S&P Global
Market Intelligence
451 Research
Business Impact Brief

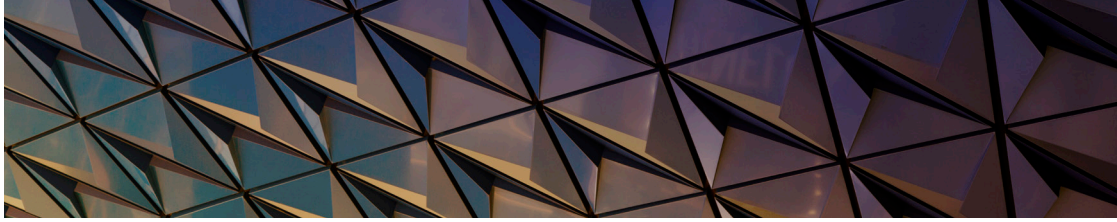

Advanced AI and data solutions hold potential for increasingly complex security and business use cases

The Take

Organizations face unprecedented security challenges as attack surfaces expand and adversaries adopt AI-driven tactics on a larger scale and with greater sophistication. 451 Research data shows that, on average, security teams are unable to investigate 45% of security analytics alerts each day, and 18% of organizations miss more than 75% of alerts. While advanced generative and agentic AI solutions show promise in easing this burden, they depend on universal, low-friction access to large, disparate enterprise datasets.

Figure 1: SecOps teams continue to struggle with alert overload

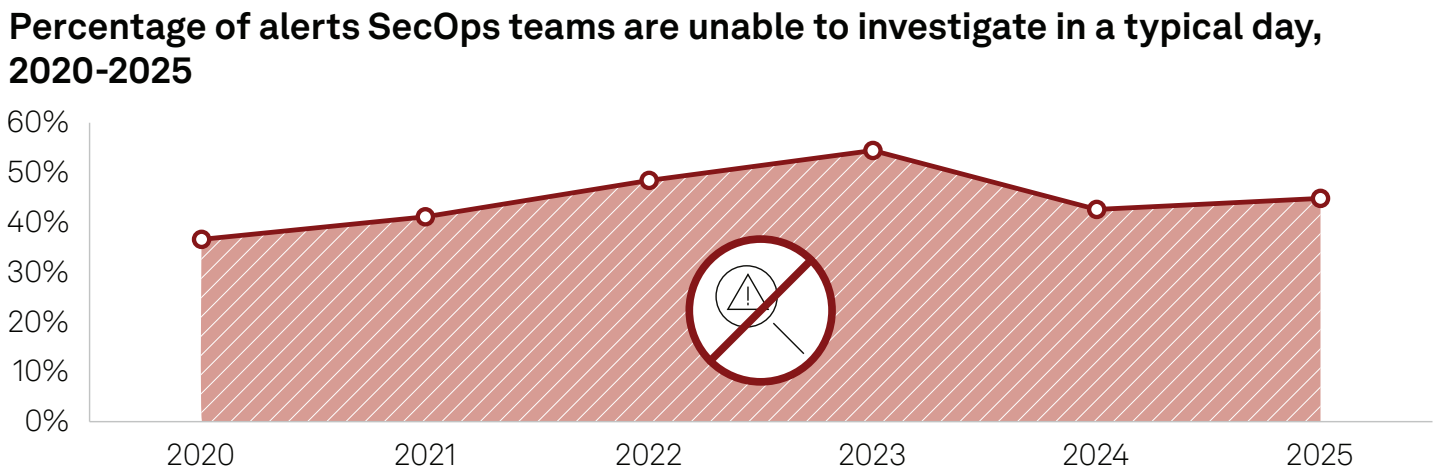

Q. What percentage of SIEM/security analytics alerts are you unable to investigate in a typical day?
Base: All respondents.
Source: 451 Research's Voice of the Enterprise: Information Security surveys: Vendor Evaluations 2020 (n=162) and 2021 (n=240); Security Operations 2022 (n=181); Security Analytics & SecOps 2023 (n=190); SecOps 2024 (n=243) and 2025 (n=313).

A core driver of security operations center "alert fatigue" is the inability of traditional analytics platforms to effectively detect issues while scaling to accommodate ever-expanding data sources and volumes. Without a robust data foundation to manage and make sense of this data, teams are overwhelmed by alert volumes and miss risky incidents, struggling to detect, analyze, prioritize and respond. At the root, this is a data problem: Analysts cannot manage what they cannot see.

Recognizing this, many organizations are pursuing a federated data foundation — often called a data fabric — that provides a single, governed source of truth without forcing data movement or transformation. Once universal data access is in place, AI analysis layers can be applied to unlock new security and broader business use cases. The goal is to offload manual, repetitive work — such as alert analysis and triage, preliminary response planning and remediation workflow creation — so analysts can focus on high-risk, high-impact issues while AI handles the drudgery. Over time, more sophisticated agents will take on increasingly complex tasks as they learn, improve and earn trust.

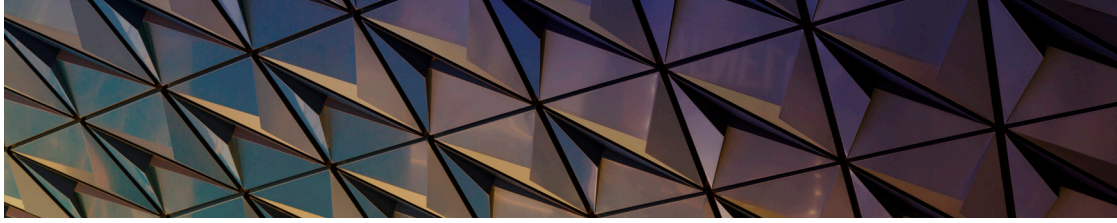

Business impact

Threat detection remains a stubborn challenge for SOC teams, but it is a natural fit for AI. Applying advanced AI to the ingestion and correlation of growing volumes of security data could ease analyst burden in multiple ways, especially when analysis can be performed without moving or transforming data. With data readily accessible, AI-driven detection tools can continuously process and correlate high volumes of signals and threat intelligence from diverse sources, enrich alerts with context and escalate the most critical threats for human review. That includes addressing missed events by automatically processing incoming alerts in real time, filtering false positives and triaging genuine alerts for deeper analysis. Improved data hygiene and availability, combined with AI-enhanced analytics, can reduce alert volumes by improving accuracy and curbing false positives.

AI-enhanced incident triage and response analysis solutions are already in the market, and agent-driven remediation is gaining traction. The first widely adopted GenAI application in security — chatbot-style AI agents introduced about two years ago — has delivered value by providing natural language query, investigation and suggested remediation capabilities to analysts at all experience levels. The next wave of agentic AI is automating routine tasks such as phishing analysis, script and code review, compliance reporting, and boosting orchestration and automation systems (hyperautomation).

Human-in-the-loop will remain standard practice for the foreseeable future. Despite rapid improvement, AI systems require human oversight for advanced threat analysis, risk balancing and remediation approval. An AI agent might recommend the fastest path to stop a threat, but that path may entail unacceptable business impact; a more nuanced approach may reduce risk with minimal disruption. For now, humans will choose and approve the approach, with AI executing the plan. Trust in full autonomy is not yet sufficient, but it will improve as systems continue to learn and become more reliable. Like junior analysts, AI agents will initially focus on basic tasks; as accuracy improves and agents collaborate, more advanced use cases can be automated.

AI-driven analytics will enable newly accessible use cases. GenAI systems can rapidly process both structured data and nontraditional sources such as documents, audio, video and observability data, creating new opportunities in security and beyond. For example, these systems can quickly assess the risk of a suspicious document, URL, email or executable, then correlate multiple data points to determine whether context matches a new or ongoing incident and prioritize accordingly.

Looking ahead

Adversaries are increasingly using advanced AI to accelerate and scale attacks, with well-resourced groups targeting multiple victims simultaneously. A recent Anthropic report highlighted how actors had manipulated Claude Code to perform reconnaissance, find vulnerabilities, exploit resources, move laterally, harvest credentials, analyze data and exfiltrate information with minimal human involvement. To counter this new generation of AI-powered attacks, AI-enabled defenses will be essential, meeting machine-speed offense with machine-speed defense. Humans will remain in the loop, but AI-automated security is on its way to becoming table stakes.

Success will hinge on the ability to process ever-increasing volumes of security data in real time, link signals through context and prioritize incidents by risk. Once incidents are identified, automated analysis and triage should free analysts to focus on the highest-impact issues, select the most appropriate remediation plans and approve execution.

While some worry that AI will replace security analysts, we do not share that concern. AI will be crucial to counter adversaries at machine speed, but humans will still be needed for the most sophisticated incidents. What will change is the role of the SOC analyst. As AI systems evolve into institutional knowledge centers — teaching best practices to new analysts while executing routine work — analyst maturity will accelerate, and senior staff will gain bandwidth to devote to the most interesting, high-impact projects.