CYBER MADNESS

Motions & Tailwinds Report 2025-2026

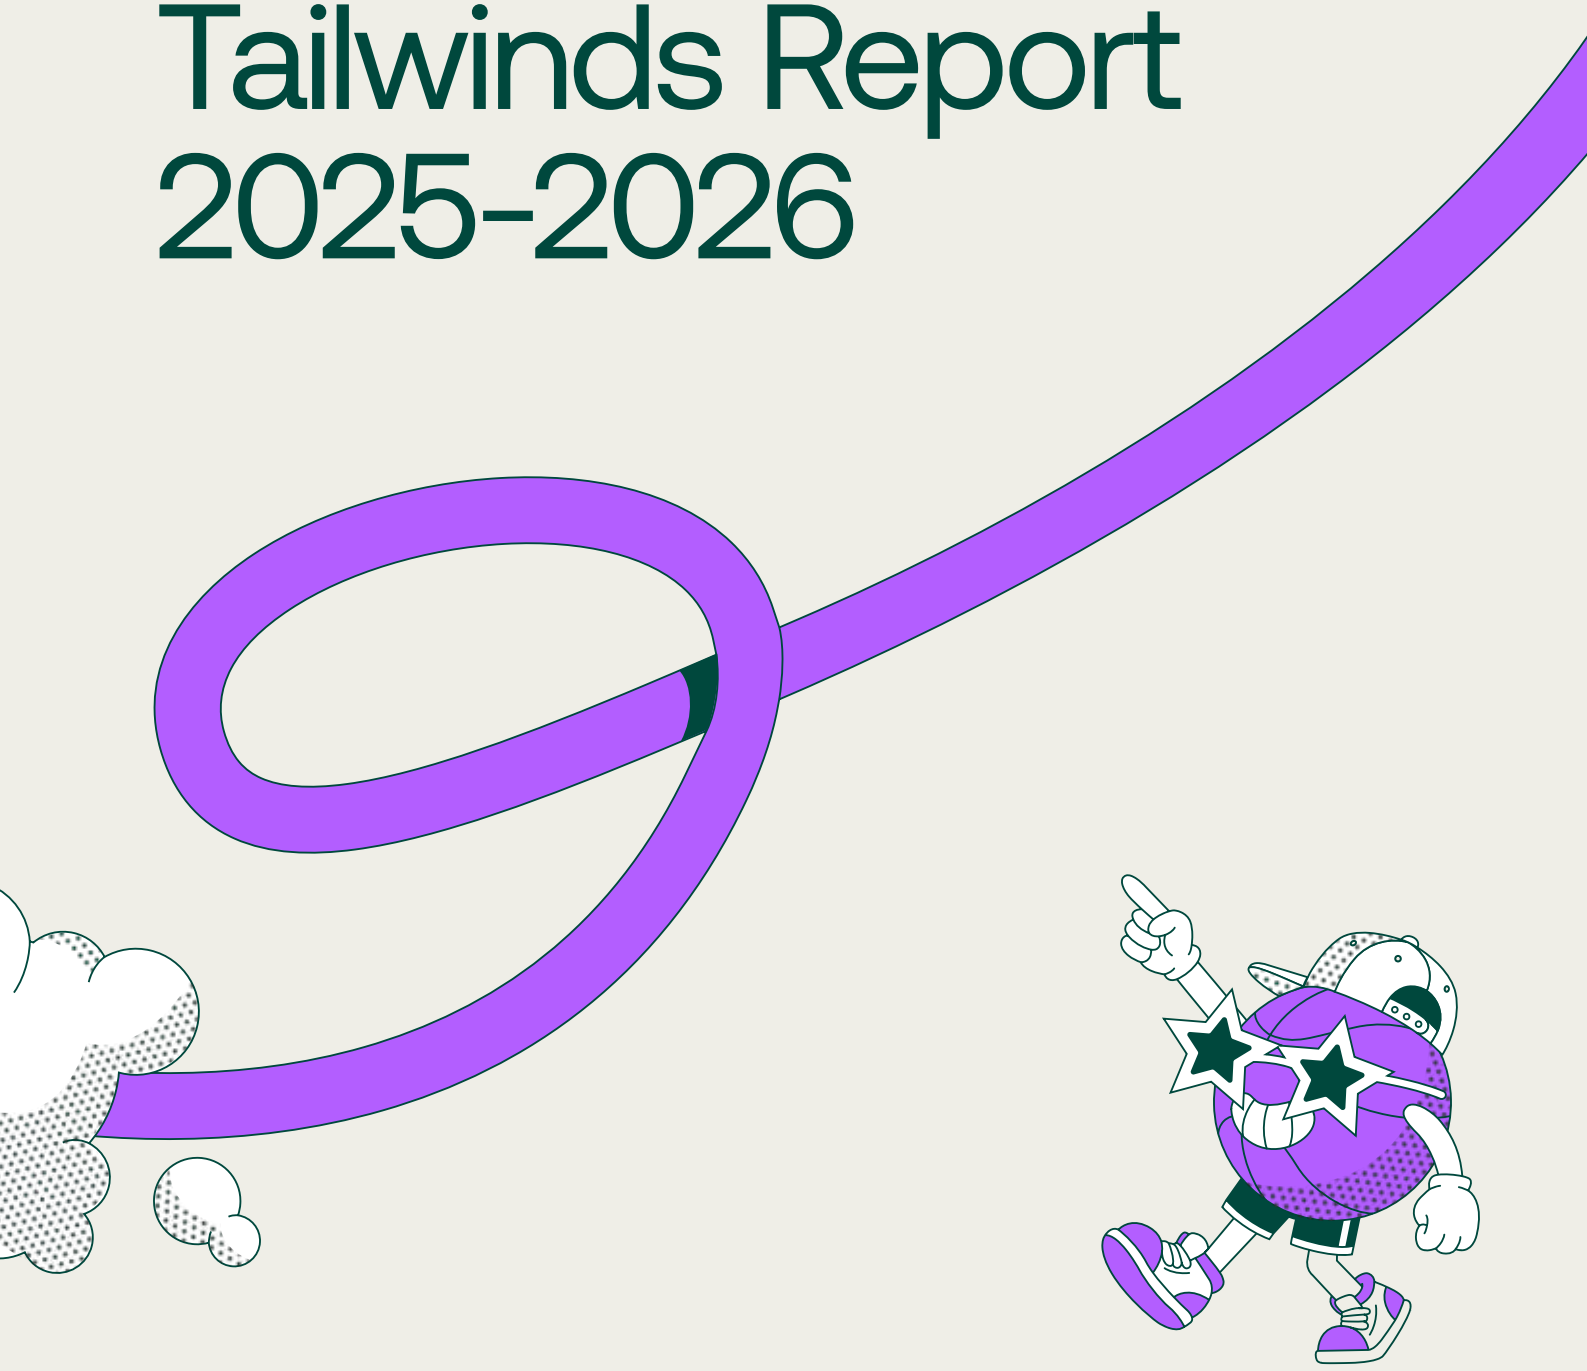

DEUTSCH&CO. X SACR

01

Motions & Tailwinds - Main Themes

Top Motions	Top Tailwinds	Disrupted Categories
Consolidation	**AI vs. AI Warfare**	**Identity, Privilege and Access Risk**
A strong market shift demanding fewer tools with broader platform capabilities, driving vendors to focus on measurable cost-efficiency like "Cost Per Investigated Alert."	The accelerated operationalization of AI by adversaries, which forces defenders to adopt AI-Native solutions and automation cycles to keep up.	Non-human identities are rapidly outnumbering human users across cloud and SaaS environments. This expanded the attack surface and accelerated the shift toward just-in-time access and zero-standing privilege models. Now vendors are beginning to build identity and access control layers specifically for autonomous agents.
Contextualization	**Shadow AI Management**	**AI-Native SecOps Platforms**
Vendors are actively making business, behavioral, and operational context the primary engine for security tools, underpinning decisions like alert triage and privilege grants.	The rapid, unmanaged spread of employee use of AI tools, which is becoming a top operational risk by bypassing security and data controls.	A new class of platforms integrating AI and autonomous agents to perform Tier-1 work in Security Operations Centers.
	LLM Marketing Fatigue	**LLM Security and Governance**
	Deep market skepticism towards generic AI claims, leading to a strong demand for transparency, measurable KPIs, and auditability in all AI-driven products.	Tools and solutions focused on managing the specific security and safety risks of Large Language Models, including measuring Hallucination Rate and Guardrail Accuracy.
	The Rise of Agents	**Security Data Pipeline and Platform Consolidation**
	The focus has moved from securing AI models to governing autonomous agents operating across enterprise systems. Security now centers on behavior, access, and execution.	The new architecture of vendor-agnostic security data pipelines, decoupling storage and compute to allow detections to run wherever the data resides.
		AI Coding and Application Security
		The AppSec domain is shifting to move critical controls upstream into the developer workflow and secure code that is written or assisted by AI.

02

2025 - A Year in Review

AI stopped being an experiment in 2025 and became operational. Across the security stack - in SOC workflows, developer environments, and adversary toolkits - AI moved from pilot projects into everyday use. So, as attackers and defenders suddenly had the ability to iterate faster and test more ideas, traditional security processes became clearly outpaced.

That shift forced real changes in how teams operate. Agentic triage began cutting through alert noise, federated detections improved signal across fragmented environments. Just-in-time privilege models started shrinking blast radius, and continuous red teaming replaced assumptions with constant validation. At the same time, resilience and recovery stopped being aspirational goals and became measurable expectations for boards and leadership teams.

This report relies on three signal types: direct feedback from CISOs, observable vendor roadmap shifts, and the operational lessons of major incidents throughout 2025. It focuses on what truly changed - and what those changes mean for 2026. The report is organized in three parts: the trends that defined 2025, five technology pivots grounded in clear signals and measurable KPIs, and practical recommendations for the year ahead.

Method

This report is based on the accumulated annual research work of the SACR and Deutsch & Co research teams, which included 200+ CISO interactions and recorded interviews, 100+ founding team interviews, and internal reviews of 50+ teams backed by industry defining VCs. It is meant to give an overall review of 2025's motions and trends, to help security professionals, founders and investors navigate the tailwinds of 2026.

About the Authors

Software Analyst Cybersecurity Research (SACR) is an independent research and advisory organization focused on helping CISOs, founders, investors, and security teams understand where cybersecurity is heading. Through in-depth industry reports, analyst research, vendor analysis, and shorter thought pieces, SACR analyzes emerging technologies, market shifts, and vendor strategies across key security domains. The firm was founded by **Francis Odum**, Founder and Chief Cybersecurity Analyst, who built SACR into one of the largest independent cybersecurity research platforms in the market. He is recognized for his work with over 60,000 security professionals worldwide and for establishing SACR as a trusted brand among CISOs and leading vendors.

Deutsch & Co is a private equity firm focused on investing in and building category leaders across cybersecurity and AI. The firm's investment strategy is grounded in proprietary research, including hundreds of annual interviews with industry leaders, buyers, and practitioners, used to identify emerging market gaps and define new categories. By combining research-driven insights with strategy, positioning, and branding, Deutsch & Co partners with companies to help shape and lead the categories they operate in. The firm was founded by **Roei Deutsch**, who serves as CEO.

- **Agentic triage** refers to a security approach in which autonomous or AI-driven agents actively analyze, prioritize, and route risks or alerts without requiring constant human interpretation.
- **Federated detections** are security detections executed across distributed data sources without centralizing the underlying data.
- **Just-in-time privilege models** grant temporary elevated permissions only when needed and revoke them immediately after use.
- **Continuous red teaming** is the ongoing simulation of adversarial attacks used to identify weaknesses and improve the safety and robustness of AI models.

Layout:

03

Overall Trends: What Dominated 2025

Throughout 2025, recurring themes emerged from private CISO forums, executive briefings, and peer discussions, revealing both the risks that dominated security leaders' attention and the priorities shaping their outlook for 2026. Here are the highlights.

Cross-sector themes:

01

The Rise of Agents

In 2025, the industry conversation shifted from "security for AI/LLMs" to securing agentic platforms themselves. Early discussions focused on model risks such as prompt injection, hallucinations, and data leakage. But as enterprises moved beyond experimentation and began embedding AI into real workflows, the conversation rapidly shifted. LLMs stopped being static assistants or chatbots and began evolving into agents, capable of taking actions across systems, chaining tools together, and operating with increasing autonomy.

This shift reframed the security challenge: the problem was no longer just securing models, but governing the behavior, access, and execution of autonomous systems interacting with enterprise infrastructure. **By the end of 2025, it became clear that the next phase of cybersecurity will center not simply on AI safety, but on securing agentic systems operating inside real enterprise environments.**

As organizations began operationalizing agentic systems inside real enterprise environments, a deeper challenge quickly emerged: governing the identities and actions of machines operating at scale. Traditional IAM and SSO architectures, designed for predictable human logins, struggle to manage agents that execute high-frequency, non-deterministic tasks using long-lived credentials across multiple systems. As a result, 2026 will see the emergence of Agentic Identity Access Platforms (AIAP): a new identity control layer that acts as an "SSO for Agents," brokering task-scoped, ephemeral identities based on agent intent. **In the agentic era, identity security will shift from verifying who is acting to continuously governing why an action is occurring and how long access should exist.**

Key industry moves and company launches

Companies evolved in the past year				
KEYCARD Access Governance	Fabrix Access Governance	AgentField Access Governance	scalekit Access Governance	token Access Governance
ONYX Agent monitoring & runtime	Straiker Agent monitoring & runtime	Singulr AI Agent monitoring & runtime	EVÉ Agent monitoring & runtime	
Geordie Agent monitoring & runtime	Manifold Agent monitoring & runtime	Virtue A Agent monitoring & runtime	Aiceberg Agent monitoring & runtime	
Aurascape Agent monitoring & runtime	Capsule Agent monitoring & runtime	Judgment Labs Agent monitoring & runtime	zenity Agent monitoring & runtime	
NOMA Agent monitoring & runtime	MINT (in stealth)			
Moves by legacy companies (M&A/ launches)				
okta Launched "Okta for AI Agents	paloalto Launched a module for real-time monitoring of agentic AI endpoints and acquired Protect AI	Microsoft Launched Entra Agent ID		
CROWDSTRIKE Launched visibility and control of AI agents and acquired Pangea and SGNL				
Companies acquired in the past year				
pangea Agent monitoring & runtime	Prompt: Agent monitoring & runtime	aim security Agent monitoring & runtime	Invariantlabs Agent monitoring & runtime	
PROTECT AI Agent monitoring & runtime	LAKERA Agent monitoring & runtime	Calypso AI Agent monitoring & runtime	CYATA Access Governance	

> *"Amid the excitement of embracing the next wave of generative AI, companies are moving quickly to deploy agentic use cases, often overlooking the critical need to secure these systems and control the sprawl of non-human identities."*

Arnab Bose Former Chief Product Officer — okta

> *"AI agents are quickly becoming a new class of workforce in the enterprise, but they require more complex identity lifecycle management than human users."*

Itamar Apelblat CEO & Co-Founder — token

> *"AI agents don't just authenticate, they take action, call APIs, chain workflows, and make decisions. Securing them requires treating identity as a runtime control plane, not a one-time configuration."*

Ido Shlomo CTO & Co-Founder — token

> *"Enterprises are moving beyond simple AI chatbots to fully autonomous agents -but with this evolution comes an exponential rise in security and safety risks. The threat vector has escalated from basic prompt injection attacks to mass data exfiltration, supply chain attacks, and even autonomous chaos."*

Ankur Shah CEO & Co-Founder — Straiker

02

Shadow AI Was the Most Consistent CISO Concern of 2025

Across CISO dinners, closed-door forums, and year-end briefings, Shadow AI emerged as the most cited operational risk. **AI tools proved useful enough that adoption quickly outpaced security's ability to inventory and govern them.** Employees uploaded sensitive data into copilots, browser extensions, and SaaS AI features with little visibility into data flow or retention.

This wasn't reckless behavior - AI was becoming embedded in everyday workflows, making its use inevitable. Blocking it outright didn't work; it simply pushed activity into unmonitored channels. The real concern was the absence of control, logging, and policy enforcement. Shadow AI became a symptom of a deeper issue: security teams losing visibility into how work actually happens.

By 2025, the conclusion was unavoidable: AI is here. Security teams must navigate it - without losing control.

Security leaders described a tension between innovation and risk. Development teams moved quickly, leadership encouraged experimentation, and security teams were asked to approve systems they barely understood yet.

The most effective organizations responded by narrowing scope. They permitted AI in defined domains, such as SOC triage, documentation, and internal tooling, while enforcing guardrails on data access, actions, and auditability.
The lesson for 2026 is not to slow AI, but to effectively constrain the boundaries in which it operates.

03

Autonomous Offense vs. Defense: AI-Driven Cyber Warfare

2025 was a transition year: security programs began a slow pivot away from non–AI-native stacks, just as adversaries operationalized AI tooling at scale. The resulting shift in attack patterns - faster iteration, lower-cost experimentation, and compressed time-to-breach - made legacy approaches objectively insufficient. For CISOs, this became an additional, concrete driver for modernization: when both sides are automated, advantage accrues to the party that closes the loop first. **In practice, "AI vs. AI" is simply competing automation cycles: faster learning, tighter feedback, and broader execution.**

That acceleration also shifted the core risk from data theft to operational paralysis. Incidents like the Jaguar Land Rover attack (estimated $2.5B impact), SaaS-jacking campaigns targeting major SaaS platforms, and the Salt Typhoon espionage activity showed attackers increasingly optimizing for continuity disruption and deep infrastructure access - not just exfiltration. The implication is blunt: security failures are now business-stopping events, not merely confidentiality breaches.

This reality resulted in two new priorities for security teams: **resilience and recovery**.

Resilience

Organizations moved away from periodic assurance toward continuous validation, treating security posture as something you prove, not assume. Continuous red teaming and autonomous attack simulation began shifting from "advanced program" to an operating baseline, **continuously exercising real attacker paths and measuring whether controls still prevent, detect, and contain under evolving tactics**. Just as importantly, teams started validating execution, not just documentation: playbooks were drilled through repeatable technical exercises to surface latency, ownership gaps, and brittle dependencies - turning response into a practiced capability rather than a binder on a shelf.

Recovery

In 2025, cyber resilience increasingly became the board's yardstick for security performance: assume incidents will occur, and measure how quickly the business can restore critical operations. That reframed investment toward verified, attack-resistant recovery - **tiered RTO/RPO commitments, dependency-aware restoration sequencing, and routine proof that backups are recoverable and operationally usable**. Teams also reduced human latency by operationalizing response and recovery workflows through orchestration and automation, turning playbooks into executable procedures rather than documentation. Mature programs now manage recovery readiness like an SLO: restoration time for Tier-0 services, recovery success rate, and time-to-restore under realistic conditions - continuously validated and extended across critical third- and fourth-party dependencies.

- **Shadow AI** is the unauthorized or unsanctioned use of AI tools within an organization without formal oversight from IT or security teams.

Select industry moves and company launches

Companies launched in the past year

ARMADIN | Novee | Gambit | BOLD. | Quantro | cogent

terra | RUNSYBIL | TENZAI | XBOW

Moves by legacy companies (M&A/ launches)

Microsoft
(Expanded security Copilot agents)

PENTERA
(Acquired EVA)

paloalto
(Launched Prisma AIRS 2.0)

SentinelOne
(Upgraded Purple AI "Athena" - agentic auto-triage and auto-investigation)

Companies acquired in the past year

EVA

> *"We're moving from the analyst's hands on the steering wheel to autonomous actions that are effective, safe, and reliable. We need full autonomy; it's not optional, it's foundational."*

George Kurtz CEO — CROWDSTRIKE

> *"I believe within the next few years virtually all cyberattacks will be AI-based - swarming, tailored, and relentless. They will be untethered to human limitations and capable to execute on a scale we have never witnessed before."*

Kevin Mandia CEO & Founder — ARMADIN

- A 2025 cyberattack that forced Jaguar Land Rover to halt production for several weeks, disrupting global supply chains and causing an estimated **$2.5B economic impact.**
- A Cyberattack in which attackers hijack **SaaS accounts or cloud applications** through stolen credentials, tokens, or misconfigured integrations, allowing persistent access to organizational system
- The cyber-espionage campaign linked to the **Chinese state-backed group "Salt Typhoon"**.

> *"Attackers don't wait for your annual pentest. Neither should your defense. What security teams actually need are high-signal findings they can trust: novel vulnerabilities that are proven exploitable."*

Ido Geffen CEO & Co-Founder

> *"Fully autonomous testing tools promise efficiency but introduce security risks and inaccuracies in production environments. Traditional pentesting tools force testers into manual workflows, limiting scalability. Resolve this tension by enabling autonomous pentesting to scale through human-governed AI execution."*

Shahar Peled CEO & Co-Founder

04

Contextualization Is Everywhere

Across virtually all sectors, vendors, customers, and domain experts are signaling the shift toward contextual security.
Alert triage and enrichment are increasingly executed through the lens of business-critical context; just-in-time privileges are being issued and revoked based on usage context and behavioral pattern recognition; and secure coding is shifting toward context-aware guidance - surfacing prescriptive fixes directly from pull requests and code reviews.
Context is no longer merely an input to security decisions; it now underpins them. Its influence is pervasive today and is poised to become even more dominant over the coming year.

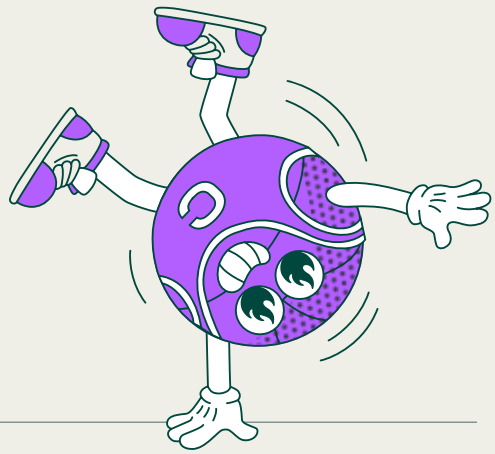

- Quote source: George Kurtz, CEO, CrowdStrike
- Quote source: Kevin Mandia, CEO & Founder, Armadin
- Quote source: Ido Geffen, CEO & Co-Founder, Novee
- Quote source: Shahar Peled, CEO & Co-Founder, Terra Security

05

The Great Security Platform Convergence

By 2025, security tooling fragmentation was becoming harder to justify: **many teams were running broad stacks of overlapping tools across endpoint, identity, cloud, network, and detection**. That sprawl added unnecessary costs and operational burden (more integrations, more policy surfaces, more triage), yet left teams with hard-to-trace visibility gaps.

At the same time, we started seeing clearer consolidation signals. Among others, Google's acquisition of Wiz suggests that leading point solutions are increasingly being treated as "platform-grade" capabilities. Despite ongoing concerns around vendor lock-in, we expect 2026 to extend this direction: **more stack rationalization and more vendor consolidation, with a stronger emphasis on fewer tools that drive prioritized remediation over noisy detection**.

This trend is amplified by the previous topic - contextualization - as the trend of putting AI & organizational context in the core of each security product, brings the different tools - and traditionally-defined cybersecurity quadrantable categories - even closer together.

Key industry moves and company launches

Moves by legacy companies (M&A/launches)	Companies acquired in the past year	
Google (Acquired Wiz)	WIZ	onum
CROWDSTRIKE (Acquired Onum)	observo.ai	chronosphere
SentinelOne (Acquired Observo AI, Prompt security)	KOI	ARMIS.
paloalto (Acquired CyberArk, Chronosphere and Koi)	Moveworks	CYBERARK
servicenow. (Acquired Armis, Moveworks and Veza)	veza	Prompt:

> *"We're seeing a trend towards more consolidation, more platformization. You cannot respond fast if you've got 70 different vendors who have different data, different logs, different APIs running."*
>
> **Nikesh Arora** Chairman & CEO

> *"What was a market littered with dozens of companies is quickly consolidating to several vendors."*
>
> **George Kurtz** CEO

06

LLM Marketing Fatigue Set in Quickly

By mid-2025, CISOs were openly dismissive of generic AI and LLM claims. Many vendors added "AI" to their messaging without fundamentally changing their products, and hype alone no longer inspired trust. CISOs demanded both proof of the advantages of the LLM through clear KPIs and a measurable way to audit, understand, and control its behavior. Vendors that couldn't provide precise answers, validation, or evidence quickly lost credibility. Adoption would no longer be driven by marketing claims - security leaders wanted results, not buzzwords.

During 2026, the term "AI agent" or "agentic LLM" is prone to suffer the same fate, as CISOs now see past the AI hype, demanding demonstrable value and explainability.

- Quote source: Nikesh Arora, Chairman & CEO, Palo Alto Networks.
- Quote source: George Kurtz, CEO, CrowdStrike

04

Distinguished Tech Pivots

The following five areas capture the main structural changes that defined 2025 and now shape 2026 planning and execution.

01

The Autonomous SOC: The Emergence of AI SOC

What happened in 2025

2025 was the year AI stopped being a toy in the SOC and started carrying real Tier-1 load. Leading vendors and early adopters put agentic systems in the alert path: they triage and cluster alerts, crush duplicate noise, enrich incidents from security data lakes using retrieval-augmented reasoning, and auto-draft response playbooks for analysts to review and approve. The conversation at the front of the market shifted from "should we add a copilot?" to "how far are we willing to let agents act on their own?". Fear of fully autonomous response, from hallucinated actions, missed or downgraded real attacks, to uncontrolled blast radius and opaque decision trails, remained a hard constraint, reinforced by new AI and cyber regulations that insist on human oversight and accountability for high-impact security actions. **Exactly where the line for true autonomy should sit is still very much unresolved, and the category will only become non-optional once there is hard evidence that SOC performance improves dramatically, without increasing the probability or impact of false negatives at scale**.

At the same time, it is hard to point to any single KPI that will unambiguously improve, as every gain in SecOps will be met by corresponding adaptations along the attackers' kill chain, including adversaries' own use of agentic AI to probe, evade, and poison automated defenses.

Vendor signals

Vendors turned autonomy into a dial. The same agents can run in "recommend only", "human approval", or "auto-execute within guardrails" modes, effectively moving SOC operators into an oversight role over policy and risk. In parallel, products started to expose more AI plumbing - model choices, lineage, and audit trails - to give security and compliance teams real visibility into what the agents do and to make their behavior auditable.

Security leader signals

Leaders adopted AI "inside the fence" first: alert triage, incident writeups, control mapping, and compliance documentation - areas where agents could be wrong without taking production down. The real question quietly shifted from 'will AI replace analysts?' to 'will AI finally let analysts do analyst work?' They also insisted on guardrails: SOC change control, red-teaming of agents, sandboxed and scoped execution, as well as actions that are signed, logged, and rollback-capable.

- Recent SOC studies show AI assistance reducing investigation time and increasing Tier-1 accuracy when embedded directly in analyst workflows. See: The agentic SOC: SecOps evolution into agentic platforms, Omdia Tech, 2025.
- AI regulations (EU AI Act; NIST AI RMF) that require meaningful human oversight for high-impact AI decisions.
- Research on agentic security systems stresses that even a small increase in missed true attacks can outweigh large productivity gains. See: LLMs in the SOC: An Empirical Study of Human-AI Collaboration in Security Operations Centres, 2025

Why it matters

Measurable productivity gains in SOCs depend on pushing more work to AI while keeping core metrics: time to detect, time to respond, dwell time, and error rates - flat or improving, and without raising the risk of bad changes. Trust rests on grounded outputs and strict guardrails on what agents can touch, but above all on uncompromising auditability: every recommendation and action is logged, explainable, and traceable back to data, model, and approver.

KPIs to watch

- Share of alerts triaged by agents, and the analyst approval / override rate on those decisions
- Time-to-first-draft for incidents (from alert to usable writeup) and the associated human-rated quality scores
- Agent-initiated actions with full audit trail and successful rollback rate (including tested and actual rollbacks)

2026 outlook - Motions & Tailwinds

We'll see a shift from assistive workflows to semi-autonomous response for low-risk changes, governed by policy and confidence thresholds. For example, automatically disabling a non-privileged account following a high-confidence login anomaly, with full auditability and rollback. The rising efficiency of attackers using AI, combined with the productivity boost from AI-powered SOCs, will drive widespread adoption, commoditization, and a shift in differentiation toward user experience and vertical flavors (e.g., "AI SOC for X industry"). In parallel, the traditional boundary between detection and cloud/ops will erode as SOC teams will gain the tooling and permissions to execute a larger share of routine operational actions.

Key industry moves and company launches

Companies launched in the past year

MATE	7AI	LEGION	Conifers	prophet
zenyard	Exaforce	Cotool	TENEX.AI	KENZO SECURITY
Bricklayer AI	fig	ANOMALI	surf.ai	

Moves by legacy companies (M&A/ launches)

Microsoft (Embedded agentic triage directly into M365 E5)	SentinelOne (Upgraded to "Athena" with agentic auto-triage and auto-investigation)	CROWDSTRIKE (Launched 7 autonomous SOC agents)
paloalto (Launched AgentiX autonomous investigation and response framework)	Google (Launched Agentic SOC agents within Google SecOps)	

Companies acquired in the past year	
wirespeed	

"A year ago, every meeting started with: "Does AI actually work in security?"
Today, it's: "How do I operationalize AI agents in my SOC?"

Lior Div CEO and Co-Founder — 7AI

"The old approach of configuring and maintaining endless playbooks doesn't scale. Attackers are already using AI to launch bigger and faster campaigns. Security teams need tools that don't just keep up but actually learn and improve continuously."

Asaf Wiener CEO & Co-Founder — MATE

"This is not about eliminating jobs. It's about ensuring an analyst doesn't have to spend time triaging and investigating alerts, because who wants to do that all day, every day? Instead, they can focus on the 4% of issues that truly matter to an organization."

Kamal Shah CEO — prophet

"What we actually want to measure is that the AI is accurate, comprehensive, and that it takes on work that is actually valuable. If it is, measuring how many equivalent analyst hours are done by the AI is a great metric to start with."

Ely Abramovitch CEO & Co-Founder — LEGION

- Quote source: Lior Div, CEO & Co-Founder, 7AI, 2025.
- Quote source: Asaf Wiener, Co-Founder & CEO, Mate, 2025.
- Quote source: Kamal Shah, CEO, Prophet Security
- Quote source: Ely Abramovitch, Co-Founder & CEO, Leigon

02

LLM Security and Governance

What happened in 2025

By 2025, LLMs were running at scale inside enterprises - dev tools, productivity suites, and SOC consoles - so security, risk, and compliance had to lock in. LLM risk management became concrete: control frameworks and checklists scored model safety, hallucination risk, and data exposure for high-impact use. **Regulators raised the bar with AI and privacy rules that expect human oversight, logging, and strong data controls around high-risk workloads**. Guardrails carried a double mandate: security (no PII spills, no prompt-injected SQL against production) and safety (keeping models inside acceptable behavior and policy in high-stakes workflows). Psychological jailbreaks, socially engineered prompts and other forms of manipulation emerged as a visible attack surface. Model and security vendors answered with hard controls - behavioral policies, input/output filtering, policy-based access, continuous red-teaming, and even AI-driven red-teaming in production to iteratively strengthen guardrails - to keep prompt injection, data leakage, and other LLM failure modes inside a visible, auditable fence, with some programs also experimenting with behavioral red-teaming to probe manipulative, multi-turn attacks.

Vendor signals

Platforms now ship with guardrails that enforce security and safety policies, backed by explainability, observability, and real-time monitoring dashboards. Compliance reporting increasingly maps directly to ISO/IEC 42001 controls rather than ad-hoc "responsible AI" checklists. Model risk management is no longer a parallel track; it is wired into mainstream MLOps pipelines as a first-class stage for validation, approvals, and ongoing monitoring. Psychological maneuvering has become a problem, and behavioral red-teaming started showing up.

Security leader signals

By 2025, many organisations effectively ran two AI policies: the formal one in the handbook, and the shadow policy that actually lived in Slack. Security leaders refused to accept that reality: they now demand governed LLM behavior with hard guardrails against toxic, biased and hallucinated outputs, plus audit trails and human review on high-impact decisions. At the same time, they expect LLMs to snap into their existing Zero Trust and data security stack, with strict data leakage controls, purpose-based and least-privilege access, continuous monitoring for adversarial misuse, and formal AI governance over both sanctioned and shadow AI.

Why it matters

LLMs introduce risks - hallucinations, bias, prompt injection - that traditional security controls miss. Visibility, Monitoring, Governance, and Auditability must be first-class concerns across the entire AI lifecycle, from data collection and training through deployment, inference, and eventual retirement.

KPIs to watch

- Hallucination rate in production
- Guardrail intervention frequency and accuracy
- Time to detect and remediate adversarial inputs

2026 outlook - Motions & Tailwinds

LLM security and safety will be pulled into mainstream AppSec: teams will bake in guardrails by default, treat model attestations as compliance evidence, and push behavioral monitoring earlier into the SDLC.

Key industry moves and company launches

Companies launched in the past year

solidcore.ai | MEIBEL | Virtue AI | Aiceberg | Aurascape | Runlayer | Pillar

Moves by legacy companies (M&A/ launches)

paloalto: Acquired Protect AI	snyk: Launched AI Trust Platform, acquired Invariant Labs	Microsoft: Announced Prompt Shields, updates of Azure Foundery	f5: (Acquired CalypsoAI)
aws: Launched Amazon Bedrock Guardrails and Automated Reasoning checks	ALICE: Launched AI safety and security platform	netskope: Launched GenAI DLP and Shadow AI visibility	CROWDSTRIKE: (Acquired Pangea)

Companies acquired in the past year

PROTECT AI | pangea | LAKERA | splx AI | Calypso AI | Prompt: | aim security | Invariantlabs

> *"Enterprises now operate in a world where anyone who knows how to talk knows how to hack."*
>
> **David Haber** Founder and CEO — LAKERA

> *"AI is here to stay, and enterprises must implement strategies to monitor and protect AI use. Traditional security offerings were not designed for the ways AI applications operate."*
>
> **Moinul Khan** Co-Founder & CEO — Aurascape

> *"Everyone talks about AI, but AI is really only as useful as the tools and the resources it has access to."*
>
> **Andrew Berman** CEO — Runlayer

> *"Enterprises we're working with have 50 to 200 LLM applications today. That number could double, and then quadruple, in just the next few years."*
>
> **Eric Chiu** CO-Founder and CEO — solidcore.ai

03

Decoupling Data from SIEM: The Security Data Layer

What happened in 2025

Security data finally started breaking out of the SIEM jail: Organisations implemented vendor-agnostic pipelines as a control plane, then pointed them at SIEM, XDR and open lakehouses so detections could run wherever the data actually lives. Procurement stopped buying "platforms by logo" and started buying "signals by dollar," with cost per event and cost per investigated alert becoming hard gates in renewals.

Vendor signals

Vendors promoted bring-your-own-lake ingestion, late-binding schemas, replay on cheap storage and tiered economics. They pushed source-agnostic, federated detections over a shared telemetry fabric, offering cross-product correlation and summarisation that run directly on whatever lake or SIEM holds the data.

Security leader signals

Platform teams rationalised duplicate ingestion, normalised data ownership and demanded transparent price performance and workload portability.

Why it matters

Signal density per dollar is increasingly the measure of detection sustainability. **Decoupling storage, compute and analytics unlocks choice, flexibility and optimisation**.

- Quote source: David Haber, Founder & CEO, Lakera.
- Quote source: Aurascape launch announcement (Business Wire, 2025).
- Quote source: Andrew Berman, CEO, Runlayer
- Quote source: Eric Chiu, co-founder and CEO, Solidcore.ai

KPIs to watch

- Cost per investigated alert and cost per retained terabyte
- Coverage of priority telemetry sources and replay SLA

2026 outlook - Motions & Tailwinds

"Query once, detect anywhere" will solidify into the default pattern: a single detection definition fanning out across SIEM, XDR and lake engines, running as close as possible to where the data already sits. Data retention will be policy-driven and explicitly tied to business risk, with hot, high-value telemetry kept close and expensive, and long-tail data pushed to cheap tiers but still replayable on demand.

Key industry moves and company launches

Companies launched in the past year					
vega	Artemis	DATABAHN	Beacon	Sawmills	ziggiz
Moves by legacy companies (M&A/ launches)					
paloalto (Acquired Chronosphere)	CISCO (Launched Cisco Data Fabric)	SentinelOne (Acquired Observo AI)	CROWDSTRIKE (Acquired Onum)		
Google (Added Bindplane-powered data processing pipelines for filter before ingestion)					
Companies acquired in the past year					
onum	observo.ai	datable	chronosphere		

> *"Enterprises aren't just overwhelmed by data volume; they're being outpaced by its complexity"*
>
> **Nanda Santhana** Co-Founder and CEO DATABAHN

- Analysts note a shift toward cost-per-event and cost-per-investigation economics in modern cloud and security programs.See: 2024 State of the Cybersecurity Market: $14B, Key Trends & Data, Returnonsecurity, 2025.

> *"Security is, at its heart, a data problem, and legacy, rules-based data pipeline platforms simply weren't built for today's ever-growing attack surface and data-rich security operations."*

Tomer Weingarten Co-Founder and CEO — SentinelOne

> *"The current operating model of the SIEM - the dominant technology in this domain for the last two decades - is not only 'crazy expensive,' but is also increasingly causing AI-native security operations to fail."*

Shay Sandler Co-Founder and CEO — vega

> *"I think we are discussing AI too much and losing the context. AI is changing our lives, but perhaps not yet. We want to show the market that data is the only place where all tools, all attacks, and everything are together."*

Pedro Castillo Co-Founder & former CEO — onum

> *"I hear fewer questions about "what platform should we choose?" and more about "how do we manage our data so each tool gets what it needs to do its job well?"*

Gal Tal-Hochberg Co-Founder & CEO — Beacon

04

Identity as the New Perimeter Attack Surface: From Users to NHIs & Agents

What happened in 2025

Identity became the first perimeter of security across cloud and SaaS - who or what you are mattered more than which network you sat on. Non-human identities exploded - API keys, service accounts, tokens and agents outnumber human users, often with tens of non-human identities per human user, turning poorly kept secrets into one of the steepest, fastest-growing risk curves in the stack. At the same time, heavy-friction IAM and PAM workflows are still pushing developers to bypass controls, fuelling shadow access and unmanaged NHIs that quietly escaped central governance. **Just-in-time and just-enough privilege finally moved from slideware into mainstream programs, as Zero Standing Privileges are slowly becoming the expected pattern for admins, developers and high-risk NHIs**.

Vendor signals

Vendors combined identity threat detection, entitlement visibility and automated access brokering. They enhanced graph based context (HR, device, workload) to score access risk.

Security leader signals

Leaders focused on toxic privilege combinations and orphaned rights that drive blast radius in a breach. Adoption was guided by measurable reductions in excessive entitlements and high risk access paths.

Why it matters

When breach impact correlates with privilege sprawl, identity must serve as the perimeter but only if entitlement right sizing and machine identity lifecycle control are continuous.

KPIs to watch

- Reduction in standing admin privileges and high risk paths
- Mean time to deprovision machine identities
- Percentage of access granted via JIT (with time bounds)

2026 outlook - Motions & Tailwinds

Identity is steadily consolidating onto unified platforms that cover workforce, customer, partner, and machine identities across on-prem and cloud, instead of living in separate stacks. At the same time, IGA and PAM will converge into a single control plane, so the same policies that govern joiners/movers/leavers also drive just-in-time privilege and approvals on production changes. Identity risk scores will plug directly into CI/CD, deployment, and change-control gates, turning "who is this, and how risky are they?" into a non-negotiable release criterion.

- See: Machine Identities Outnumber Humans Increasing Risk Seven-Fold, infosecurity magazine, 2025.
- Consulting guidance increasingly treats just-in-time access and zero standing privileges as the default for admins and high-risk workloads. See: The agentic reality check: Preparing for a silicon-based workforce, Deloitte Insights, 2025.

Key industry moves and company launches

Companies evolved in the past year

VENICE	cyki	Fabrix	Orchid	aizome.	Silverfort
LINX	OASIS	outtake	Astrix	Abnormal	imper.ai
AUTHMIND Agent monitoring & runtime	HushSecurity (Access Governance)	(ConductorOne)	token Access Governance		

Moves by legacy companies (M&A/ launches)

okta (Acquired Axiom)	paloalto (Acquired CyberArk)	Microsoft (Launched Entra Agent ID)	CROWDSTRIKE (Launched Falcon Privileged Access)
IBM (Acquired HashiCorp)	GitGuardian (Launched its NHI Governance product)	SailPoint (Launched NHI governance features)	Saviynt (Expanded its platform to cover NHI + AI agents)

Companies acquired in the past year

Axiom	CYBERARK	Otterize	zilla SECURITY	veza	sgnl
STACK IDENTITY	strongdm	Vault	CYBERARK Access Governance		

> *"Identity is under relentless attack, and adversaries are going straight for the keys to the kingdom — privileged access. From social engineering to sophisticated insider abuse, they're escalating privileges to access the most sensitive systems and data."*
>
> **Michael Sentonas** President — CROWDSTRIKE

> *"We're at a pivotal moment in identity security. The unseen dark matter of identity is overtaking what organizations can manage or even see. It's no longer about control - it's about context."*
>
> **Roy Katmor** CEO & Co-Founder — Orchid

> *"For years, companies assumed the root of identity security was making access as convenient as possible. But what has changed is the scale and dynamism of modern environments. Humans can manage things manually, but organizations operating at today's speed, especially with AI agents, need systems that can handle constant change."*
>
> **Rotem Lurie** CEO & Co-Founder — VENICE

05

AI coding and Application Security

What happened in 2025

Application security continued to move upstream, as adversaries increasingly targeted the environments where software is made, not just where it runs. Developer endpoints, IDEs, CI runners, package ecosystems, and build credentials became the soft underbelly of otherwise well-hardened production stacks. Campaigns like NX and Shai Hulud reflected a broader shift in adversary strategy: **rather than attacking hardened production systems directly, attackers targeted the systems that created them**.

In parallel, AI-written code became the default operating mode for many engineering teams, expanding the developer plane. The volume of AI-generated code paths and dependency decisions now outpaces already-stretched security checks. As a result, teams tuned out generic "AI security" messaging and demanded auditable evidence at decision time: what was generated, what shaped it, what data it touched, and whether it cleared a defensible ship bar.

This points to a new generation of AppSec: continuous, context rich, policy-driven tooling for AI-assisted delivery, with automated provenance and verification across the creation layer.

- Quote source: Michael Sentonas, President, CrowdStrike.
- Quote source: Roy Katmor, CEO, Orchid Security
- Quote source: Rotem Lurie, Co-Founder & CEO, Venice

Vendor signals

Vendors started bundling developer security into a single story across AppSec, code scanning, secrets scanning, and supply chain controls. Key patterns included:

- Security checks earlier in the workflow, inside IDEs and pull requests
- Guardrails for AI coding, like policy checks, safe prompts, and blocked risky patterns
- Better visibility into open source and build dependencies, including provenance and signing
- Contextualised tooling - for posture and remediation alike
- Detection for repo access abuse, token misuse, and suspicious CI activity

Security leader signals

Security leaders wanted to keep developer velocity high while making AI-driven changes reviewable and enforceable. Priorities included clear ownership and defined scope of AI coding tools used across repos and environments, hard controls on secrets and tokens in repos and CI, mandatory security gates embedded throughout the SDLC, and evidence that AI use does not bypass review, testing, or approval standards.

Why it matters

If AI can write more code, teams will ship more code. That raises the odds of vulnerabilities, insecure defaults, and dependency risk, unless controls scale with output. **The builder is now part of the perimeter. If developer identities, endpoints, repos, and pipelines are compromised, production will fall downstream**.

KPIs to watch

- Mean time to rotate or revoke exposed secrets and tokens found in code or CI logs
- Percentage of critical repos covered by branch protection, signed commits, and required reviews
- Mean time to patch vulnerable dependencies after disclosure
- Percentage of PRs that pass security gates before merge (and how often gates are bypassed)
- Rate of high severity findings introduced per release (not just total findings)

2026 outlook - Motions & Tailwinds

AppSec will move from “scan and report” to “gate and verify.” AI coding will push policy into authoring, build, and merge before code is built and merged -not after the fact. Expect security controls to plug directly into IDEs, PRs, CI/CD, and change management so teams can answer, in real time:

- Who wrote or generated this change?
- What was the source of the code and dependencies?
- What risk do we accept, and why?

In 2026, the goal is not to stop AI written code. The goal is to secure the builder. Secure developer identities, endpoints, and secrets, and you reduce downstream production risk.

- A supply chain attack (August 2025) in which attackers hijacked the Nx build platform, a developer tool with 5M+ weekly downloads, to steal developer credentials, tokens, and keys at scale.
- A self-replicating npm supply chain worm that compromised 500+ packages

Key Industry Moves and Company Launches

Companies evolved in the past year

Archipelo | Clover | symbiotic | DamSecure | echo | sola

Moves by legacy companies (M&A/ launches)

Checkmarx
(Acquired Tromzo)

ANTHROP\C
(Claude Code Security)

snyk
(Rebranded its platform as the "AI Security Fabric")

Google
(Wiz launched Wiz Code)

GitHub
Launched Copilot coding agent with built-in security validation

Companies acquired in the past year

TROMZO | Mayhem | Trag | XEOL by herodevs | DAZZ.

> *"In a world where AI is transforming software development, the biggest security risk isn't just in the code - it's in how the code is written."*

Matthew Wise Co-Founder and CEO — Archipelo

> *"The problem is clear: AI has pushed engineering velocity far beyond what reactive security tools were built to handle."*

Alon Kollmann Co-Founder and CEO — Clover

> *"Making security a positive experience for developers is key to growing their cyber judgement and knowledge. By integrating AI-powered training into their workflow and using their current work as the reference point, developers learn in a way that's impactful, helping them better understand and resolve security vulnerabilities without disrupting productivity."*

Edouard Viot CTO and co-founder — symbiotic

- Quote source: Matthew Wise, CEO & Co-founder of Archipelo
- Quote source: Alon Kollmann, CEO & Co-Founder, Clover
- Quote source: Edouard Viot, CTO, Symbiotic Security

05

Newly emerging tailwinds

While there were no meaningful public motions to reference in these categories, having seen many teams working in stealth on solutions to these and given the strong technological shift that enables / require them - we make an educated guess that these categories would be prominent in next year's report:

01

AI DLP

It is the next generation of DLP: instead of relying mostly on fixed rules like "block credit card numbers" or "detect files labeled confidential," it uses AI models plus business context to understand what sensitive data is, where it is going, why the user is sending it, and whether the action is risky.

02

Agentic Runtime Security

Given the inability to pre-define guardrails for agents, new tools will "run" next to agents and logically & contextually make sure they do as intended and reported, stopping them when they don't. These dynamic guardrails - which we call in this report (coining a phrase!) "runners" - are an inevitable development in cybersecurity.

03

Unified Agentic Defense Platforms (UADPs)

As autonomous AI agents scale, the traditional, siloed security stack is beginning to break down. In 2026, we will see the emergence of Unified Agentic Defense Platforms (UADPs): This is a new architecture that converges data security, identity governance, and AI security into a single control plane. Rather than managing DSPM, DLP, AI security, and identity as separate tools. UADPs address the core challenge of the agentic era: governing the intersection of who or what is acting, what data is being accessed, and why the action is occurring. Static security rules will increasingly give way to real-time, behavior-driven defenses that evaluate agent intent and intervene at machine speed. The result will be significant consolidation across the security stack, as standalone tools collapse into a unified platform designed to secure autonomous systems operating across enterprise environments.

- Coining a new phrase: "Runners" also referred to as Agentic Runtime Security) are a new class of tools designed to provide dynamic security for autonomous AI agents. They operate alongside agents in real-time, functioning as dynamic guardrails. Their purpose is to logically and contextually verify that an agent is operating as intended, auditing all actions for compliance and traceability in realtime, and immediately stopping or containing the agent if its behavior deviates from its designated purpose.

06

Cross Cutting Recommendations For CISOs

These recommendations reflect lessons from 2025 as AI adoption accelerated across security and engineering environments. As AI capabilities expand, organizations must adopt them in ways that remain measurable, controllable, and auditable. Risk is reduced when clear operational boundaries are defined, when the full lifecycle from development through production is hardened, and when real incidents are systematically converted into stronger preventive controls. Sustained executive support depends on translating resilience into business-relevant outcomes that leadership can track and fund.

DON'T

Roll out AI everywhere all at once

Start with a narrow, high-frequency security use case where outcomes are measurable and the blast radius is controlled (e.g., phishing triage, alert summarization, investigation drafts). Running AI in a contained domain allows teams to validate reliability, measure productivity gains, and understand operational risks before expanding its role. Once the system consistently improves triage speed, investigation quality, or analyst workload, the scope can gradually expand to additional workflows with higher operational impact.

DO

Build guardrails before you build autonomy

Require guardrails across the entire AI lifecycle, from development to production, including strict data policies, versioned prompts and tools, approval workflows for sensitive actions, and full logging of AI decisions and evidence. AI systems operating in security environments must function under controlled autonomy, where their behavior can be audited, explained, and constrained.

DON'T

Collect security data you can't afford to investigate

Don't treat telemetry ingestion as a free resource. Uncontrolled data pipelines create noise and drive up cost without improving detection quality. Instead, treat each telemetry source as a measurable investment: track the ingestion, storage, and compute costs alongside the detections and investigations it enables.

- Quote source: Nanda Santhana, Co-Founder & CEO, DataBahn.
- Quote source: Tomer Weingarten, CEO & Co-Founder, SentinelOne.
- Quote source: Shay Sandler, Co-Founder & CEO, Vega
- Quote source: Pedro Castillo, Founder & CEO, Onum
- Quote source: Gal Tal-Hochberg, Co-Founder & CEO, Beacon

DO

Make identity risk visible where work happens

Ensure that identity risk is visible inside the operational environments where access decisions are actually made, not only within security dashboards. Instead of relying solely on centralized IAM or governance tools, make sure there is an integration of identity risk signals into developer and platform workflows. This allows engineers and platform teams to see the security implications of privileges, tokens, and access paths at the moment they create or modify them.

DON'T

Ignore non-human identities

Don't assume API tokens, service accounts, and machine identities are low-risk or temporary. In modern cloud and SaaS environments, these identities often hold persistent privileges and can create powerful attack paths if left unmanaged. CISOs should ensure that non-human identities are governed with the same rigor as human credentials, including automated discovery, rotation policies, and lifecycle controls across the SaaS and API ecosystem.

DO

Quantify cyber resilience in business terms

Translate cyber resilience into clear, measurable outcomes that leadership and the board can understand and track over time. While no universally accepted KPI exists for resilience, CISOs should prioritize indicators that demonstrate tangible improvements in operational performance and risk reduction. The specific metric matters less than its transparency, repeatability, and connection to business impact. Frame progress through trend lines and measurable changes in exposure, rather than relying solely on abstract technical metrics.

DON'T

Try to fight Shadow AI

Employees will adopt AI tools whether security teams approve them or not. Instead of attempting to block usage entirely, CISOs should focus on discovering, monitoring, and governing AI adoption before it evolves into unmanaged risk.

Closing Perspective

Security improvements did not come from adding more tools this year, it came from making them work together. The programs that improved fastest were those that connected signals, shared context across controls, and enabled teams to act quickly on what they saw.

AI did not change that principle. It accelerated everything - attacker experimentation, alert volume, and response timelines - but the difference between strong programs and weak ones remained coordination. Systems that share context and support decisive action outperform stacks that simply accumulate tools and signals.

With the goal being to reduce time between risk appearing and controls taking effect, faster detection and containment, quicker restoration of services, and a smaller exposure window for revenue-critical systems became both the method, and the metric.