

AI for security: Agentic AI will be a focus for security operations in 2025

Analysts - Scott Crawford, Mark Ehr

Publication date: Thursday, March 20 2025

Introduction

It is no surprise that security has been seen as a primary use case for generative AI — for recognizing and acting on significant findings in the mountains of detection telemetry that can overwhelm human analysis. The promise of AI at the service of security doesn't end there. Automation has already had an impact on security operations, through domains such as security orchestration, automation and response. Coupling advances in automation with ongoing innovation in AI holds promise for overstressed security operations centers. Vendors are embracing these advances, but are their customers ready to as well?

The Take

Already, generative AI co-pilots and assistants have appeared as part of security operations (SecOps) technology platforms, looking to apply insight at scale to the daunting volume of threat data. We are seeing capability that enables intelligent agents to reason, come to independent decisions, use tools, and take actions in response to inputs and findings. This is part of the landscape of agentic AI: AI empowered with agency to act. While still some way from the ideal of the autonomous security operations center (SOC), these innovations will likely make themselves more evident in SecOps tech in 2025. Security organizations should expect to benefit — provided they are prepared to embrace them, and understand the limitations and tradeoffs they must consider to do so.

Context

Cybersecurity has always been an asymmetric fight. Attackers can focus on specific opportunities, but defenders must stretch resources across their estates as judiciously as possible. These already-thin assets — among which the most important are the people with the expertise to recognize and respond to attacks — are stretched even further by the sheer scope and scale of the threat landscape. Threat detection and response technology has evolved to help them get a handle on this challenge, but there is always room for further improvement.

The need for continued investment in the effectiveness of threat detection and response remains high. For the last several years, our 451 Research Voice of the Enterprise: Information Security

surveys have been asking SecOps practitioners what percentage of security alerts they are unable to handle in a typical day. While it has declined slightly in the last year, that figure remains at about 50%, reflecting not only the continued flourishing of cyberattacks, but also increased scale and complexity in the IT environment.

The emerging shape of intelligent SOC automation

The cybersecurity market segment of security orchestration, automation and response has gained prominence alongside other disruptive trends in SecOps, such as extended detection and response and behavioral analytics applied to threat detection. Vendors focusing their strategies on building out SecOps platforms have acquired players in this space. Others built out approaches to SecOps automation organically, while some continue to specialize in security automation as independent vendors.

While security automation can offload a number of tasks, such as gathering context for an investigation from various telemetry sources or automating the creation and management of incident cases, the deterministic nature of automation programming often means that these implementations must be maintained and adapted as requirements change. This can inhibit the benefits realized from automation, when the investment in maintenance and developing new automations and workflows must be kept up. It also poses a threat to automation's business value, and could even pose some risks to security effectiveness as attackers begin to see value in more AI-driven, non-deterministic tactics, techniques and procedures.

With the rise of generative AI, we have seen the introduction of functionality such as co-pilots — digital assistants that can bring the insight and context necessary to take effective action to the attention of human analysts. This functionality has the potential to handle larger volumes of data than people can, recognizing and surfacing significant activity more quickly and effectively. It can also provide context to detected behavior and help analysts become more knowledgeable about the threats they face.

Coupling this capability with the ability to take action could unite the values of generative AI with the ambitions of security automation. It is this aspect of taking action that is at the heart of agentic AI.

Examples

Contrary to what some headlines might suggest, agentic AI has been a branch of AI research for the last four decades. The recent high-profile disruption of generative AI has brought to the attention of the market a new generation of AI that can take action in ways that build on recent innovations in machine learning and reasoning.

Agentic capability may be deployed as focused functionality, equipped to perform specific tasks based on training and on the analysis and interpretation of input. A distinctive aspect of agentic functionality is its ability to take action based on its own assessment of appropriate steps. In a SecOps context, this may mean recognizing the need to gather additional context about an event from a variety of telemetry sources. It may further conclude that additional actions may be indicated, such as the need to create a case file or implement threat mitigation to shield an exposed vulnerability. A key aspect of such a sequence of activity is that it may not necessarily be deterministic: The agentic functionality may make its own determination of a best course of action.

This introduces a dynamic aspect to threat detection and response that has heretofore largely been human-centric. To be clear: People will necessarily need to remain engaged with such functionality, to assure that it performs to expectations, informs and augments the knowledge of analysts to better equip them to deal with the all-too-human nature of the "gamesmanship" of security, and make sure it aligns with human objectives and intent. It can't be overlooked that

this ability incrementally opens the door to the potential of what some think of as "autonomous" security operations in greater degrees as the technology is seen as increasingly reliable.

One example that indicates the direction of this coming change: Whereas today AI functionality may act as an assistant to a human analyst, using the analyst's access and credentials to perform tasks, more autonomous agents in the near future may act on their own behalf, logging their actions under their own identities and access privileges. That, at least, is the aim of those making the investment in this approach to innovation.

Agentic models and workflows


Agentic models can be deployed as focused agents purposed to specific tasks or actions in a workflow. This reduces the training burden and resource requirements for individual agents, which can make them more modular and readily deployed. As a result, they may be deployed in patterns that take advantage of their ability to take actions in sequence. A series of agents may move processes through a sequential workflow, taking output from "upstream" inputs or other agents, and producing output that is, in turn, consumed by successive agents, people or other functionality downstream.

This can also introduce some complexity in agentic workflows. Tasks may be directed to agents that specialize in certain tasks — such as obtaining telemetry from vertically specific operational technology and having the training to handle it in those specific use cases. Other tasks may benefit from parallelism, with several agents operating on the same or similar tasks simultaneously, or dividing large tasks into smaller groups for faster execution. This offers additional advantages, such as a variety of agents analyzing an environment for security vulnerabilities. Each different approach may incorporate different trainings or specializations that, together, yield results that can be combined into a more comprehensive analysis that includes multiple perspectives on a problem.

As workflow complexity increases, supervisory agents may be engaged to oversee more involved processes, or to engage additional agents or workflows as needed in more comprehensive task sets. We expect to see patterns emerge in SecOps that clarify where the elaboration of agentic workflows yields the greatest benefit. In more than a few cases, we expect providers to shield users from unnecessary complexity, giving them insight into overall progress and steps toward results as needed to meet their requirements for visibility and human engagement, as well as action.

We expect to see such innovations combine with other advances in process and workflow automation for security, taking advantage of achievements in robotic process automation and engaging emerging disciplines such as detection engineering in SecOps where appropriate. This larger trend is often referred to as "hyperautomation," applied as an overall term to highlight how advances in intelligent security automation are changing the nature of SecOps tech. The intersections of automation with AI are evident in our research, such as the expected growth in rules-based AI from our [Voice of the Enterprise: AI & Machine Learning, Use Cases 2025](#) survey, as depicted in Figure 1.

Figure 1: In 2025, organizations expect to see significant growth in rules-based AI



Source: 451 Research's Voice of the Enterprise: AI & Machine Learning, Use Cases 2025.

Q. What is the expected status of each AI type in your organization 12 months from now? Please include in your answers types of AI that you have already invested in, and will continue to invest in over the next 12 months.

Base: All respondents (n=1,006).

© 2025 S&P Global.

Implications

In considering this evolution, we see that organizations will need to consider its implications. Keeping people engaged with more intelligently automated processes to make sure they are executed as expected and informing them with evolving insight, such as changes in the regulatory environment that have an impact on security implementation, is just part of that consideration.

Agentic functionality may introduce costs in terms of computing resources to determine appropriate courses of action. Perhaps counterintuitively, it may take more time to make dynamic determinations compared with traditional, deterministic approaches, with actions measured in minutes compared with more deterministic systems. This is in part because, with deterministic techniques, humans assess and determine a course of action prior to the implementation of automation. This time for assessment thus doesn't factor into workflow processing time, as humans have taken it on "out of band." With agentic functionality, agents perform that assessment within the workflow, which may introduce latency into a process. As agentic techniques become accommodated to specific workflows — as well as mature in their implementation over time — this latency may be reduced. Early in adoption, however, it should be a consideration for initial deployments.

There's also a tradeoff to be considered with automation that can function with autonomy. It means that organizations will have to consider that they will be giving up at least some of the detail of control. Agentic functionality essentially means asking technology to dynamically produce and execute a plan, given the information with which it has been equipped. This, however, highlights the human role of guidance and influence over agentic actions. Regulations such as data privacy, for example, may require the need to audit, review and validate agentic functionality in ways that AI may not be equipped to incorporate without human involvement. Even when agentic or hyperautomated functionality can adapt to such new or emerging situations, it may still take unnecessary steps in processes that humans can identify and optimize. What organizations realize in return should, in the aggregate, be outcomes that are at least directionally

correct — and they will generate future knowledge that the organization will retain.

We plan to cover these trends in more detail this year, looking at specific technology and service providers and their offerings, given that we expect this trend to be one of the most significant shaping SecOps tech in 2025.