




451 Research Vanguard Report

April 2025

Agentic AI in SecOps

At the threshold of harnessing intelligent action

Commissioned by

S&P Global
Market Intelligence

©Copyright 2025 S&P Global. All Rights Reserved.



SentinelOne

Introduction

The cybersecurity challenge is vast, but most security operations (SecOps) teams are not. The volume and variety of data gathered by security tools overwhelm those teams. The human expertise needed to recognize, prioritize and respond to threats is in limited supply, and it is often difficult to find, train and retain personnel. This leads to gaps in threat detection, increased exposure to threat impact and the analyst fatigue that can plague security organizations.

Our research reflects this reality. Nearly half of respondents to a recent 451 Research SecOps survey say their security teams are unable to investigate more than 50% of their security alerts on a typical day.

Figure 1: Proportion of SIEM/security analytics alerts that teams are unable to investigate in a typical day



Q. What percentage of SIEM/security analytics alerts are you unable to investigate in a typical day?
Base: Respondents currently using SIEM/security analytics, abbreviated fielding (n=243).
Source: 451 Research's Voice of the Enterprise: Information Security, SecOps 2024.

It's hardly surprising, then, that enterprises believe AI and machine learning (ML) could enhance SecOps through more sophisticated threat recognition and rapid processing of large data volumes. Recent advances in generative AI allow users to interact with technology more naturally and integrate it into workflows. And AI is evolving beyond assistance to take independent action with human oversight.

This emerging capability is "agentic" AI: AI empowered with the agency to reason, decide, choose tools and act independently. It holds great potential for augmenting the expertise of security teams, and we are only at the beginning of this evolution. To use agentic AI effectively, security teams must be prepared to embrace it, mature in its application and mitigate associated risks. They are not without allies in this quest, however: Security technology providers are eager to enable these advances in their products and make them useful to security organizations.

The Take

Historically, automation relied on static rule-based systems, but advances in AI have equipped current models with capabilities for assessing inputs, evaluating them against training and inference data, and producing useful outputs. This has enabled emerging agentic AI to apply new skills that model dynamic reasoning, decision-making and tool use to take initiative and act. In SecOps, this agentic functionality is complemented by advances in security automation, which can be augmented by agents to significantly improve outcomes. Using agents to deal with the high volume of detection telemetry and response tasks allows personnel to focus on the more complex and nuanced security issues where human expertise has the greatest impact, such as understanding the human gamesmanship that is necessary to respond to constantly emerging threats.

To succeed in security operations, agentic AI requires a solid foundation. While static rule-driven process automation has proven effective, agentic AI introduces dynamic decision-making and action based on learning and inputs. Although fully autonomous functionality may be the goal, human interaction will likely remain necessary to ensure proper directions and outcomes as implementations mature.

Agentic AI also needs a responsive architecture that provides access to high-quality input data, which must often be collected, normalized, integrated and managed across multiple security technology segments. It requires tools that enable agents to act, such as agent-amenable APIs for specific inputs and outputs, and the memory to maintain necessary context.

SecOps technology providers are developing such systems. Their aim is to harness innovation with advanced AI models to improve threat detection and response while aligning with human objectives, leveraging human expertise to guide and inform the technology, iteratively reviewing its insights and actions, and ensuring human oversight and control.

Use cases

Previously, generative AI primarily focused on search, information synthesis and content generation (hence the term “generative” AI). Although agentic AI has been a research area for years, current trends can leverage recent innovations in generative technologies such as reasoning, planning and accessing domain-specific external information storable as memory. When combined with the ability to interface with tools and execute tasks, actions can be determined less by static code and more by agents’ dynamic choices, governed by policies. This evolution represents a progression toward increasing autonomy, balanced with human interaction to ensure governance and alignment with human objectives while allowing human expertise to flourish.

Agentic actions can be integrated into various workflows tailored to specific use cases and objectives. In SecOps, we may see patterns resembling linear sequences, routing, parallelism and orchestration. As agentic technology matures, additional patterns may emerge, but these examples demonstrate the utility of such workflows in SecOps.

Linear sequence

A linear sequence is the most straightforward example. It represents a process that can be broken down into well-defined subtasks where a sequence of agents can receive input, determine and take actions, and produce output that is then used as input by the next agent in sequence, continuing until an ultimate result is achieved.

In threat investigation, for example:

- An agent receives an alert and determines what additional information is required.
- This output may be supplied to a subsequent agent that is specialized in gathering specific information, such as activity context or environment configuration details.
- The next agent in the sequence, trained on policy, determines the next appropriate steps.
- Successive agents in the sequence may isolate a threat by reconfiguring the environment or opening an incident case.

Humans may be engaged depending on the nature and impact of workflow output, such as reviewing and escalating findings, supervising outcomes for alignment with objectives, alerting investigators of a case or assigning an incident to an analyst or team.

In cases where actions are triggered by policy thresholds, such as indicators of a denial-of-service (DoS) attack, agents may autonomously engage high-availability resources in the environment, adhering to policy, which may include human confirmation and supervision.

Routing

In some SecOps workflows, certain tasks may require specialized functionality or access to domain-specific information handled by agents designed for that purpose. Agents earlier in a workflow sequence can be trained to identify such cases and delegate tasks to specialized agents as appropriate.

Examples include:

- **Specialized agents focused on particular threat tactics** or trained on content that associates tactics with certain threat actors may enhance investigations by providing crucial insight and potential correlations with other investigations that may otherwise go unrecognized. These agents may represent capabilities offered by partners, such as threat intelligence or incident response providers, adding depth and reach to the investigation.
- **Investigation and response** may also require routing tasks to agents with specialized expertise in certain environments. This is often the case in operational technology, where environment reconfiguration to isolate an incident may require expertise in industry-specific technology and unique functionality from specific technology providers or original equipment manufacturers. This specialization can complement human insight, helping to prevent critical safety issues and protect vital dependencies.

Parallelism

In some cases, workflows can incorporate parallel tasks. When several tasks are similar, multiple agents can work on multiple instances of the same or similar tasks to accelerate the production of aggregate output.

- **Assessment of vulnerabilities throughout an environment** is an example of this pattern. In vulnerability assessment, various agents may take different analytical approaches that surface distinct findings. Human expertise combined with decision-making agents may determine the best distribution of such tasks based on factors such as agent specialization. Aggregator agents may assemble completed results that represent more comprehensive findings with fewer coverage gaps.

Orchestration

While routing tasks to individual agents can be straightforward, more complex scenarios may require the coordination of multiple workflows, each involving several agents. In these cases, supervisory agent functionality may play a more comprehensive systems role.

- **Investigations at scale** offer an instructive illustration. The dynamic use of elastic infrastructure to increase the number of agent workflows responding to complex scenarios across software, networking, infrastructure and access control domains — in concert with service provider workflows — supports the growing adoption of agentic technology use cases.

In all cases, human oversight may be engaged to ensure that outcomes align with expectations, enhance workflows with the insight that only human experts can provide and address areas where agents have not yet been adapted.

Conclusion


In SecOps, the goal of agentic AI is to increase the effectiveness and efficiency of threat recognition, response and mitigation. Agentic capability aims to complement assistive generative AI with action-enabled functionality, blending evolving autonomy with human expertise across diverse technology landscapes and threat tactics.

This journey has just begun, however. While agentic AI builds on previous security automation and orchestration approaches to reduce human effort, it may require more extensive computing resources to determine appropriate actions. It may take longer to make dynamic determinations compared to traditional, deterministic software, with actions measured in minutes rather than seconds. Regardless, exploring agentic functionality as it develops is essential to determine where it yields the greatest benefits, mature the technology and optimize its potential as we progress toward a future in which agentic AI is likely to play an increasing role.

Understanding the real impact of autonomy provides a useful perspective. While agentic AI may bring us closer to the ideal of autonomous security operations, organizations will give up at least some control over details. Agentic functionality means asking technology to dynamically produce and execute plans based on information it has been given. This highlights the importance of human guidance and influence over agentic actions. Regulations governing data privacy, for example, require auditing, reviewing and validating agentic functionality in ways that AI may not be equipped to incorporate without human involvement. Even when agentic functionality adapts to new situations, it may still take unnecessary steps that humans can identify and optimize. In return, organizations should realize outcomes that are at least directionally correct, and in the process they will generate valuable knowledge.

Many organizations believe that their environments are not yet ready to embrace this potential. For example, in a commissioned survey conducted by 451 Research, respondents indicate that their top concern about adopting generative AI for SecOps is that it would require a greater degree of technological maturity than they currently possess, or think they can achieve.

Figure 2: Top concerns about adopting generative AI for SecOps



Q: What are your greatest concerns about adopting generative AI for SecOps? Select all that apply.

Base: All respondents (n=606).

Source: S&P Global Market Intelligence 451 Research Security Operations custom survey, 2024.

This is one reason why SecOps technology vendors are investigating the potential of agentic AI and plan to incorporate it into their offerings. They aim to ease and accelerate adoption and support the growing maturity that will follow from experience. Those building platforms to integrate data and architecture across multiple functions have an imperative to build the necessary foundations. Ideally, they should offer platforms that make it easier for organizations to tailor agentic functionality to their specific needs and integrate it into existing workflows. This approach should optimize the value of this new functionality without requiring SecOps teams to radically retrain analyst behavior.

These developments are happening now, and security teams should investigate the current blending of generative and agentic AI implementations in SecOps to better understand how they can benefit from emerging innovations.



SentinelOne (NYSE:S) is a global leader in AI-powered cybersecurity, enabling modern enterprises to protect, detect, and respond at machine speed. At the heart of our AI strategy is Purple AI—an advanced security analyst designed to help teams investigate threats, automate responses, and stay ahead of attackers. From real-time threat detection to intelligent automation, SentinelOne is redefining what's possible with AI in cybersecurity.

To learn more about SentinelOne and its AI-powered capabilities:

<https://s1.ai/AIpowered>

About the author



Scott Crawford

Research Director, Information Security

Scott Crawford is research director of the Information Security channel at S&P Global Market Intelligence 451 Research, where he leads the industry analyst team covering innovation, disruption and strategic players in cybersecurity and cyber risk. Scott joined S&P Global through its 2019 acquisition of 451 Research, where he has led the Information Security channel since 2015. In addition to directing the Information Security channel's research efforts, Scott covers forces and events shaping cybersecurity. He maintains a focus on areas including security operations, cyber risk management, the intersection of AI/machine learning and cybersecurity, and related interests.

As a practitioner, Scott was the first information security officer for the Comprehensive Nuclear-Test-Ban Treaty Organization's International Data Centre in Vienna, with a background including systems and security management at the University Corporation for Atmospheric Research in Boulder, Colorado. His private-sector experience ranges from startups to leading industry players such as IBM, where Scott was a senior strategist with IBM Security.

Scott holds a Bachelor of Arts in molecular, cellular and developmental biology from the University of Colorado and a postgraduate Master of Science from the University of Salford (UK), with additional graduate study in telecommunications at the University of Colorado and in information systems at the University of Denver.

About S&P Global Market Intelligence

At S&P Global Market Intelligence, we understand the importance of accurate, deep and insightful information. Our team of experts delivers unrivaled insights and leading data and technology solutions, partnering with customers to expand their perspective, operate with confidence, and make decisions with conviction.

S&P Global Market Intelligence is a division of S&P Global (NYSE: SPGI). S&P Global is the world's foremost provider of credit ratings, benchmarks, analytics and workflow solutions in the global capital, commodity and automotive markets. With every one of our offerings, we help many of the world's leading organizations navigate the economic landscape so they can plan for tomorrow, today. For more information, visit www.spglobal.com/marketintelligence.

CONTACTS

Americas: +1 800 447 2273

Japan: +81 3 6262 1887

Asia-Pacific: +60 4 291 3600

Europe, Middle East, Africa: +44 (0) 134 432 8300

www.spglobal.com/marketintelligence

www.spglobal.com/en/enterprise/about/contact-us.html

Copyright © 2025 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global keeps certain activities of its divisions separate from each other to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain nonpublic information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its websites, www.standardandpoors.com (free of charge) and www.ratingsdirect.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.