybersecurity governance and named to Fast Company's Most Innovative Companies of 2026 [58]. - Holistic AI, by contrast, leads with the deeper assurance-and-auditing feature framing: algorithmic auditing, bias testing, and end-to-end governance for enterprises running AI at scale [59]. - Underneath both, Gartner's AI TRiSM (AI Trust, Risk and Security Management) market guide models the space as four layers — AI governance, AI runtime inspection and enforcement, information governance, and the infrastructure stack — with the top two consolidating into a distinct new market segment [60].

My read. - I read AI Governance & Assurance as a genuinely bifurcated category whose winner will be whoever bridges the two buying centers rather than picking one. - The security-org buyer and the compliance-org buyer evaluate, budget, and renew differently: the CISO wants runtime controls that plug into the SOC; the GC wants defensible audit evidence and regulatory-alignment artifacts, and sits outside the SOC entirely. A vendor that sells only to one center is selling half a category — strong in its lane, structurally capped at the other. - The public signals are consistent with each contender currently anchoring one side: a public-sector, procurement-and-compliance motion on one hand; a deeper algorithmic-assurance feature motion on the other. - The category-level claim I am willing to stand behind is that the bridge between the two centers is the contested high ground. Gartner formalizing AI TRiSM as a four-layer market, with the governance and runtime-inspection layers consolidating, is the analyst evidence that the bridge is real and being built. - To be precise about scope: the two buying centers here both buy *governance of AI systems themselves* — model-risk, bias, safety, and regulatory-conformity evidence for the models and agents an organization runs — not monitoring of the people who use them. That keeps this category cleanly distinct from the data-governance and insider-risk platforms that secure employees and data flows. The AI-governance buyer is evaluating the model, not the workforce.

Conditional prediction. - Suppose a single AI-governance vendor surfaces both motions at once, in named-outlet coverage or its own disclosed go-to-market: a security-org motion (runtime inspection / SOC integration) *and* a compliance-org motion (audit evidence, procurement availability such as Credo's Carahsoft public-sector route). If that happens, and Gartner's AI TRiSM coverage continues to consolidate the governance and runtime-inspection layers into one named segment, the bifurcation resolves toward a bridging winner and the category coheres. - If instead the public signals stay split — governance vendors visibly anchored to one buying center, analyst coverage keeping governance and runtime inspection as separate layers — the category stays

bifurcated and “AI governance” remains two products under one word. - I am grounding this on vendor go-to-market signals, named partnerships, and Gartner’s analyst category formation, not on when any AI regulation takes effect.

Sources. [57] [58] [59] [60]

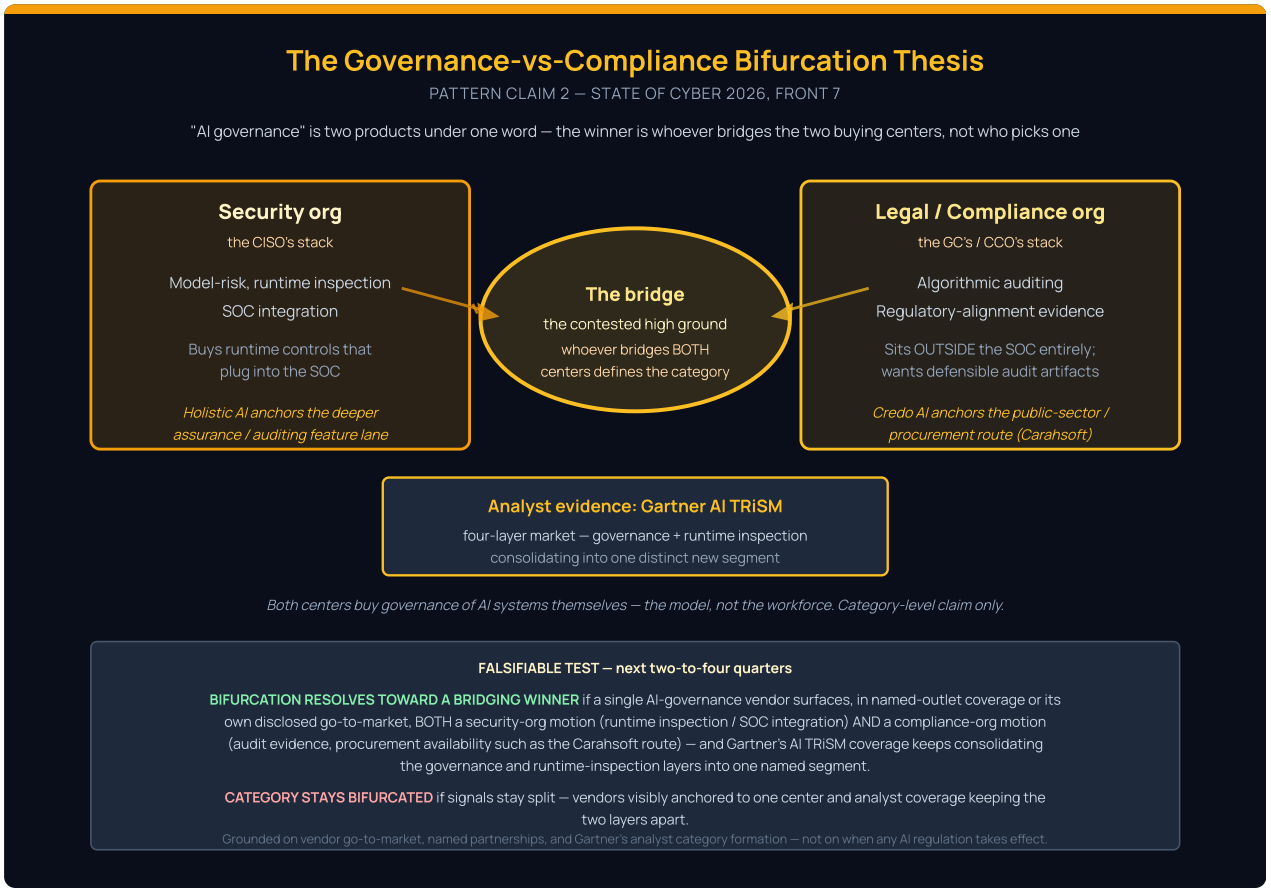

Pattern Claim 3 — The AI-SOC Layering Thesis

Observation. - The AI-Native SecOps vendors are funding and shipping as a layer that runs *on top of* the existing security-operations stack, not as a rip-and-replace of the SIEM. - Prophet Security, a pioneer of the agentic AI SOC, raised strategic investment from Amex Ventures and Citi Ventures announced February 25, 2026, bringing it to roughly $41M across three rounds. Its own framing is that its agentic platform "works alongside security analysts and automates manual, time-sensitive workflows" across investigation, threat hunting, detection engineering, and incident response [61]. - Anvilogic closed a $45M Series C led by Evolution Equity Partners on April 17, 2024 (bringing total funding to ~$85M). The stated purpose was to expand generative-AI features *across the SOC lifecycle* and to enable security-data-lake adoption — a detection-engineering layer over multiple data platforms rather than a replacement SIEM [62]. - Tines, the security-automation platform, raised a $125M Series C to a ~$1.125B valuation and is broadening from a security-automation pure-play into a horizontal agentic-workflow platform spanning IT and operations [63]. - Across these, the common architectural posture is an analytic or agentic layer that consumes the telemetry the SIEM/XDR layer already collects. That is the same boundary the report's taxonomy draws, where AI-native SecOps sits on top of the SIEM layer rather than displacing it [64].

My read. - I read AI-Native SecOps as forming as a *layer-dependent* category — an agentic analytic layer that needs the SIEM/XDR stack underneath it, not one that retires it. - The economics point this way. The data-collection, retention, and normalization work the SIEM does is expensive and entrenched. The AI-SOC vendors' own pitches (first-pass triage, detection engineering, investigation acceleration) are explicitly about *acting on* that data faster, which presupposes the underlying stack keeps existing. - The funding patterns reinforce it: capital is flowing to vendors that augment the SOC lifecycle and to an automation layer going horizontal, not to anyone publicly claiming to replace the SIEM. - The cross-front seam matters here, and I am keeping it clean. AI-augmented detection that ships *into* an incumbent SIEM/XDR is Front 5's territory; the contenders in this front are the *standalone, AI-native* layers, and even they sit on top of the stack. - This connects directly to Front 5's agentic-SOC analysis — the autonomous-triage motion characterized there (including vendors like Dropzone) is the same on-top-of-the-stack pattern viewed from the SOC-platform side rather than the AI-native-pure-play side [65].

Conditional prediction. - If the AI-SecOps independents continue to raise and ship against an *augment-the-SOC-lifecycle* framing — disclosed funding rounds, named-outlet enterprise deployments, and analyst category placement that positions them as a layer over the SIEM rather than a replacement — the layer-dependent thesis holds and the category forms on top of the stack. - If instead a credible AI-native vendor raises specifically to displace the SIEM, and a named-outlet enterprise win documents a SIEM rip-and-replace by an AI-native platform, the thesis weakens and the category becomes a replacement contest. - I am grounding this on vendor funding, named-outlet enterprise deployments, and analyst coverage, not on hypotheticals.

Sources. [61] [62] [63] [64] [65]

Winners

Because this front has no Gravity vendor and most contenders are early-stage privates, I am framing winners at the level of *categories and structural positions*, not naming any individual vendor as a winner. Three positions are winning.

The platform incumbents are winning the absorption race. The clearest advancing position in this front is not a pure-play at all — it is the platform that buys or builds the AI-security capability and cross-sells it into an installed base. Two strands of public evidence point the same way: the M&A record (Check Point/Lakera at ~$300M; Palo Alto Networks/Protect AI into Prisma AIRS; Cisco/Robust Intelligence), and the embedding motion across Microsoft Defender for Cloud AI, CrowdStrike's Falcon protections for AI, and Prisma AIRS. Together they are, in my read, the strongest evidence available that the durable AI-security control is consolidating into platforms faster than any independent can graduate [52] [55]. The winning *position* is platform-with-distribution, not any single named acquirer.

The compliance-and-procurement go-to-market is winning a defensible lane in governance. The sub-market position that is converting fastest on public signals is governance sold through procurement and regulatory-alignment channels. The public-sector route is exemplified by Credo AI's Carahsoft Master-Government-Aggregator partnership and the analyst recognition clustering around it [57] [58]. My read is that this lane wins because it attaches AI governance to a buying center with budget authority and a mandate, independent of the SOC: legal, compliance, and public-sector procurement. I am attributing the win to the *lane*, not declaring the vendor that currently occupies it the category winner. That is the open bridging question of Pattern Claim 2.

The augment-the-SOC architectural posture is winning AI-Native SecOps. Within AI-native security operations, the position drawing capital and shipping into enterprises is the layer that sits *on top of* the SIEM/XDR stack — agentic triage, detection engineering, and automation that act on existing telemetry rather than replacing it. The funding flowing to SOC-lifecycle augmentation and to automation going horizontal is, in my view, the public signal that this architectural posture is the winning one in the sub-market [61] [62] [63]. The win is the *posture*, characterized at the category level — not a verdict on which AI-SOC vendor survives.

Losers / Casualties

This front has *no* individual casualty I am willing to name, and I will not manufacture one — manufacturing a casualty in a pre-Gravity market of sub-$60M-raised privates would be exactly the directional verdict the front's discipline forbids. The two consolidation events of the period — Check Point/Lakera and Palo Alto Networks/Protect AI — are *exits, not deaths*: the category's most-funded independents were acquired because they were strong, and reading either as a casualty would invert what the public record shows [52] [53]. The P2-0 currency scan confirmed no deaths and no renames across the front's tracked vendors as of mid-2026.

What is losing here is not a vendor but a *position*: **the standalone-AI-security-pure-play-as-a-durable-independent path is the losing structural bet of 2026**, and I state that at the category level. The honest read is not that any specific independent is failing — several are well-capitalized and on-thesis — but that the *path* of remaining independent all the way to Gravity is the one the public M&A record keeps foreclosing, as the consolidation exits demonstrate. The casualty, such as it is, is an *architecture of the market*: the assumption that a buyer's safest AI-security purchase is a standalone pure-play. The contested vendors therefore belong under Watch, not here.

Watch

The first pure-play graduation-or-absorption is the signal to watch above all others. The Pre-Gravity Window thesis resolves on whichever comes first. On the graduation side: a pure-play AI-security acquisition above ~$200M as a standalone valuation, or a pure-play crossing to a public listing or a disclosed $100M-plus standalone private round. On the absorption side: a fourth named independent absorbed into a platform, which leaves the window a way station [56]. Either signal will surface in named-outlet M&A coverage, SEC filings, or a dated funding announcement.

The governance bridging-vendor is the wildcard motion to watch in AI Governance & Assurance, tracked at the category level only. Whether any single governance vendor surfaces *both* a security-org motion and a compliance-org motion — bridging the two buying centers Pattern Claim 2 describes — is the open category question. Three public signals are worth monitoring: vendor go-to-market disclosures, named partnerships (the Carahsoft public-sector route is one such dated signal), and Gartner's AI TRiSM coverage consolidating or holding apart the governance and runtime-inspection layers [57] [60]. I am watching the *category's* convergence, not handicapping individual vendors.

The AI-SOC layer-versus-replacement question is the architectural motion to watch in AI-Native SecOps. The layering thesis stays intact as long as the AI-native vendors keep raising and shipping against an augment-the-SOC framing. It weakens the moment a credible AI-native platform raises

specifically to displace the SIEM and a named-outlet enterprise win documents a rip-and-replace [62] [65]. This is a category-level architectural watch, cross-referenced to Front 5's agentic-SOC analysis — not a survival call on any one AI-SOC vendor.

7.7 The Campaign Ahead

This is the youngest front in the report, and paradoxically the one whose decisive signals are easiest to watch — because in a pre-Gravity market, every move that matters is a public funding round, a public acquisition, or an analyst drawing a category boundary for the first time. None of the five signals below requires privileged access; each has a falsifiable threshold and a public source class anyone can monitor through H2 2026. Critically, not one of them is a regulatory date: this front's analytic trap is to mistake a policy deadline for a market event, and I am holding every threshold to a public *market* signal instead.

1. **The first >$200M AI-security pure-play acquisition.** Signal: whether any standalone AI-security pure-play (AI App Sec, AI-Native SecOps, or AI Governance) is acquired by a platform incumbent above ~$200M as a standalone valuation, in named-outlet M&A coverage. Threshold: a documented >$200M pure-play AI-security acquisition marks the absorption pattern reaching nine-figure-plus scale, and closes the question of whether the category exits as features or as franchises. Continued sub-$200M tuck-ins leave the window in its current narrow-but-open state. Primary source: acquirer SEC filings (Form 8-K) and named-outlet M&A coverage [66].

2. **The first pure-play crossing to Gravity.** Signal: whether any AI-security pure-play files a public S-1, completes an IPO, or discloses a $100M-plus standalone private round at a standalone valuation — the inverse signal that the window stayed open long enough for a pure-play to graduate. Threshold: a public listing or a disclosed $100M-plus standalone round validates that an AI-security pure-play can reach Gravity independent of a platform. A period passing with no pure-play crossing that line, while absorption continues, confirms the window is closing on the absorption side. Primary source: SEC EDGAR S-1 filings and dated funding announcements in named outlets [67].

3. **A Gartner AI-security / AI-TRiSM Magic Quadrant or Market Guide refresh.** Signal: whether Gartner's AI TRiSM coverage advances from Market Guide toward a Magic Quadrant, and whether its four-layer model continues to consolidate the AI-governance and AI-runtime-inspection layers into a single named market segment. Threshold: a Gartner Magic Quadrant for AI security/TRiSM, or a Market Guide refresh that names a consolidated governance-plus-runtime segment, confirms the analyst-recognized category is cohering. A held or fragmented coverage map signals the category is still pre-consolidation. Primary source: the published Gartner AI TRiSM Market Guide / Magic Quadrant and vendor announcements citing placement [60].

4. **AI-governance vendor funding rounds and bridging go-to-market.** Signal: whether an AI-governance vendor raises a disclosed round explicitly to span both buying centers — runtime/security-org capability *and* compliance-org audit/procurement reach (e.g., a public-sector-aggregator route like Credo's Carahsoft partnership) — surfacing in named-outlet funding coverage or the vendor's own disclosure. Threshold: a disclosed governance round funding a documented two-buying-center motion confirms the bifurcation is resolving toward a bridging winner.

Governance rounds that stay anchored to a single buying center keep the category bifurcated. Primary source: dated funding announcements in named outlets and vendor partnership press releases [68].

5. **AI-SecOps platform funding and product-page positioning changes.** Signal: whether the AI-native SecOps vendors keep raising against an augment-the-SOC-lifecycle framing and keep their product-page heroes positioned as a layer over the SIEM/XDR stack. The inverse signal is any credible AI-native vendor raising specifically to *displace* the SIEM, with a named-outlet enterprise rip-and-replace win to match. Threshold: continued funding and product-page positioning around SOC-lifecycle augmentation confirms the layering thesis; a funded SIEM-replacement claim plus a documented enterprise rip-and-replace weakens it. Primary source: dated funding announcements, vendor product-page hero snapshots, and named-outlet enterprise-deployment coverage [69].

References & Footnotes

1. Prisma AIRS (AI Runtime Security) markets defense against 30+ prompt-injection and jailbreak techniques, 1,000+ sensitive-data-pattern scanning, AI-agent discovery across cloud/SaaS/endpoints, and per-agent identity with RBAC and audit trails — illustrating the AI App Sec runtime-defense primitive. Palo Alto Networks, "Securing the AI Enterprise — Introducing Prisma AIRS 3.0," and Prisma AIRS AI Runtime Security datasheet, accessed 2026-06-16. https://www.paloaltonetworks.com/blog/2026/03/prisma-airs-3-0-autonomous-ai/ ↩

2. Microsoft Security Copilot agents — the phishing-triage agent reported as surfacing 6.5x more malicious alerts than analysts alone, and a Security Analyst Agent running multi-step investigations across Defender and Sentinel telemetry — illustrating the AI-Native SecOps primitive (AI doing security operations). Microsoft Security Blog, "Security Copilot in Defender: empowering the SOC with assistive and autonomous AI," accessed 2026-06-16. https://techcommunity.microsoft.com/blog/microsoftthreatprotectionblog/security-copilot-in-defender-empowering-the-soc-with-assistive-and-autonomous-ai/4503047 ↩

3. Cross-front boundary policy — AI App Sec owns the AI *runtime* (prompts, agent tool calls, output filtering) while DLP owns *egress* regardless of channel; AI-augmented detection shipping into existing SIEM/XDR is characterized in Front 5, with AI-Native SecOps sitting as an analytic layer on top of the SIEM substrate rather than replacing it. State of Cyber Security Markets 2026, the report's taxonomy §7 (Part-2 Delta, Layer 3) and §2.6. ↩↩

4. Pre-Gravity / absorption-into-platforms finding — Palo Alto Networks acquired pure-play AI-security startup Protect AI for a reported figure above USD 500M (officially undisclosed; reported at USD 500M-700M in named-outlet coverage) covering model-manipulation, prompt-injection, and training-data-poisoning defense, and Cisco acquired Robust Intelligence, an AI-security pure-play; both folded into platforms. CrowdStrike's Falcon AIDR and Continuous Identity for AI Agents extend the same incumbent-absorption pattern. Yahoo Finance, "Palo Alto Networks Acquires Protect AI For $500M+," accessed 2026-06-16, https://finance.yahoo.com/news/palo-alto-networks-acquires-protect-181130061.html ; Cisco Blogs, "Fortifying the future of Security for AI: Cisco Announces intent to acquire Robust Intelligence," accessed 2026-06-16, https://blogs.cisco.com/news/fortifying-the-future-of-security-for-ai-cisco-announces-intent-to-acquire-robust-intelligence ↩

5. Agentic AI Security market USD 1.65B (2026) → USD 13.52B (2032) at 42.0% CAGR; threat-detection-and-response largest segment at 23.1% of 2026 spend; North America 41.9% of 2026. MarketsandMarkets, "Agentic AI Security Market worth $13.52 billion by 2032," 2026-05-22, accessed 2026-06-16.

https://www.prnewswire.com/news-releases/agentic-ai-security-market-worth-13-52-billion-by-2032-marketsandmarkets-302761232.html ↩

6. Broader AI-in-cybersecurity lens — MarketsandMarkets puts the market at USD 25.53B (2026) → USD 50.83B (2031) at 14.8% CAGR; Grand View Research estimates USD 25.35B (2024) → USD 93.75B (2030) at 24.4% CAGR; the spread is a definitional artifact of how much embedded-AI tooling each firm counts (much of which this report characterizes inside Fronts 5–6). MarketsandMarkets, "Artificial Intelligence in Cybersecurity Market," accessed 2026-06-16, https://www.marketsandmarkets.com/Market-Reports/artificial-intelligence-ai-cyber-security-market-220634996.html ; Grand View Research, "AI In Cybersecurity Market Size, Share I Industry Report, 2030," accessed 2026-06-16, https://www.grandviewresearch.com/industry-analysis/artificial-intelligence-cybersecurity-market-report ↩

7. AI governance lens — MarketsandMarkets sizes the market at USD 0.89B (2024) → USD 5.78B (2029) at 45.3% CAGR; Research and Markets puts 2026 at ~USD 0.61B → USD 2.63B (2030) at 44.3% CAGR. Governance is the thinnest-scoped and most divergent slice. MarketsandMarkets, "AI Governance Market," accessed 2026-06-16, https://www.marketsandmarkets.com/Market-Reports/ai-governance-market-176187291.html ; Research and Markets, "AI Governance Market Report 2026," accessed 2026-06-16, https://www.researchandmarkets.com/reports/5951966/ai-governance-market-report ↩

8. KuppingerCole stood up an *Emerging AI Security Operations Center (SOC)* report category in 2026 and named Microsoft an overall leader — the creation of an analyst category is itself a public signal that AI-Native SecOps has crossed from feature to market. Microsoft Security Blog, "Microsoft named an overall leader in KuppingerCole Analyst's 2026 Emerging AI Security Operations Center (SOC) report," 2026-05-06, accessed 2026-06-16. https://www.microsoft.com/en-us/security/blog/2026/05/06/microsoft-named-an-overall-leader-in-kuppingercole-analysts-2026-emerging-ai-security-operations-center-soc-report/ ↩

9. Microsoft Security Copilot hosts third-party partner agents (OneTrust Privacy Breach Response, Tanium Alert Triage, BlueVoyant SecOps Tooling, Aviatrix Network Supervisor, Fletch Task Optimizer), making the incumbent platform the distribution channel for AI-security capability. Microsoft Security Blog, "Microsoft unveils Microsoft Security Copilot agents and new protections for AI," 2025-03-24, accessed 2026-06-16. https://www.microsoft.com/en-us/security/blog/2025/03/24/microsoft-unveils-microsoft-security-copilot-agents-and-new-protections-for-ai/ ↩

10. Agent identity as the contested control point — CrowdStrike's Continuous Identity for AI Agents (powered by its SGNL acquisition) grants, denies, and revokes agent access in real time with no standing privileges, announced at Identiverse 2026 (June); Palo Alto's Prisma AIRS 3.0 assigns per-agent identity with RBAC and audit trails. CrowdStrike, "CrowdStrike Unveils Continuous Identity for AI Agents," 2026-06-15, accessed 2026-06-16, https://www.crowdstrike.com/en-us/press-releases/crowdstrike-unveils-continuous-identity-for-ai-agents/ ; Palo Alto Networks, "Securing the AI Enterprise — Introducing Prisma AIRS 3.0," accessed 2026-06-16, https://www.paloaltonetworks.com/blog/2026/03/prisma-airs-3-0-autonomous-ai/ ↩

11. Regulatory backdrop for the AI Governance & Assurance slice (cited as context, not as a market-timing mechanism) — the EU AI Act, the NIST AI Risk Management Framework, and ISO 42001 (AI management system standard). Cited here as the frameworks the governance buyer must demonstrate alignment to; this front grounds no forward prediction on a regulatory effective date. ISO, "ISO/IEC 42001:2023 Artificial intelligence management system," accessed 2026-06-16. https://www.iso.org/standard/81230.html ↩

12. HiddenLayer homepage, hero "The most comprehensive security platform for AI" and supporting line "Backed by patented technology and industry-leading adversarial AI." https://www.hiddenlayer.com/ (accessed 2026-06-16). ↩

13. HiddenLayer homepage, capability framing: securing "agentic, generative, and predictive AI applications" via "AI Discovery, AI Supply Chain Security, AI Attack Simulation" and related runtime protection. https://www.hiddenlayer.com/ (accessed 2026-06-16). ↩

14. HiddenLayer has raised ~$56M in total, anchored by a ~$50M Series A (September 2023) led by M12 (Microsoft) and Moore Strategic Ventures; independent, ~169 employees as of March 2026; founded Austin, 2022 by Chris Sestito, Tanner Burns, and Jim Ballard. https://www.hiddenlayer.com/ ; https://www.crunchbase.com/organization/hiddenlayer (accessed 2026-06-16). ↩

15. Check Point AI Security page, verbatim pillars: "AI Is Moving From Language To Action. Are You Ready?"; "Secure AI usage across browsers, SaaS, and copilots."; "Protect model inputs and outputs at the application edge."; "Control tool calls, file access, and autonomous runtime behavior." https://www.checkpoint.com/ai-security/ (accessed 2026-06-16). ↩↩

16. Check Point acquired Lakera (LLM/AI application security) for a figure reported at ~$300M, announced 2025-09-16 and closed 2025-10-22, establishing it as the foundation of Check Point's Global Center of Excellence for AI Security. https://cxotoday.com/media-coverage/check-point-software-technologies-completes-acquisition-of-lakera-to-secure-the-worlds-ai-transformation/ (accessed 2026-06-16). ↩

17. Palo Alto Networks Prisma AIRS page, hero "Introducing Prisma AIRS 3.0" and tagline "Secure the entire enterprise AI lifecycle and get the visibility, assurance and runtime governance to deploy autonomous AI agents safely."; platform statement "The Prisma AIRS platform secures all AI agents, apps, models and data from development to deployment." https://www.paloaltonetworks.com/prisma/prisma-ai-runtime-security (accessed 2026-06-16). ↩↩

18. Palo Alto Networks announced intent to acquire Protect AI (ML/AI model security) on 2025-04-28 and completed the acquisition on 2025-07-22 (PANW fiscal Q4 FY2025; fiscal year ends July 31), folding the capability into Prisma AIRS. https://www.paloaltonetworks.com/company/press/2025/palo-alto-networks-completes-acquisition-of-protect-ai (accessed 2026-06-16). ↩

19. Anvilogic homepage, hero "Scale your SOC, Keep your stack." and platform descriptor "Agentic SecOps platform that onboards, searches, detects, and investigates across every data source — without moving a single byte." https://www.anvilogic.com/ (accessed 2026-06-16). ↩

20. Anvilogic homepage, deployment pillar "Standalone on your security data lakes or alongside your SIEM"; lifecycle AI framing across AI-enhanced detection and investigation, automated detection tuning with ML recommendations, and agentic triage; named data-lake support for Snowflake, Databricks, and Azure. https://www.anvilogic.com/ (accessed 2026-06-16). ↩

21. Anvilogic raised a $45M Series C led by Evolution Equity Partners; total funding ~$85M; founded 2019, independent. https://www.anvilogic.com/ ; https://www.crunchbase.com/organization/anvilogic (accessed 2026-06-16). ↩

22. Tines homepage, hero "The intelligent workflow platform" and core line "Securely scale AI and automation. Integrate agents, teams, and tools with speed and control." https://www.tines.com/ (accessed 2026-06-16). ↩

23. Tines homepage, horizontal-positioning pillars "No matter how many tools you change, you can keep Tines in place" and explicit scope across "Security, IT Ops, Infrastructure, Engineering, and Product teams"; three workflow types — human-led, deterministic, and agentic. https://www.tines.com/ (accessed 2026-06-16). ↩

24. Tines raised a $125M Series C lifting its valuation to ~$1.125B (unicorn); ~$272M raised in total; based in Dublin. https://www.tines.com/ ; https://www.crunchbase.com/organization/tines (accessed 2026-06-16). ↩

25. Prophet Security homepage, hero "AI SOC Agents that Accelerate Detection and Response" and pillar "Investigate every alert in minutes. Hunt threats before they surface. Close detection gaps before attackers find them." https://www.prophetsecurity.ai/ (accessed 2026-06-16). ↩

26. Prophet Security homepage, capability framing: AI SOC agents that "work every alert end-to-end," "hunt before attackers move," and "strengthen detections before gaps become incidents." https://www.prophetsecurity.ai/ (accessed 2026-06-16). ↩

27. Prophet Security has raised ~$41M total across a Bain Capital Ventures seed, a ~$30M Accel-led Series A, and a strategic round joined by American Express Ventures and Citi Ventures dated 2026-02-25; founded by Kamal Shah and Arun Sreeshankar. https://www.prophetsecurity.ai/ ; https://www.crunchbase.com/organization/prophet-security (accessed 2026-06-16). ↩

28. Credo AI homepage, hero "Govern AI Everywhere." and tagline "One platform to discover, assess, and govern every AI agent, model, and application — continuously and in context." https://www.credo.ai/ (accessed 2026-06-16). ↩

29. Credo AI homepage, pillars "Discover and catalog every AI system — agents, models, and applications" and "Continuous, contextual risk assessment for bias, security, privacy, and compliance"; framework mapping to EU AI Act, NIST AI RMF, ISO 42001, SOC 2, GDPR, HITRUST. https://www.credo.ai/ (accessed 2026-06-16). ↩

30. Credo AI: January 2026 Carahsoft partnership (public-sector procurement); #6 Applied AI on Fast Company Most Innovative Companies 2026; Gartner "Cool Vendor in AI Cybersecurity Governance"; total funding ~$41.3M, anchored by a $21M Series B (July 2024); Palo Alto CA, founded 2020. https://www.credo.ai/ ; https://www.crunchbase.com/organization/credo-ai (accessed 2026-06-16). ↩

31. Holistic AI homepage, verbatim positioning line "The end-to-end AI governance platform trusted by global enterprises running AI at scale." https://www.holisticai.com/ (accessed 2026-06-16). ↩

32. Holistic AI homepage, pillars "Identify," "Protect," "Enforce"; coverage spanning AI model discovery and inventory, bias and safety testing (40+ specialized tests), shadow-AI discovery across cloud and code, and regulatory mapping to the EU AI Act, NIST AI RMF, and ISO 42001. https://www.holisticai.com/ (accessed 2026-06-16). ↩

33. Holistic AI: UK-origin (London, founded 2020), algorithmic auditing and AI governance. The widely circulated ~$200M funding figure does NOT belong to this London AI-governance vendor (holisticai.com); it refers to a similarly-named but distinct company (a Paris-based generative-AI-model startup). The London governance vendor's disclosed funding is in the low tens of millions (a reported ~$35M Series B, May 2024); no fresh post-baseline round was confirmed in the bounded scan, and no funding figure is treated as load-bearing here. Position rests on the dated vendor surface (blog activity into May 2026) and named-outlet / buyer-guide corroboration. https://www.holisticai.com/ ; https://www.crunchbase.com/organization/holistic-ai (accessed 2026-06-16). ↩

34. Check Point Software Technologies, "Check Point to Acquire Lakera, Setting the Standard for AI Security," September 16, 2025, https://www.globenewswire.com/news-release/2025/09/16/3150869/0/en/Check-Point-Acquires-Lakera-to-Deliver-End-to-End-AI-Security-for-Enterprises.html (accessed 2026-06-16). Lakera (AI-application security, red-teaming) acquired for a reported ~$300M; announced 2025-09-16. ↩

35. Palo Alto Networks, "Palo Alto Networks Announces Intent to Acquire Protect AI," April 28, 2025, https://www.paloaltonetworks.com/company/press/2025/palo-alto-networks-announces-intent-to-acquire-protect-ai–a-game-changing-security-for-ai-company (accessed 2026-06-16). Protect AI folded into the Prisma AIRS AI-security runtime offering; acquisition intent announced 2025-04-28. ↩

36. Cisco, "Cisco Completes Acquisition of Robust Intelligence," 2024, https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2024/m08/cisco-completes-acquisition-of-robust-intelligence.html (accessed 2026-06-16). Robust Intelligence (AI model security) acquired 2024; built into Cisco's AI-security narrative. Cisco's primary characterization in this report is in Front 5; this is a cross-front consolidation event referenced here. ↩

37. HiddenLayer (private): the most prominent independent AI-application-security vendor not yet absorbed into a platform incumbent. $50M Series A, September 2023; ~$56M raised to date; M12 (Microsoft's venture fund) among investors. HiddenLayer, "HiddenLayer Raises $50 Million Series A," 2023, https://www.hiddenlayer.com/news/hiddenlayer-raises-50m-in-series-a-funding-to-safeguard-ai (accessed 2026-06-16). ↩

38. Microsoft Learn, "AI security posture management — Microsoft Defender for Cloud," 2026, https://learn.microsoft.com/en-us/azure/defender-for-cloud/ai-threat-protection (accessed 2026-06-16); CrowdStrike, "Falcon for AI" product page, https://www.crowdstrike.com/en-us/platform/ (accessed 2026-06-16); Palo Alto Networks, "Prisma AIRS" product page, https://www.paloaltonetworks.com/prisma/prisma-ai-runtime-security (accessed 2026-06-16). Microsoft, CrowdStrike, and Palo Alto Networks are characterized as primary contenders in Front 5; here they appear only as §7.4 AI-security plays via the named AI-security product line each attaches to its platform (Defender for Cloud AI; Falcon for AI; XSIAM AI / Prisma AIRS). ↩

39. Cross-front primary-slot note: Microsoft, CrowdStrike, and Palo Alto Networks are analyzed as primary §N.3 contenders in Front 5 (Detection & Response) per the report's §N.3-slot primary map. In Front 7 they appear only as §7.4 Plays exerting embed-rather-than-acquire pressure on the independent AI-security layer, named by their AI-security-specific product families, not by their Front-5 EDR/SIEM pillars. ↩

40. Prophet Security (private): ~$41M raised across seed and Series A; a round dated February 2026 disclosed strategic participation from American Express Ventures and Citi Ventures, positioning an "agentic SOC" automating first-pass investigation. Prophet Security, funding and strategic-investment coverage, February 2026, https://www.prophetsecurity.ai/ (accessed 2026-06-16). ↩

41. Anvilogic (private): $45M Series C led by Evolution Equity Partners; detection-engineering and AI-triage layer across existing SIEM and data-lake tooling. Anvilogic, "Anvilogic Raises $45M Series C," https://www.securityweek.com/multi-data-platform-siem-anvilogic-raises-45-million/ (accessed 2026-06-16). ↩

42. Cross-front note: the agentic-SOC thesis (AI agents doing first-pass SOC investigation) is examined as a primary matter in Front 5 (Detection & Response), where Dropzone AI is the front's Wildcard analog of the same bet. This Play references that cross-front link rather than re-grading the F5 detection-and-response Wildcard or its incumbents here. ↩

43. Credo AI (private): AI-governance platform (model inventory, risk assessment, policy mapping, audit evidence); ~$41.3M raised in total, anchored by a $21M Series B (July 2024); public-sector distribution including a Carahsoft partnership placing the platform onto government procurement vehicles. Carahsoft, "Carahsoft and Credo AI Partner to Provide AI Governance to Public Sector," https://www.carahsoft.com/credo-ai (accessed 2026-06-16); Credo AI funding history, https://www.crunchbase.com/organization/credo-ai (accessed 2026-06-16). ↩

44. Holistic AI (private): AI audit, risk, and assurance tooling positioned for enterprise and regulated buyers; occupies the same governance-and-assurance wedge as Credo AI. No fresh 2026 funding round confirmed on the public record at access time. Holistic AI, company and product pages, https://www.holisticai.com/ (accessed 2026-06-16). ↩

45. HiddenLayer (private): $50M Series A, September 2023; ~$56M raised to date; M12 (Microsoft's venture fund) among investors. The most prominent independent AI-application-security vendor not absorbed into a platform incumbent. HiddenLayer, "HiddenLayer Raises $50 Million Series A," 2023, https://www.hiddenlayer.com/news/hiddenlayer-raises-50m-in-series-a-funding-to-safeguard-ai (accessed 2026-06-16). ↩

46. Tines (private): $125M Series C at a ~$1.125B valuation, led by Goldman Sachs Alternatives; workflow-automation and agentic security-operations vendor, an adjacency to AI-application security proper. Tines, "Tines Raises $125M Series C," https://techfundingnews.com/new-unicorn-from-ireland-next-gen-ai-automation-startup-tines-raises-125m-at-1-1b-valuation/ (accessed 2026-06-16). ↩

47. Anvilogic (private): $45M Series C led by Evolution Equity Partners; ~$85M raised to date. Anvilogic, "Anvilogic Raises $45M Series C," https://www.securityweek.com/multi-data-platform-siem-anvilogic-raises-45-million/ (accessed 2026-06-16). ↩

48. Prophet Security (private): ~$41M raised across seed and Series A; strategic round dated February 2026 with American Express Ventures and Citi Ventures. Prophet Security, funding and strategic-investment coverage, February 2026, https://www.prophetsecurity.ai/ (accessed 2026-06-16). ↩

49. Credo AI (private): ~$41.3M raised in total, anchored by a $21M Series B (July 2024); investors include Sands Capital and Decibel; public-sector distribution via Carahsoft. Credo AI funding history, https://www.crunchbase.com/organization/credo-ai (accessed 2026-06-16). ↩

50. Lakera (AI-application security, red-teaming): acquired by Check Point Software Technologies for a reported ~$300M; announced 2025-09-16. A consolidation exit into a platform incumbent, not a distress event. Check Point, "Check Point to Acquire Lakera," September 16, 2025, https://www.globenewswire.com/news-release/2025/09/16/3150869/0/en/Check-Point-Acquires-Lakera-to-Deliver-End-to-End-AI-Security-for-Enterprises.html (accessed 2026-06-16). ↩

51. Protect AI: Palo Alto Networks announced intent to acquire on 2025-04-28; folded into the Prisma AIRS AI-security runtime offering. A consolidation exit into a platform incumbent, not a distress event. Palo Alto Networks, "Palo Alto Networks Announces Intent to Acquire Protect AI," April 28, 2025, https://www.paloaltonetworks.com/company/press/2025/palo-alto-networks-announces-intent-to-acquire-protect-ai–a-game-changing-security-for-ai-company (accessed 2026-06-16). ↩

52. Check Point Software acquired AI-application-security and red-teaming specialist Lakera, announced 2025-09-16, reported at ~$300M; Lakera forms the foundation of Check Point's Center of Excellence for AI Security. Check Point press release, "Check Point Acquires Lakera to Deliver End-to-End AI Security for Enterprises," and CSO Online coverage, accessed 2026-06-16. https://www.globenewswire.com/news-release/2025/09/16/3150869/0/en/Check-Point-Acquires-Lakera-to-Deliver-End-to-End-AI-Security-for-Enterprises.html ↩↩↩↩

53. Palo Alto Networks announced intent to acquire Protect AI on 2025-04-28, folding it into the Prisma AIRS AI-security runtime line. Palo Alto Networks press release, "Palo Alto Networks Announces Intent to Acquire Protect AI, a Game-Changing Security for AI Company," accessed 2026-06-16. https://www.paloaltonetworks.com/company/press/2025/palo-alto-networks-announces-intent-to-acquire-protect-ai–a-game-changing-security-for-ai-company ↩↩↩

54. Cisco acquired AI-model-security firm Robust Intelligence in 2024, building it into the company's AI-security narrative. Cisco acquisition announcement, accessed 2026-06-16. https://blogs.cisco.com/news/cisco-completes-acquisition-of-robust-intelligence ↩↩

55. Platform incumbents ship AI-security capability inside platforms the enterprise already buys: Microsoft Defender for Cloud AI security posture; CrowdStrike Falcon protections for AI; Palo Alto Networks Prisma AIRS. Vendor product pages, accessed 2026-06-16. https://www.paloaltonetworks.com/prisma/prisma-ai-runtime-security ; https://www.crowdstrike.com/en-us/platform/ ; https://learn.microsoft.com/en-us/azure/defender-for-cloud/ai-threat-protection ↩↩↩

56. As of mid-2026 no pure-play AI-security vendor has reached the public market or the ~$100M-private threshold the report's vendor-tier rubric uses for Gravity; three of the most prominent independent AI-application-security vendors (Lakera, Protect AI, Robust Intelligence) have exited into platforms within ~18 months, with HiddenLayer the most prominent remaining independent. HiddenLayer newsroom and Front 7 vendor-selection / P2-0 currency scan, accessed 2026-06-16. https://www.hiddenlayer.com/newsroom ↩↩↩

57. Credo AI named Carahsoft its Master Government Aggregator on 2026-01-07, making its AI-governance platform available to the public sector via NASA SEWP V, ITES-SW2, NASPO ValuePoint, TIPS, and OMNIA Partners contracts, positioned around NIST AI RMF and OSTP alignment. GlobeNewswire / Carahsoft, "Credo AI and Carahsoft Partner to Accelerate AI Adoption Through Purpose-Built AI Governance," accessed 2026-06-16. https://www.globenewswire.com/news-release/2026/01/07/3214812/0/en/Credo-AI-and-Carahsoft-Partner-to-Accelerate-AI-Adoption-Through-Purpose-Built-AI-Governance.html ↩↩↩↩

58. Credo AI recognized as a Gartner Cool Vendor in AI cybersecurity governance and named to Fast Company's Most Innovative Companies of 2026 (Applied AI). Credo AI newsroom and Carahsoft release (industry recognition), accessed 2026-06-16. https://www.credo.ai/ ↩↩↩

59. Holistic AI positions as an end-to-end AI governance platform with algorithmic auditing and bias testing for enterprises running AI at scale; vendor page carries dated 2026 content. Holistic AI, accessed 2026-06-16. https://www.holisticai.com/ ↩↩

60. Gartner's Market Guide for AI Trust, Risk and Security Management (AI TRiSM) models the space as four layers — AI governance, AI runtime inspection and enforcement, information governance, and infrastructure and stack — with the top two layers (governance and runtime inspection/enforcement) consolidating into a distinct market segment. Gartner document 6185655, "Market Guide for AI Trust, Risk and Security Management," accessed 2026-06-16. https://www.gartner.com/en/documents/6185655 ↩↩↩↩

61. Prophet Security, a pioneer of the agentic AI SOC, announced strategic investment from Amex Ventures and Citi Ventures on 2026-02-25 (bringing total funding to ~$41M across three rounds); its platform "works alongside security analysts and automates manual, time-sensitive workflows" across investigation, threat hunting, detection engineering, and incident response. PR Newswire, "Prophet Security Accelerates the Agentic AI SOC Movement with Strategic Investments from Amex Ventures and Citi Ventures," accessed 2026-06-16. https://www.prnewswire.com/news-releases/prophet-security-accelerates-the-agentic-ai-soc-movement-with-strategic-investments-from-amex-ventures-and-citi-ventures-302696588.html ↩↩↩

62. Anvilogic closed a $45M Series C led by Evolution Equity Partners on 2024-04-17, bringing total funding to ~$85M, to expand generative-AI features across the SOC lifecycle and enable security-data-lake adoption. SecurityWeek, "Multi-Data Platform SIEM Anvilogic Raises $45 Million," and Anvilogic Series C founders' note, accessed 2026-06-16. https://www.securityweek.com/multi-data-platform-siem-anvilogic-raises-45-million/ ; https://www.anvilogic.com/learn/series-c-founders-note ↩↩↩↩

63. Tines raised a $125M Series C to a ~$1.125B valuation (~$272M total raised) and is broadening from a security-automation pure-play into a horizontal agentic-workflow platform spanning IT and operations. BankInfoSecurity, "Tines Raises $125M to Expand Security Automation, Agentic AI," and PR Newswire (Tines agents launch), accessed 2026-06-16. https://www.bankinfosecurity.com/tines-raises-125m-to-expand-security-automation-agentic-ai-a-29000 ↩↩↩

64. The report's shared taxonomy draws the cross-front boundary placing AI-native SecOps as an analytic/agentic layer on top of the SIEM/XDR substrate rather than a replacement for it. the report's taxonomy §2.6 and Front 7 §7.1; see also Front 7 vendor-selection cross-front notes, accessed 2026-06-16. ↩↩

65. Cross-front reference: Front 5 (Detection & Response / SOC platform war) characterizes the agentic-SOC / autonomous-triage motion (including vendors such as Dropzone) from the SOC-platform side; the AI-native pure-plays in this front are the same substrate-layer pattern viewed from the pure-play side. See Front 5 §5.4 agentic-SOC analysis, this report. ↩↩↩

66. Acquirer SEC filings (Form 8-K) and named-outlet M&A coverage are the public source class for any >$200M AI-security pure-play acquisition. SEC EDGAR full-text search, accessed 2026-06-16. https://efts.sec.gov/LATEST/search-index?q=%22AI%20security%22 ↩

67. SEC EDGAR S-1 filings and dated funding announcements in named outlets are the public source class for any AI-security pure-play crossing to a public listing or a disclosed $100M-plus standalone private round. SEC EDGAR, accessed 2026-06-16. https://www.sec.gov/cgi-bin/srqsb?text=form-type%3DS-1 ; https://efts.sec.gov/LATEST/search-index?forms=S-1 ↩

68. Dated funding announcements in named outlets and vendor partnership press releases are the public source class for AI-governance vendor rounds and bridging go-to-market signals (e.g., the Credo AI / Carahsoft public-sector route). GlobeNewswire / Carahsoft, accessed 2026-06-16. https://www.globenewswire.com/news-release/2026/01/07/3214812/0/en/Credo-AI-and-Carahsoft-Partner-to-Accelerate-AI-Adoption-Through-Purpose-Built-AI-Governance.html ↩

69. Dated funding announcements, vendor product-page hero snapshots, and named-outlet enterprise-deployment coverage are the public source class for AI-SecOps positioning changes (augment-the-SOC framing versus a funded SIEM-replacement claim with a documented enterprise rip-and-replace). Prophet Security, Anvilogic, and

Tines vendor pages and funding coverage, accessed 2026-06-16. https://www.prophetsecurity.ai/ ; https://www.anvilogic.com/ ; https://www.tines.com/ ↩

Disclosures

DISCLOSURE.

This is ProductBeacon's independent, impartial market research, based solely on publicly verifiable signals, published on the open web. No vendor sponsors, no paywalled data, no analyst-firm reuse. It reflects ProductBeacon's independent view and is not investment advice. Full independence and sourcing methodology: Methodology →.

AFFILIATION.

By Yohay Etsion, Head of Product (Fractional), AXIA. AXIA is an AI-powered data-security (Insider Risk Management / DLP) company. This analysis is based only on publicly verifiable material.

NOT ADVICE.

This report does not constitute investment, legal, tax, or accounting advice. No claim should be relied upon as the basis for any investment decision. The author holds no trading position in any named public security and is not compensated by any named vendor. Readers who use this report in investment contexts bear sole responsibility for their decisions.