

H1 2025 Cybersecurity Report

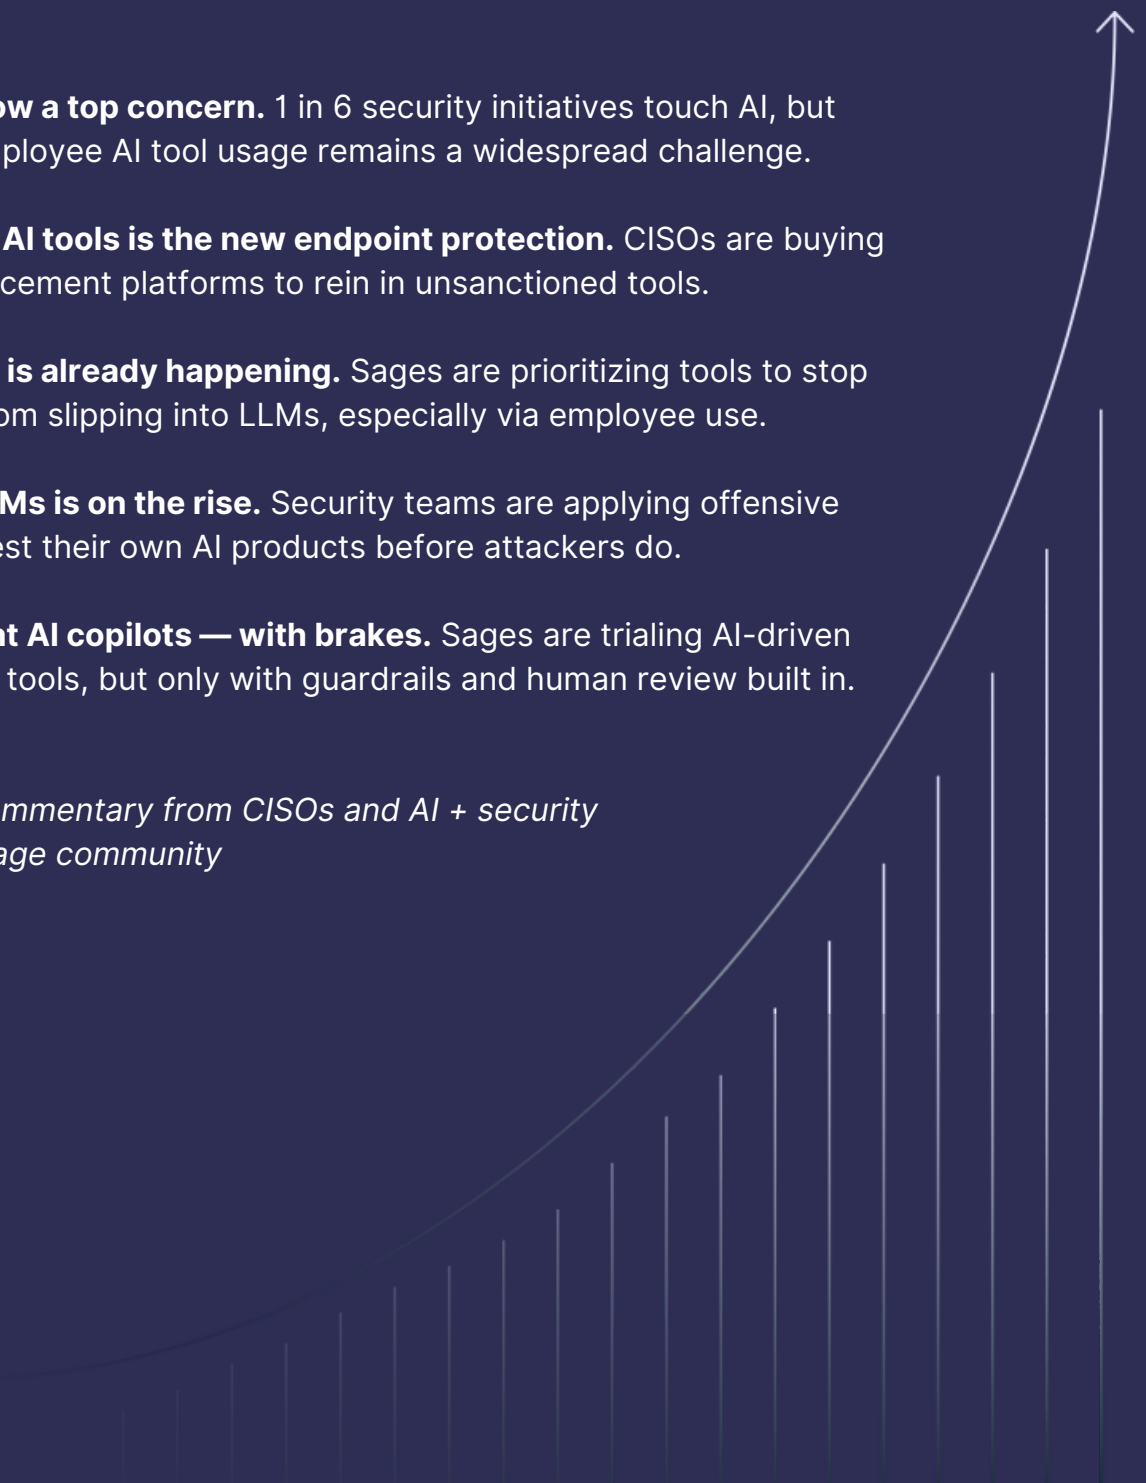
Published September 8, 2025

The Security Shift: AI Trends From Verified Leaders

Highlights

- **Shadow AI is now a top concern.** 1 in 6 security initiatives touch AI, but visibility into employee AI tool usage remains a widespread challenge.
- **Blocking rogue AI tools is the new endpoint protection.** CISOs are buying visibility + enforcement platforms to rein in unsanctioned tools.
- **AI data leakage is already happening.** Sages are prioritizing tools to stop sensitive info from slipping into LLMs, especially via employee use.
- **Red teaming LLMs is on the rise.** Security teams are applying offensive techniques to test their own AI products before attackers do.
- **SOC teams want AI copilots — with brakes.** Sages are trialing AI-driven SIEM and SOAR tools, but only with guardrails and human review built in.

Featuring expert commentary from CISOs and AI + security executives in the Sage community



H1 2025 Cybersecurity Report


At Sagetap, we track the initiatives launched by Sages (enterprise security leaders) — active projects that reflect where teams are focusing time, budget, and attention. This H1 report analyzes initiatives from January through June to surface the most important trends shaping security strategy today and reveal where the industry is headed.

In the first half of 2025, 1 in 6 security initiatives involved artificial intelligence, showing that AI is the new operating environment rather than just an emerging trend. Leaders on Sagetap are thinking about how to use AI and how to secure against it, from internal misuse and shadow tool sprawl to model poisoning and hallucination in the SOC.


We've organized the insights into four key themes:

- Securing AI Usage: Governance & Compliance
- Securing AI Workflows: Data Protection & Privacy
- Securing AI in Production: Model Lifecycle Security
- AI for Security Operations: Threat Detection & Response

Security Leaders' Top Priorities in H1



In H1, cybersecurity executives referenced AI in a growing share of initiatives, from 9% in January to a peak of 22% in June, and averaging 17% for the first half of 2025.




Note: Pie charts throughout this report reflect weighted averages based on initiative volume per month, rather than a uniform monthly average.

Securing AI Usage:

Governance & Compliance


Percentage of AI Security Initiatives Focused on Governance & Compliance



Note: January data is not shown because relevant initiatives for this theme had not yet been created following the initiative feature's release.

Sages are prioritizing control over AI tools used across their orgs. The goal: identify both approved and unapproved usage, monitor for risk, and enforce policies. Automated detection and response are key to closing gaps before they result in compliance issues.

AI Security Initiatives Focused on Governance & Compliance vs. Other Security Topics



What Sages Are Looking For

Visibility & Control Over Employee AI Usage

- Security teams want visibility into both sanctioned and unsanctioned AI tools. Monitoring usage helps identify risk early and shape internal policies.
- Leaders are implementing controls that track how data flows through AI systems, ensuring that policy violations and sensitive data use are flagged in real time.
- Teams are beginning to use guardrails to block unauthorized AI access at the point of use. This limits exposure while improving adherence to internal standards.
- Compliance teams need audit-ready reporting for all AI-related activity. These insights help prove controls are working and reduce manual effort during reviews.


Real-World Sage Perspective

"Some organizations get caught up in the 'magic' of AI capabilities but overlook the importance of clean, accessible, real-time data and proper access controls. Without that data infrastructure, even the most sophisticated AI becomes expensive guesswork. The hype around AI often overshadows the unglamorous but critical work of data governance, security frameworks, and compliance processes that actually enable AI to deliver meaningful business value safely."

Securing AI Workflows:


Data Protection & Privacy

Percentage of AI Security Initiatives Focused on Data Protection & Privacy



Sages are working to understand which AI tools employees are using — including unsanctioned ones — while also consolidating usage onto secure, governed platforms. The goal is to prevent sensitive data from leaking into AI tools and ensure privacy and control.

AI Security Initiatives Focused on Data Protection & Privacy vs. Other Security Topics



What Sages Are Looking For

Preventing Data Leakage Across AI Workflows


- Security leaders want to block PII, PHI, and IP from entering unapproved AI tools, which helps prevent privacy violations and unintentional data loss.
- Monitoring tools are being deployed to detect and stop sensitive inputs in real time and to support compliance and reduce downstream risk.
- Enterprise-ready AI platforms with built-in protections are in high demand. Sages want encryption, access controls, and reporting capabilities built in.
- Teams are consolidating AI usage onto secure platforms with policy guardrails. This transition helps ensure safer adoption and stronger oversight.

Real-World Sage Perspective

“AI doesn't automatically eliminate human error — it often just shifts the types of errors we need to manage, especially with hallucinations in complex reasoning tasks. The set-it-and-forget-it mentality around data protection and privacy is where organizations get burned. Using AI properly means understanding its limitations, implementing proper data governance, and maintaining human oversight where reasoning and accuracy are critical. Success isn't about replacing humans entirely, but creating secure workflows where AI amplifies human judgment.”


Securing AI in Production: Model Lifecycle Security

Percentage of AI Security Initiatives Focused on Model Lifecycle Security



As organizations bring LLMs into production, security is extending to the full model lifecycle. Sages are focused on protecting training data, testing model behavior, and monitoring for abuse in production. Lifecycle coverage is becoming a key focus of AI risk management.

AI Security Initiatives Focused on Model Lifecycle Security vs. Other Security Topics



What Sages Are Looking For

End-to-End AI Security with Targeted Threat Defense


- Sages want to secure models from training through runtime. That includes defending against poisoning, prompt injection, and other emerging threats.
- Security teams are investing in tools to scan models for vulnerabilities. Early detection reduces the risk of production-time misuse or unexpected behaviors.
- Teams are exploring automated red teaming to test model resilience, simulating attacks and surface weaknesses before models go live.
- Runtime monitoring is a must for deployed models. Sages need continuous visibility into how their LLMs are being used — and misused — in the wild.

Real-World Sage Perspective

“Securing AI models in production is far more complex than traditional data security because we’re now dealing with model drift, adversarial attacks, and constantly evolving regulations on top of existing security challenges. The regulatory burden has multiplied with new frameworks, bias testing, and explainability mandates that didn’t exist before. What was once straightforward data protection is now a multi-layered effort to monitor model performance, manage version control, ensure compliance across jurisdictions, and keep up with new AI-specific regulations that seem to emerge every few months.”


AI for Security Operations: Threat Detection & Response

Percentage of Threat Detection & Response Initiatives Involving AI



Some teams are exploring AI to boost speed, precision, and efficiency in threat detection and response. Sages are adopting copilots to triage alerts, surface anomalies, and automate action. But trust and control remain essential (AI must accelerate the right decisions).

Percentage of Threat Detection & Response Initiatives Involving AI



What Sages Are Looking For

Smarter, Faster Threat Detection with Human Oversight

- SIEM copilots that use AI to triage alerts are in demand. These tools help analysts reduce noise, spot key issues, and respond faster.
- SOC teams want agentic tools that handle low-level threats, but they insist on human-in-the-loop approval for high-risk or ambiguous incidents.
- SOAR platforms powered by AI are being adopted for automated response. The goal is faster coordination across tools without losing oversight.
- Explainability is essential for any AI in the SOC. Sages want clear reasoning behind every action so they can trust, verify, and fine-tune results.

Real-World Sage Perspectives

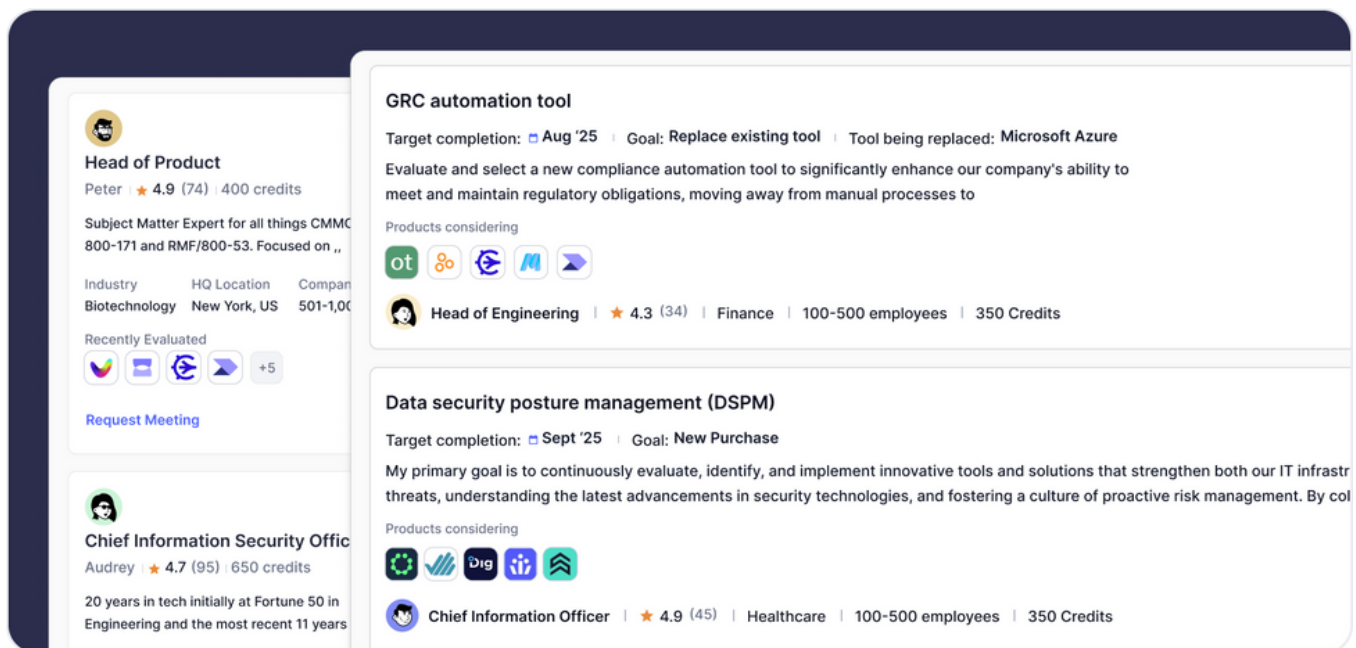
"AI in SecOps can do more than accelerate threat detection — it can reduce analyst burnout and turnover. However, we CISOs must evaluate these tools against our costs to ensure they deliver clear, measurable value in potential cost avoidance, faster and more accurate detections, and shorter response times."

"As CISOs, we need a deep understanding of AI security tools to proactively manage risk, ensure model integrity, and lead with confidence in an increasingly automated threat landscape."

This report provides a data-driven perspective on how security leaders are approaching the rise of AI in the enterprise. It highlights where AI is showing up most often — from internal usage to production systems to detection workflows — and what priorities are emerging in response. Use these insights to guide your security roadmap, assess where risks are evolving fastest, and evaluate technologies that align with how peers are adapting in real time.

Join Sagetap for personalized product recommendations that align with your team's top initiatives.

[Explore tailored solutions for free →](#)



The screenshot displays the Sagetap platform interface. On the left, there are two user profile cards. The top card is for 'Head of Product' Peter, with a 4.9 rating (74 reviews) and 400 credits. His profile mentions he is a Subject Matter Expert for all things CMMC 800-171 and RMF/800-53, focused on Biotechnology, HQ Location New York, US, and Company size 501-1,000. Below his profile are icons for 'Recently Evaluated' products and a 'Request Meeting' button. The bottom card is for 'Chief Information Security Officer' Audrey, with a 4.7 rating (95 reviews) and 650 credits. Her profile mentions 20 years in tech initially at Fortune 50 in Engineering and the most recent 11 years. On the right, there are two project cards. The top card is for 'GRC automation tool' with a target completion of Aug '25, goal of 'Replace existing tool', and tool being replaced 'Microsoft Azure'. The project description is 'Evaluate and select a new compliance automation tool to significantly enhance our company's ability to meet and maintain regulatory obligations, moving away from manual processes to'. Below the description are icons for 'Products considering' and a card for 'Head of Engineering' with a 4.3 rating (34 reviews) and 350 credits. The bottom card is for 'Data security posture management (DSPM)' with a target completion of Sept '25 and goal of 'New Purchase'. The project description is 'My primary goal is to continuously evaluate, identify, and implement innovative tools and solutions that strengthen both our IT infrastr threats, understanding the latest advancements in security technologies, and fostering a culture of proactive risk management. By col'. Below the description are icons for 'Products considering' and a card for 'Chief Information Officer' with a 4.9 rating (45 reviews) and 350 credits.



Showcase Your Expertise to Peers

Build detailed user profiles that highlight your expertise, ongoing projects, and contributions.



Engage with Trusted Anonymity

Meet anonymously until you're ready, maintaining full control over every engagement.



Learn from Peer-Driven Insights

Gain insight from real practitioners who've implemented the tools you're evaluating.



Accelerate Technology Adoption

Monitor product performance and peer outcomes at every stage of evaluation.



Manage Your Evaluation Process

Document your requirements, track evaluation recordings, and score each solution.



See the Gaps, Shape the Future

Discover personalized product recommendations that align with your enterprise goals.