Omdia
by informa techtarget

The Agentic SOC

Exploring the Practitioner Mindset as AI Permeates SecOps

Adam DeMattia | *Senior Director, Research*
David Gruber | *Principal Analyst, Cybersecurity*

February 2026

This Omdia eBook was commissioned by Google and is distributed under license from TechTarget, Inc.

Research objectives

As AI-infused SecOps solutions strive to change the operating dynamics of modern security operations, security practitioners are confronted with both opportunity and challenges in how and where to assimilate this powerful new technology. With the promise of significant improvement to security outcomes, coverage, and the efficiency of SecOps, practitioners must ultimately adopt, learn, and transform established practices to reap the benefits of an AI-powered operating environment—all while under pressure to continue to secure and govern their attack surface against an AI-enabled adversary.

To gain further insight into these trends, Google commissioned Omdia to execute a survey of 300 security practitioners and SOC managers at organizations in North America and Canada involved with or responsible for the use of AI-enabled technologies within security operations.

THIS STUDY SOUGHT TO:

- **Assess SecOps teams' adoption of agentic AI technologies.**
- **Understand how security practitioners view agentic AI** and how they educate themselves about agentic innovations.
- **Identify security "jobs to be done" where respondents trust AI agents to help** (as well as where they don't) and what's needed to build trust looking ahead.
- **Measure perceived early wins being delivered by agentic AI** within SecOps teams and the teams' expectations for the future.

Note: Totals in figures and tables throughout this eBook may not add up to 100% due to rounding or organizations choosing more than one answer to select questions.

Back to Contents

Contents

Back to Contents

89% of CISOs are pushing to accelerate the adoption of agentic AI solutions for security

Back to Contents

"Over half (54%) of cybersecurity practitioners indicated their organization is already using agentic AI (31%) or planning to use and deploy it over the next 12 to 24 months."

Rapid acceleration in agentic AI deployments

Agentic AI refers to AI systems capable of independent decision-making and autonomous behavior. These systems can reason, plan, and perform actions, adapting in real time to achieve specific goals. A key capability of agentic AI is its ability to select, access, and utilize various tools and external resources to accomplish tasks more effectively. These systems can learn, adapt, and improve their performance over time. Many agentic AI solutions interact with users and even other AI agents via generative AI interfaces.

Over half (54%) of cybersecurity practitioners indicated their organization is already using agentic AI (31%) or planning to use and deploy it over the next 12 to 24 months (23%). Another 36% are interested in agentic AI but admitted that they are still early in the adoption process. This landslide commitment to putting AI to work to improve SecOps reflects the urgency in combatting an AI-enabled adversary together with the rapid transformation of the operating infrastructure that security teams are tasked with protecting.

Which of the following best describes your organization's use of, or interest in, agentic AI technologies to augment or execute security operations processes and workflows?
(Percent of respondents, N=300)

We have no plans for or interest in agentic AI technology

We are interested in agentic AI technology but early in the adoption process

We have definitive plans for agentic AI technology and will likely deploy solutions to production within the next 12-24 months

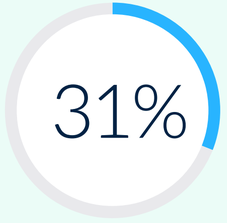

We are leveraging agentic AI technologies in production today

Back to Contents

Larger SOC teams are adopting more agressively

Looking closer at these numbers, we see that larger SOC teams are further along in the adoption of agentic AI vs. SOC teams of a smaller size. Notably, many of these larger SOC teams are not part of larger companies but rather are from companies with fewer than 5,000 employees.

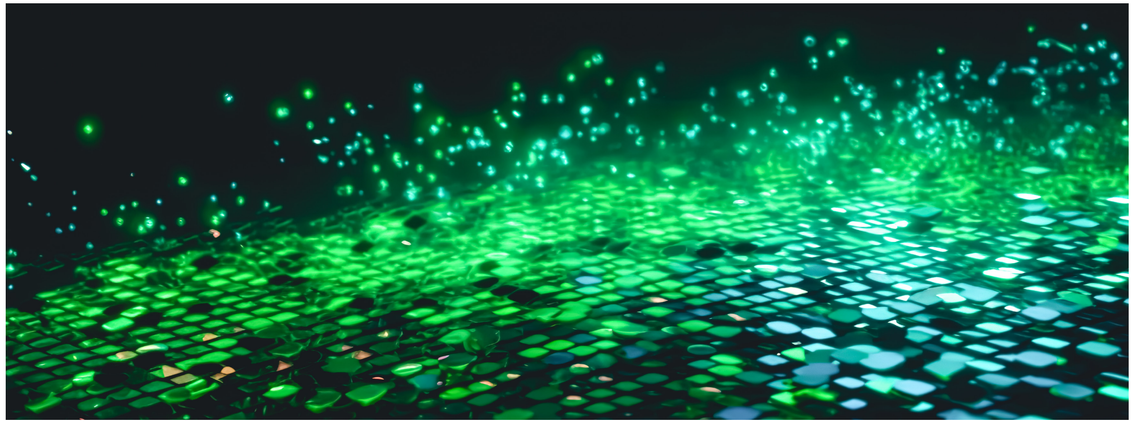

Which of the following best describes your organization's use of, or interest in, agentic AI technologies to augment or execute security operations processes and workflows?
(Percent of respondents)

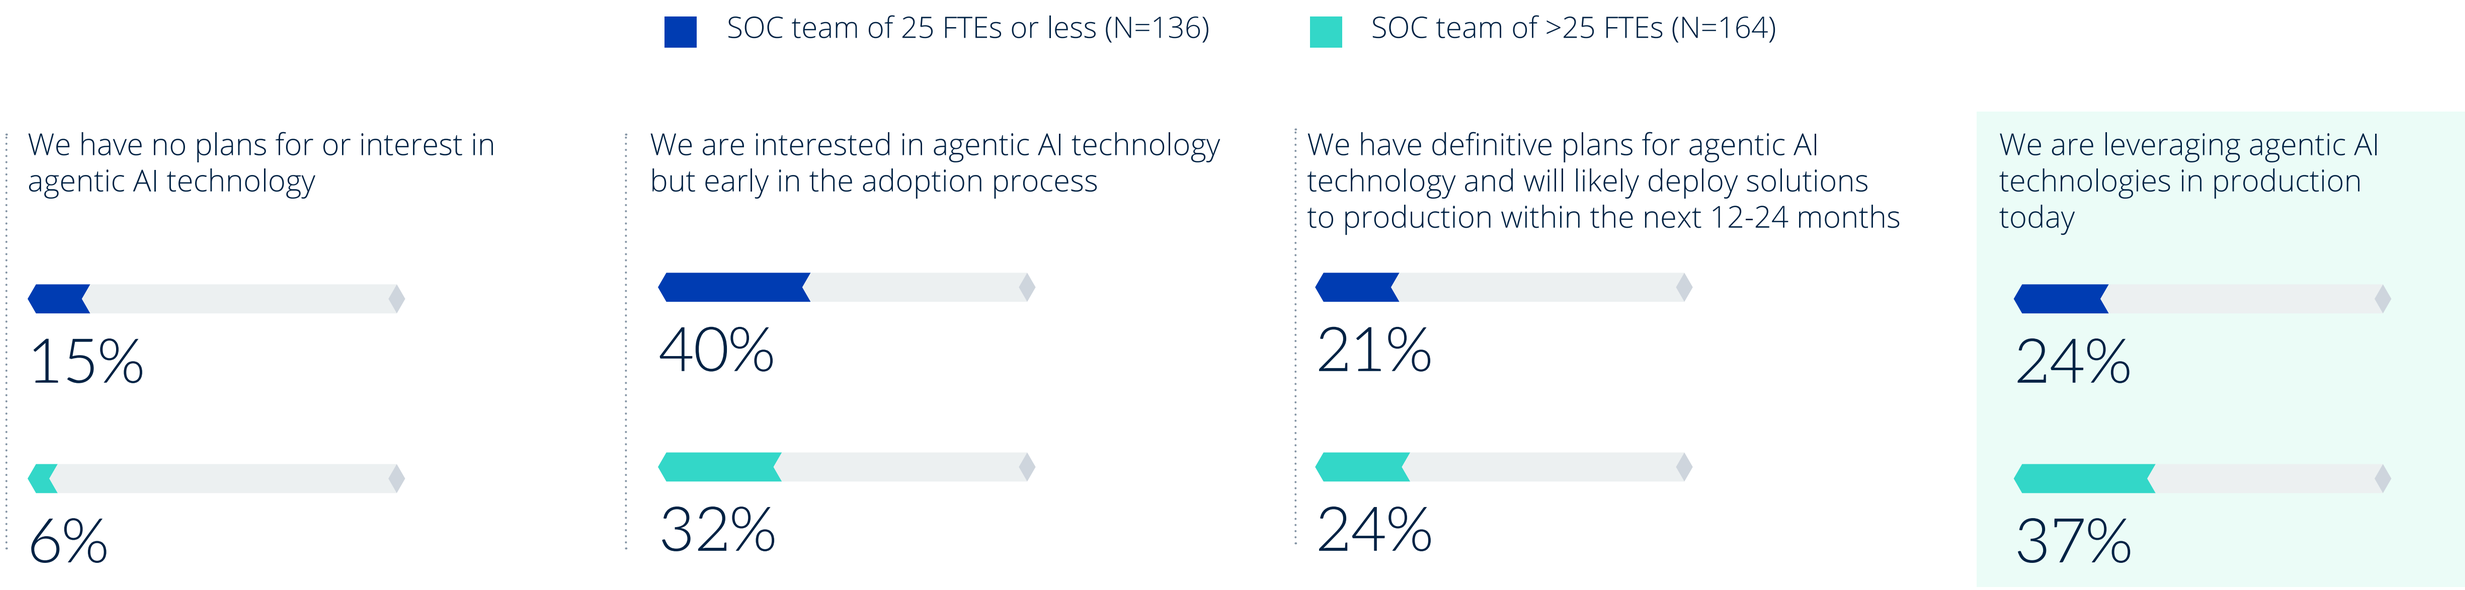

Back to Contents

Agentic AI opportunities in the SOC are multifaceted

When it comes to where the opportunity lies for agentic AI to make a difference in security operations, the survey found that most organizations could see a wide array of applications and use cases, some more horizontal, such as improving decision-making and general efficiency improvements, with other more vertical functions specific to improvements in threat detection and intelligence. Notably, the survey further revealed that experienced security professionals are more likely to indicate that faster threat detection is an area they find agentic AI to be the most promising (43% vs. 31% of less experienced professionals).

Based on this multifaceted array of applications and use cases, while we expect to see agentic AI broadly applied, it is still rapidly evolving and not fully understood. This is also true across demographic and firmographic subsets of both current and planned users.

What aspects of/opportunity areas for agentic AI do you find most promising for security operations?
(Percent of respondents, N=300, up to three responses accepted)

39%
Faster threat detection

37%
Better threat intelligence analysis and operationalization

33%
Enhanced threat hunting capabilities

36%
Improved incident response

35%
Better triage to prioritize alerts and enable faster decision-making

30%
More efficient threat investigations

34%
Automated routine or rote tasks

30%
Reduced alert fatigue and more manageable alert backlogs

Back to Contents

Use cases are similarly multifaceted

Agentic AI's transformative potential is on full display, based on the wide variety of use cases reported. Notably, when we look at specific use cases where people report AI agents are most suitable for use within SecOps, we see that those already utilizing agentic AI are more likely to indicate that threat detection and pen testing are important use cases: 2x for threat detection (34% vs. 17% of less experienced professionals) and 1.8x for pen testing (27% vs. 15% of less experienced professionals).

Which of the following use cases and tasks do you believe are most suitable for AI agents in security operations?
(Percent of respondents, N=300, up to five responses accepted)

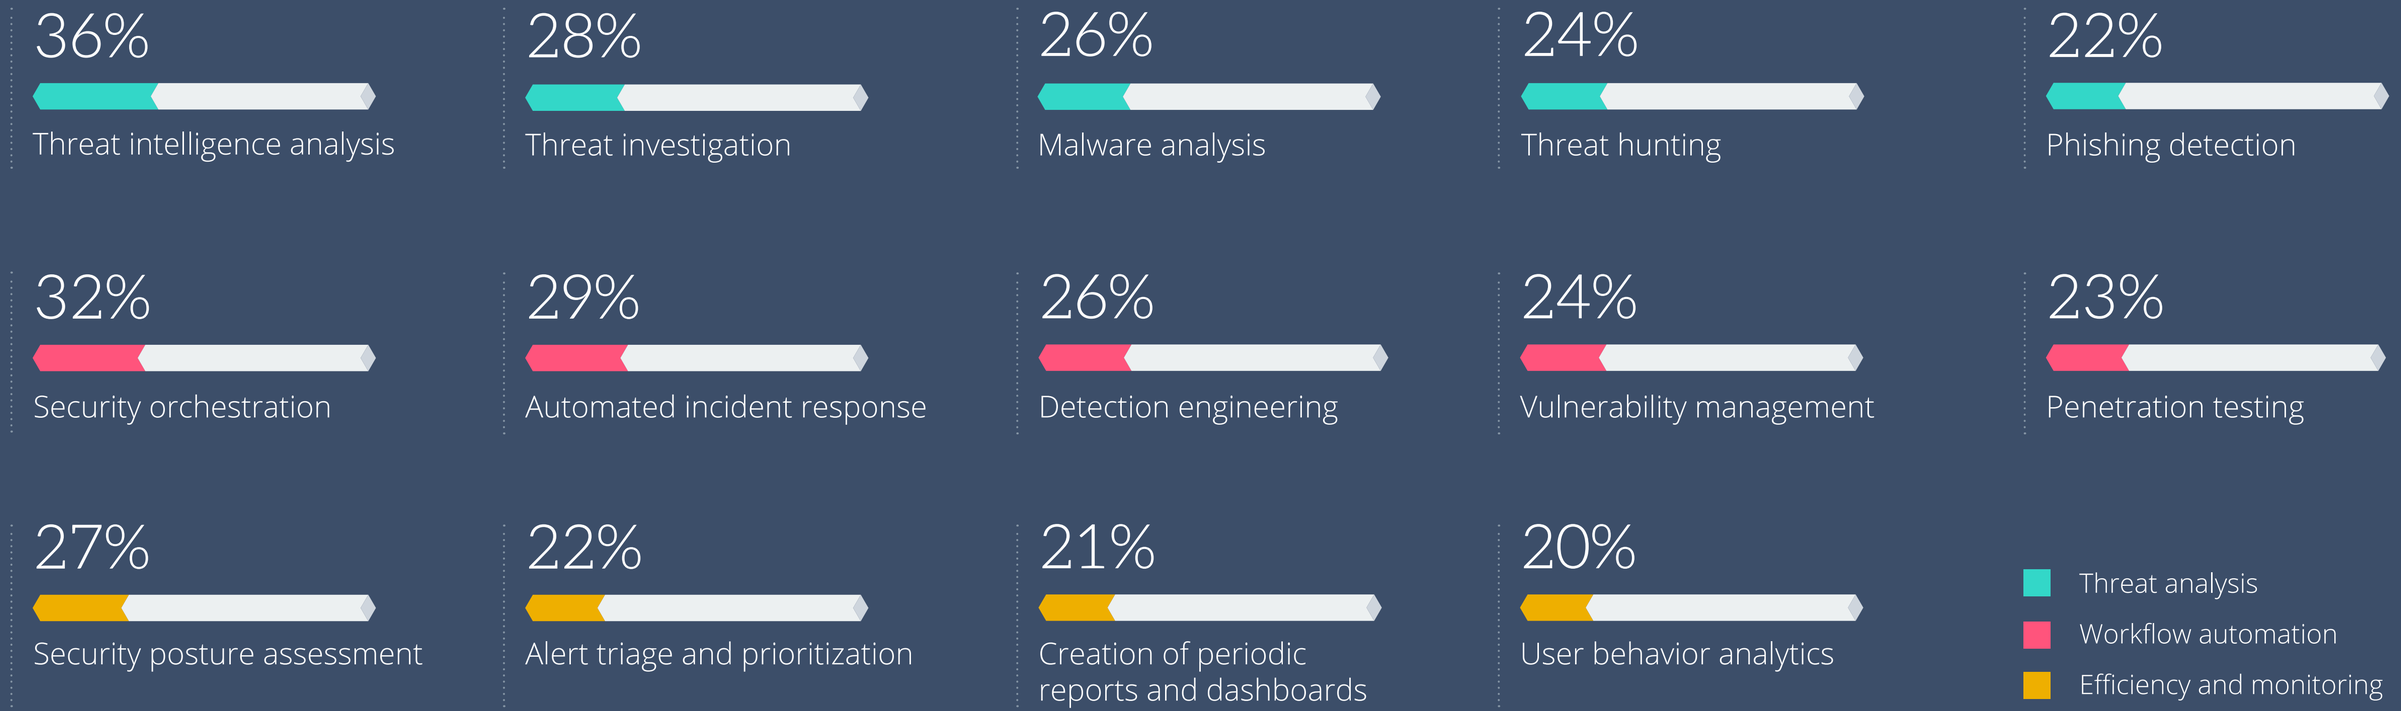

Back to Contents

“96% of adopters reported a measurable positive impact from implementing agentic AI solutions within their SecOps.”

AI agents in the SOC are delivering

96% of adopters reported a measurable positive impact from implementing agentic AI solutions within their SecOps, with half reporting the impact is significant.

Given the speed of adoption and the early maturity levels of current deployments, these numbers are not only impressive but also arguably unprecedented in the history of cybersecurity technology advancements.

You mentioned your organization has implemented early-generation agentic AI solutions in your security operations. Generally speaking, how would you rate their contribution?
(Percent of respondents, N=94)

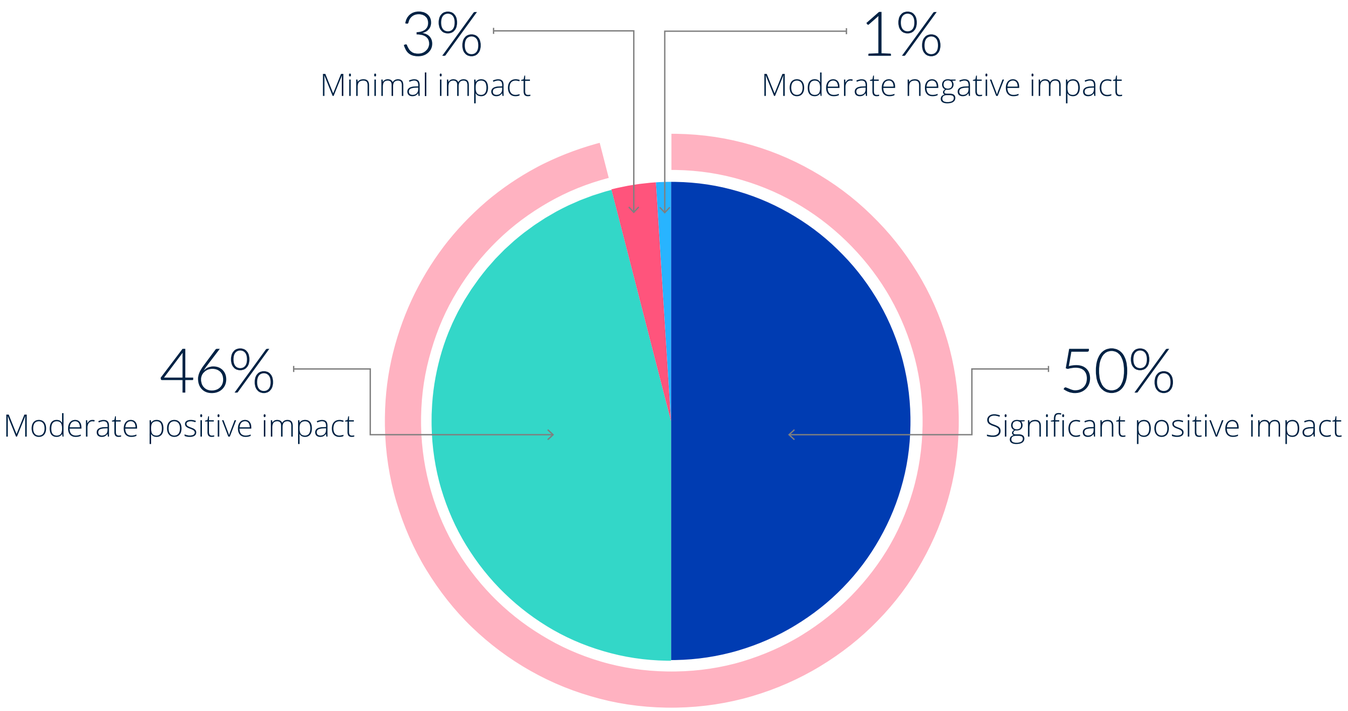

Back to Contents

Agentic AI delivers across all SOC functions

Looking closer at where and how the use of agentic AI is delivering value, **87%+ reported a meaningful positive impact in each SOC KPI/function included in the survey.** These results span not only improvements in the speed of SecOps processes but also further reflect an immediate improvement in alert fatigue and backlogs, with 88% of practitioners reporting positive improvement here

> "87%+ reported a meaningful positive impact in each SOC KPI/function included in the survey."

What impact have agentic AI solutions had in each of the following specific areas of security operations?
(Percent of respondents, N=94)

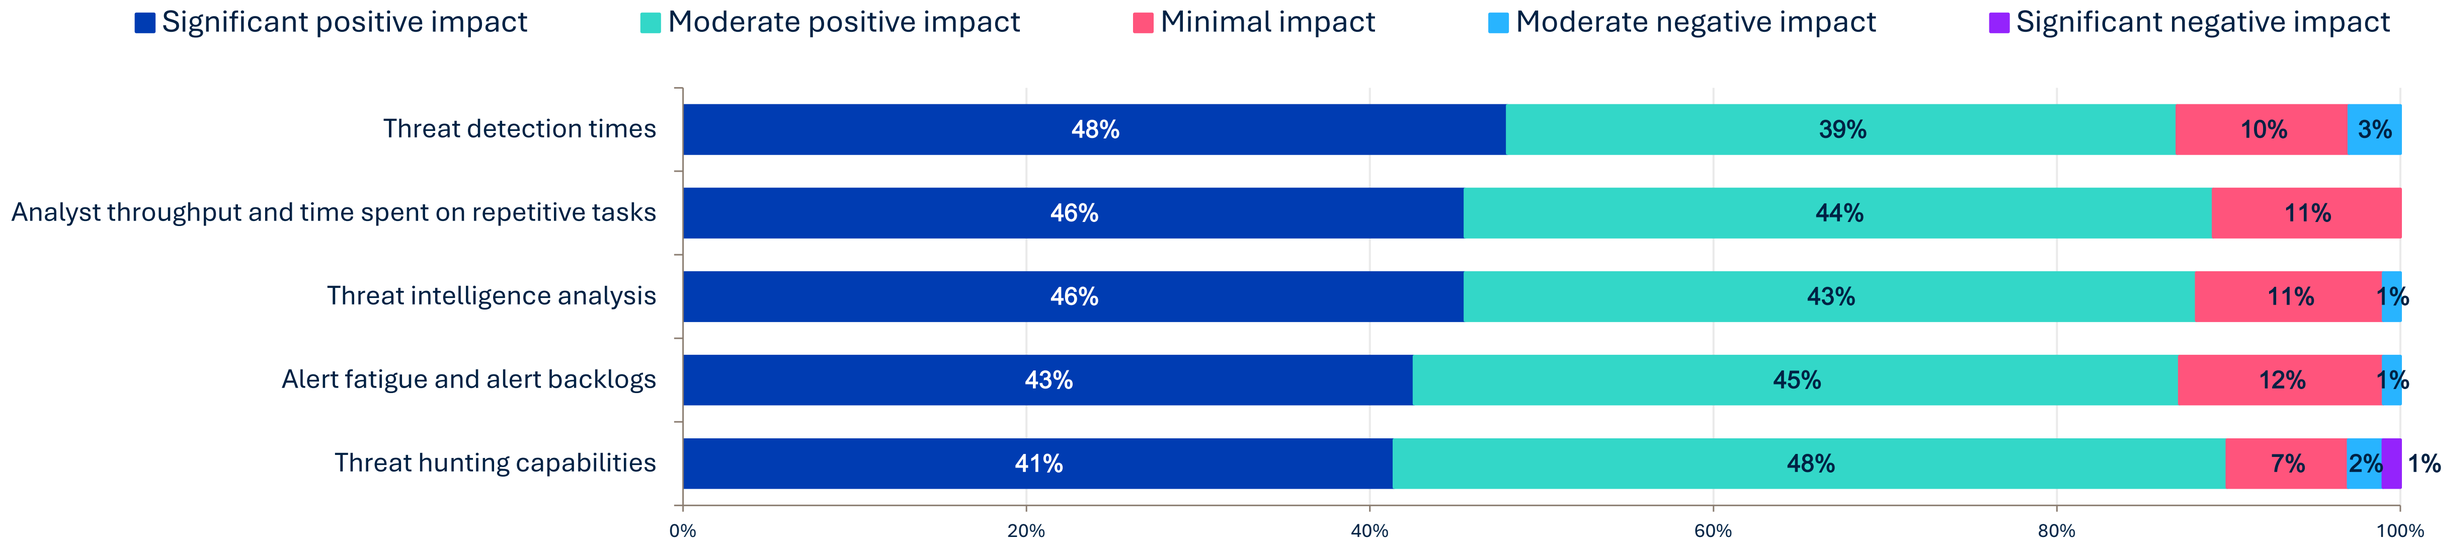

Back to Contents

Four key challenges to overcome as utilization advances

Back to Contents

Challenge 1: Security data architectures need significant redesign to accommodate agents

In total, more than four out of five (82%) security professionals indicated that moderate or complete data architecture redesign is needed for agents to be effective. Organizations racing to deploy agents without getting their security data house in order will likely see lackluster results.

Do you believe your organization's current security data architecture needs to be redesigned to effectively support agentic AI?
(Percent of respondents, N=300)

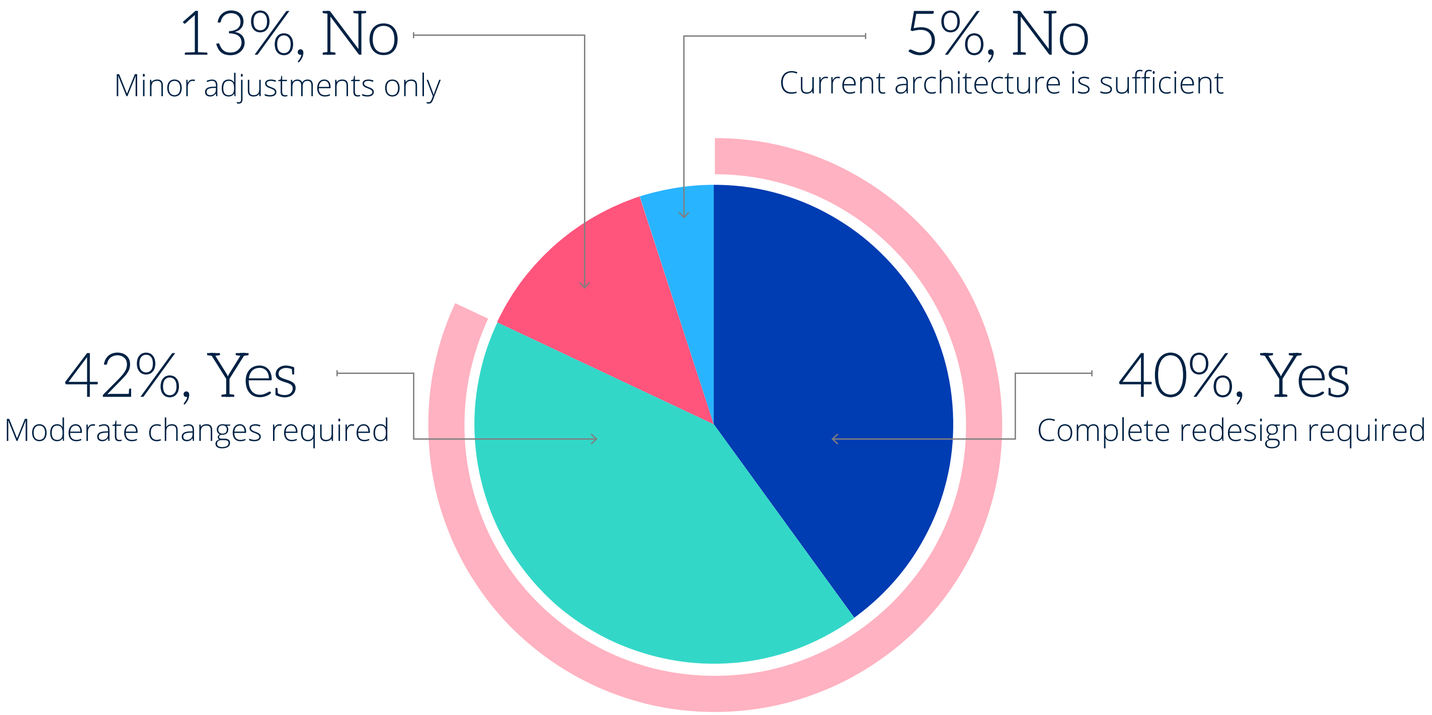

Back to Contents

Quantifying the data readiness problem

How much security data is already “AI-ready”? On average, only ~48% of the security/telemetry data is considered “AI-ready,” meaning more than half (52%) of the security data estate needs some level of transformation prior to agents being able to use it. Considering the impact data quality has on AI outcomes, this should be a top priority for organizations.

What percentage of all your organization’s security/telemetry data do you consider to be “AI-ready” (e.g., properly prepared and organized to be effectively used by AI systems)?

(Percent of respondents, N=300)

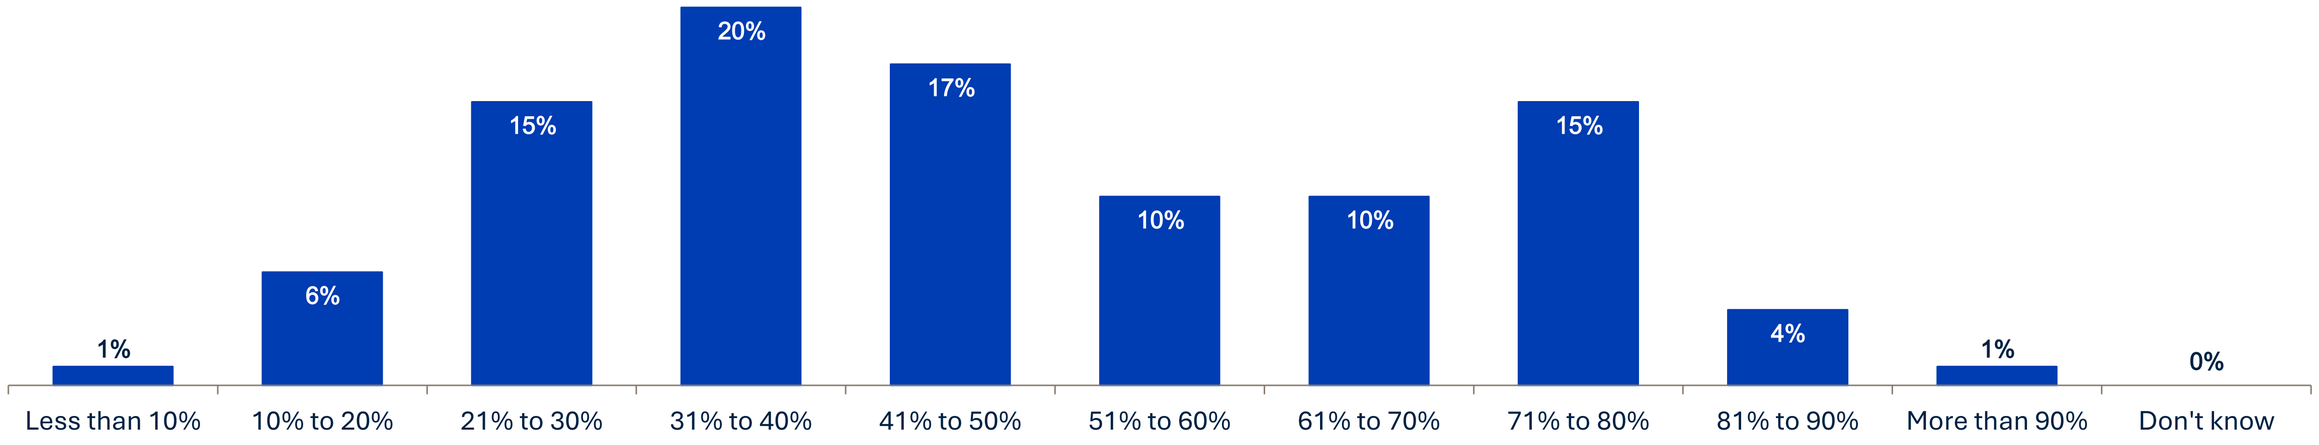

Back to Contents

As architects and engineers scale agentic usage, the need for data rearchitecture becomes clear

Looking closer at the demographics, we saw that security architects and engineers were more passionate about the need for a complete data redesign than SOC managers were, reminding us of the importance for practitioners to clearly communicate this need to management.

Percentage of respondents selecting “Yes, complete redesign required.”

Role comparison

44%
Security analysts/engineers

26%
SOC managers

Back to Contents

Leadership mandates are driving strong momentum for agentic AI adoption for SecOps, and agentic differentiation will catalyze switching behavior

Challenge 2: Determining what's real in a noisy market landscape

With much confusion in the market about what is hype versus what is real, security teams need to push vendors to both articulate and demonstrate evidence of where specific improvements are enabled, together with proof of operational reliability of AI capabilities. When AI-enabled capabilities demonstrate superior outcomes, many will rapidly replace existing solutions. Both vendors and security teams are on this AI journey together, reinforcing the need for vendor-buyer partnerships that are willing to work together to progress.

Back to Contents

Data is reinforcing the perception that agents will augment rather than replace humans in the SOC

Beyond specific vendor capabilities, there are more general challenges in the AI adoption journey. In relative terms, concerns related to using agents within SecOps **most often include a lack of trust, which will likely limit autonomy.** Trust is a journey, and things are moving fast, as demonstrated by the more favorable results we see in our survey between those already in production vs. those still in implementation and planning.

We also found it interesting to see the difference in perspective between SOC manager and security analysts/engineers, especially regarding the trust agenda. SOC managers' biggest concern is trust, while worries about a lack of skills were 2x greater for analysts/architects vs. SOC managers.

Challenge 3:
Overcoming the trust gap

What concerns you most about the potential use of agentic AI in security operations?
(Percent of respondents, N=300, up to three responses accepted)

Reliability of decisions and actions made through agentic functions

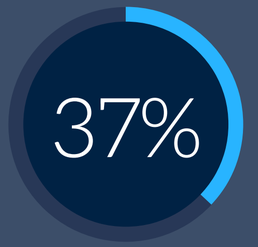

Data privacy implications

Potential for false positives or negatives

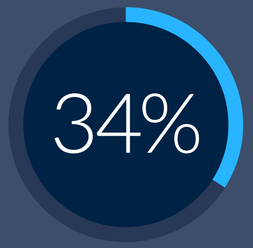

A lack of skills to manage AI systems and solutions

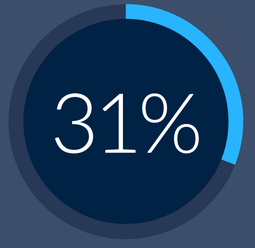

Costs of solutions

Lack of explainability in AI decisions

Job displacement

None of the above

Back to Contents

Reliability and human-centric solutions are non-negotiable as organizations evaluate agentic solutions for the SOC

What will it take to increase trust? When it comes to enabling agentic AI to execute core SecOps functions, 87% said proof of reliability and hands-on validation is critical (46%) or important (41%) to establishing trust. Beyond this proof, organizations expect solutions to provide smooth workflows that enable effective human-agent collaboration, as we better learn how machines and people work together most effectively. Beyond human collaboration, the need for inter-agent collaboration has quickly risen in importance, as organizations architect solutions tailored to their specific needs and environment.

How important are each of the following capabilities to you personally trusting agentic AI technologies to execute core SecOps functions? (Percent of respondents, N=300)

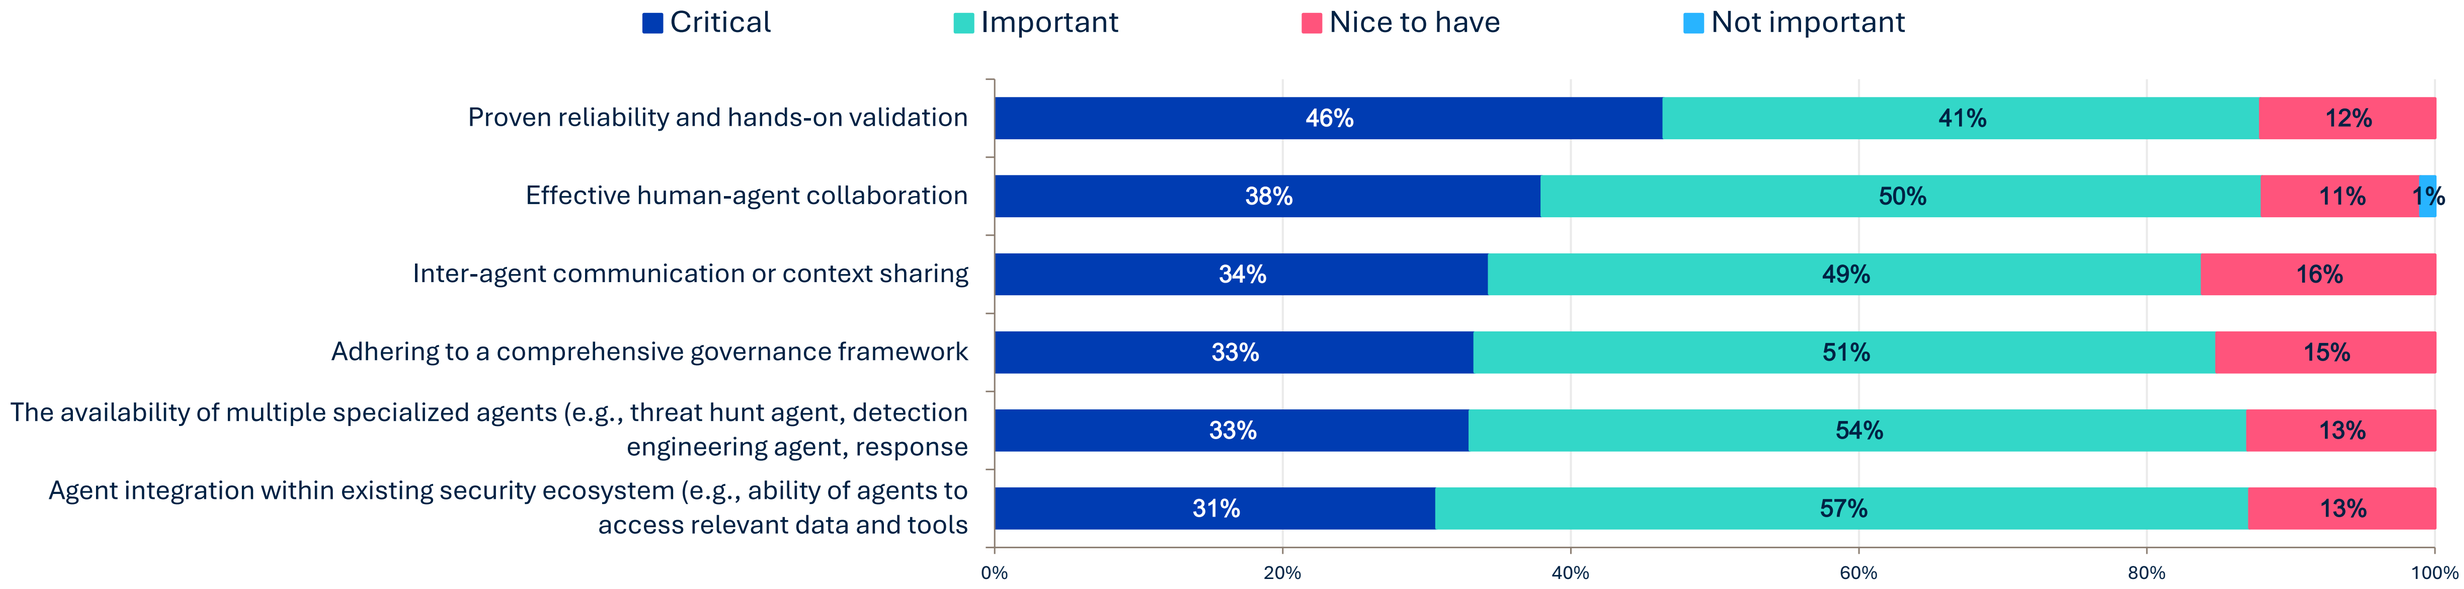

Back to Contents

AI systems management challenges are more apparent to practitioners

Beyond trust, other concerns exist, with 36% of security analysts and engineers worried about their ability to onboard the right skills to manage AI systems operating within their environment. Note that security analysts/architects were over 2x more likely to be concerned about a lack of skills than SOC managers were, reminding us, again, of the importance for practitioners to clearly communicate this need to management.

Challenge 4: A lack of skills to manage AI systems and solutions

What concerns you most about the potential use of agentic AI in security operations?

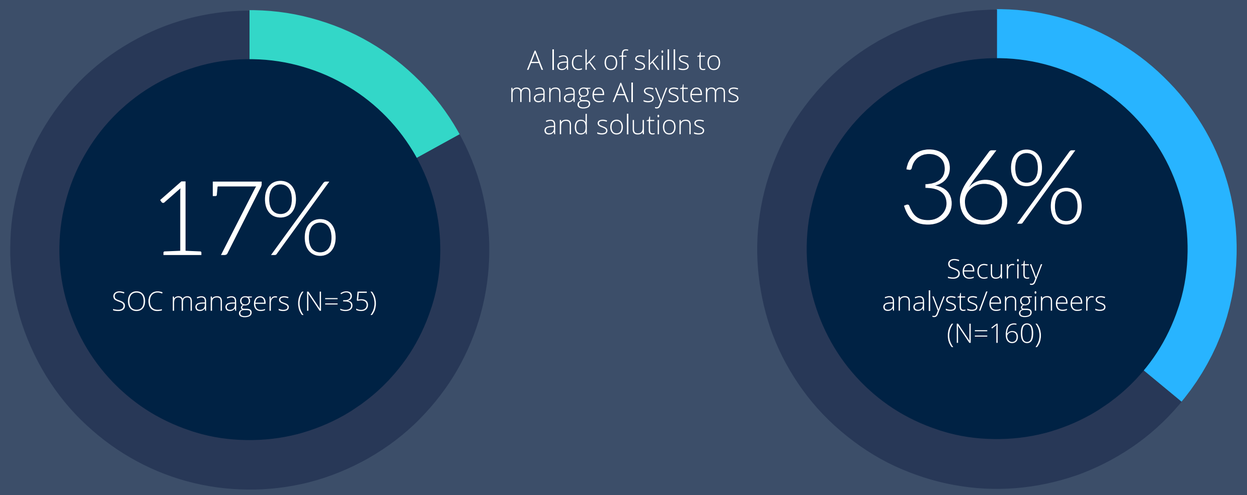

"36% of security analysts and engineers worried about their ability to onboard the right skills to manage AI systems operating within their environment."

Back to Contents

The future of SecOps with agentic AI

Back to Contents

Agentic AI to improve career prospects

Security practitioners are bullish about the impact of agentic AI on their personal career prospects over the next few years, with half (50%) believing that agentic AI will improve their career opportunities. And those more experienced were even more bullish, as 56% saw improved opportunities vs. 39% of less-experienced professionals.

Additionally, it's clear from the research that the more people learn and use agentic AI, the more optimistic they become, as 69% of those already using agentic AI are optimistic about career opportunity improvement.

Which do you think is the most likely impact of agentic AI technologies on your personal career prospects as a cybersecurity practitioner over the next few years? (Percent of respondents, N=300)

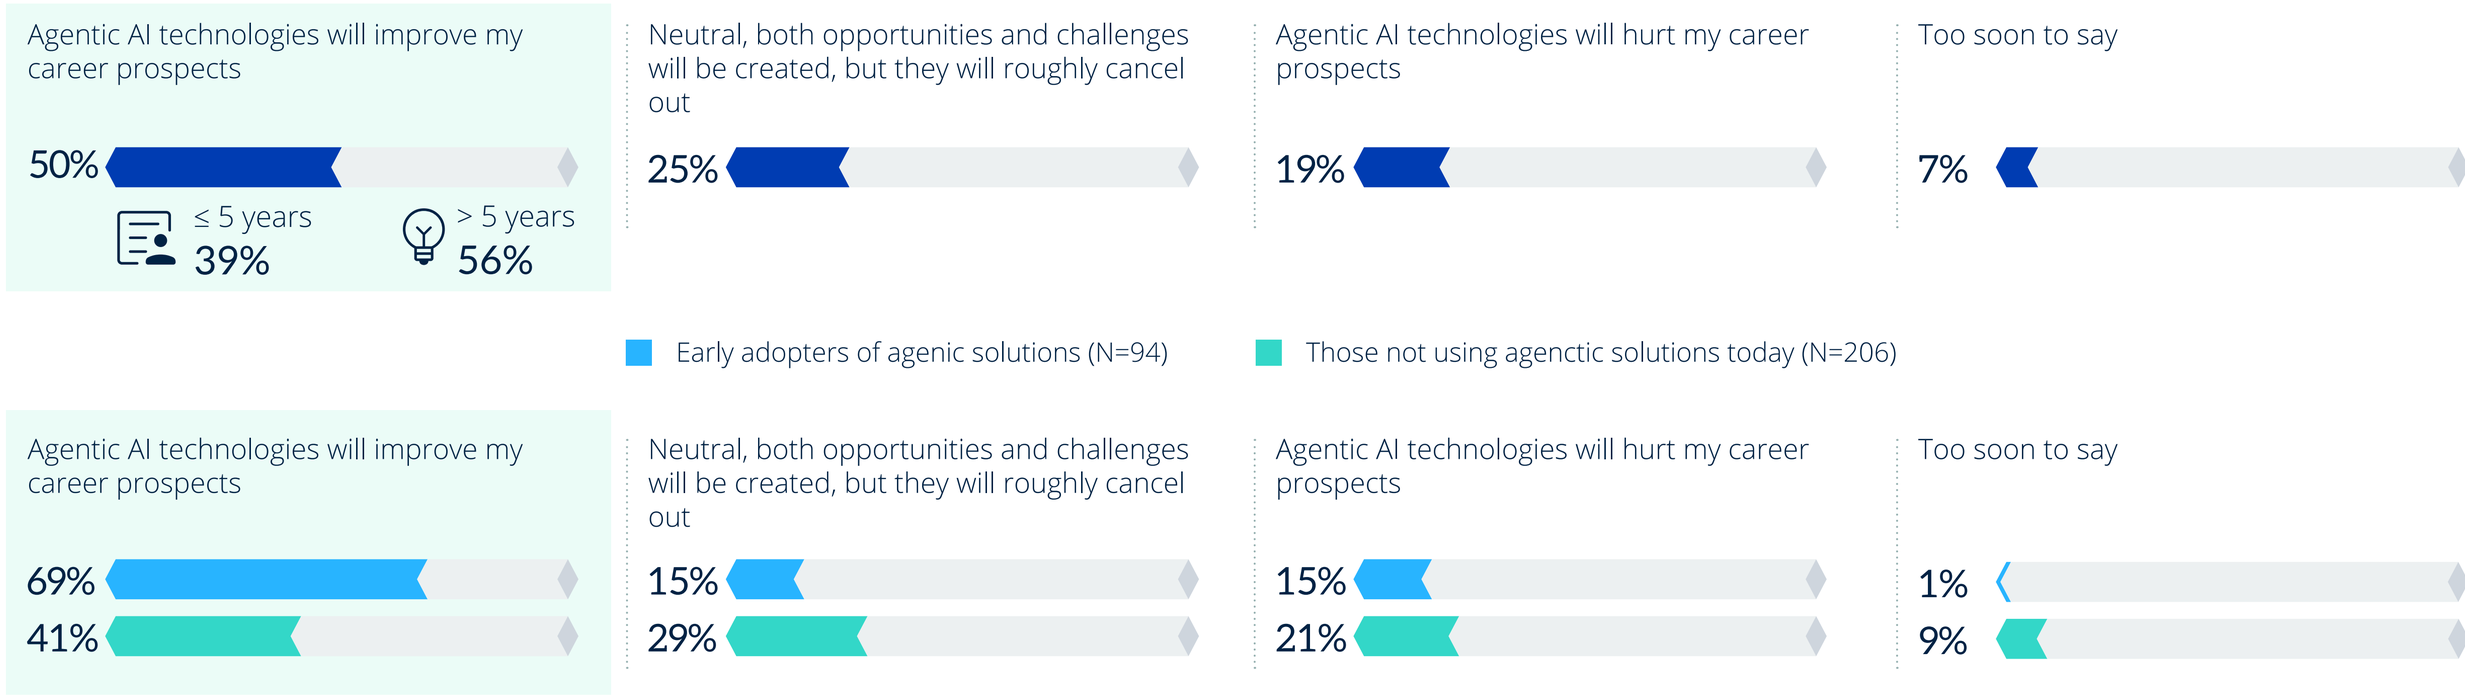

Back to Contents

Respondents are even more optimistic about agentic AI's ability to reduce enterprise risk

Beyond career opportunities, over half (56%) of cybersecurity practitioners indicated that agentic AI offers a bigger advantage to cybersecurity defenders vs. the adversary. And more experienced professionals are even more likely to indicate the same sentiment (63% vs. 44% of less-experienced professionals). Beyond experience as a security professional, those that have experience with agentic SOC solutions were 57% more likely to report agentic AI will deliver bigger benefits to defenders (rather than attackers).

When you think about the relationship between bad actors and cybersecurity defenders, which do you think will gain a bigger advantage from agentic AI technologies over the next few years? (Percent of respondents, N=300)

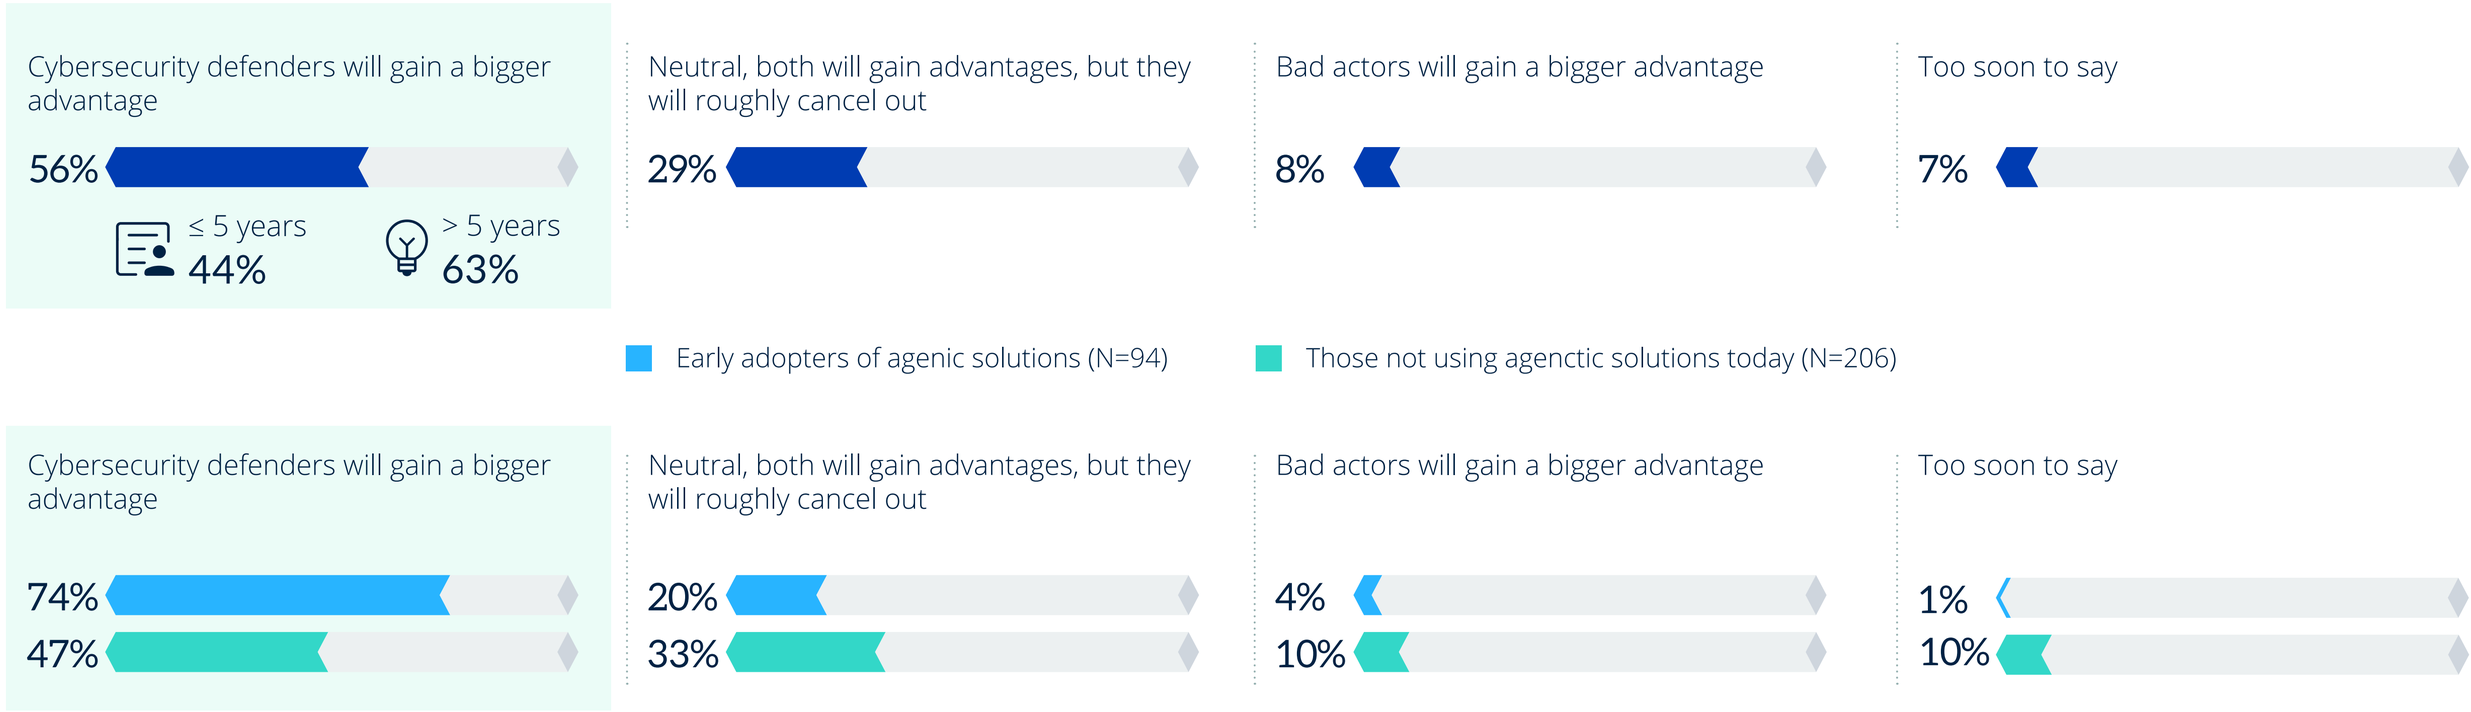

Back to Contents

What will security teams focus on as the SOC becomes more autonomous?

Beyond productivity improvement, the use of agentic AI will most likely change how SecOps teams function, including where security practitioners and architects spend their time every day. While still in early days, our survey shows that practitioners expect to spend more time on proactive preparedness, threat intel analysis, data analytics, and collaboration with the business, leaving more tactical, operations-like alert triage to machines.

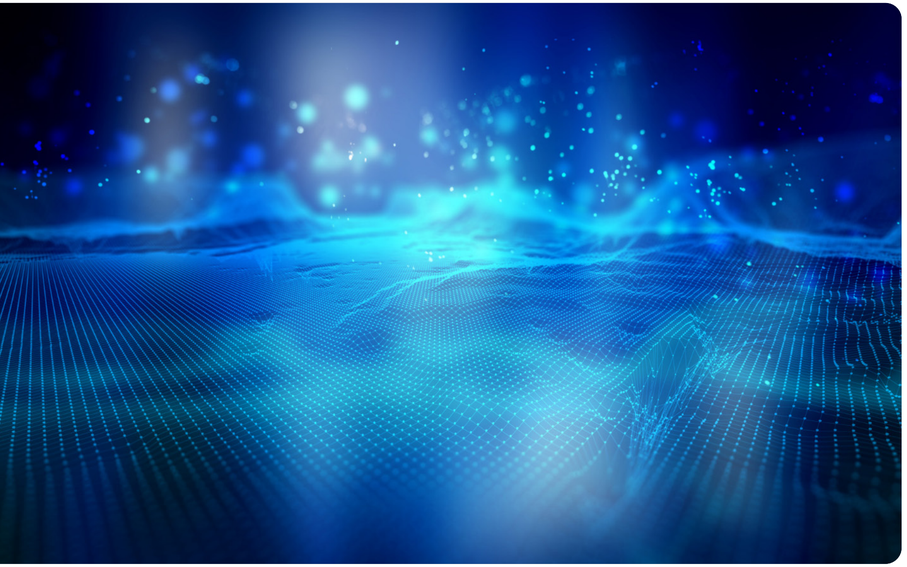

If agentic AI is widely adopted within the SOC, what do you think SOC personnel will spend more time on/focus on? (Percent of respondents, N=300, up to five responses accepted)

Back to Contents

Conclusion

Few would argue that the progress made in the past 12-18 months is remarkable in our ability to put AI to work to improve SecOps. Yet there is much to learn and do as we continue our journey to being more autonomous. As AI capabilities advance, we must rearchitect the security data layer to support a network of data-hungry agents throughout every function within SecOps. We must also understand what it takes to deploy, monitor, and manage an AI-driven operating stack. And as practitioners offload more manual, time-consuming activities to machines, we must redefine how people and machines collaborate to improve security outcomes, manage risk, and enable innovation. The big question, though, is when will we be at a tipping point where trust and capabilities mature to a level where we can refocus on other more strategic activities? One thing is certain though: This is a fast-moving train with only one choice ahead, and organizations need to get on board, or they risk getting run over.

“This is a fast-moving train with only one choice ahead, and organizations need to get on board, or they risk getting run over.”

Back to Contents

Google

ABOUT GOOGLE

Google is ushering in a new era in AI-powered SecOps productivity. Google's unique full-stack AI capabilities, powerful automation, and world-class experts help when you need it. Agentic AI and automation make everyone more productive and bend the curve on decades-old SecOps challenges to protect organizations like never before.

Learn More

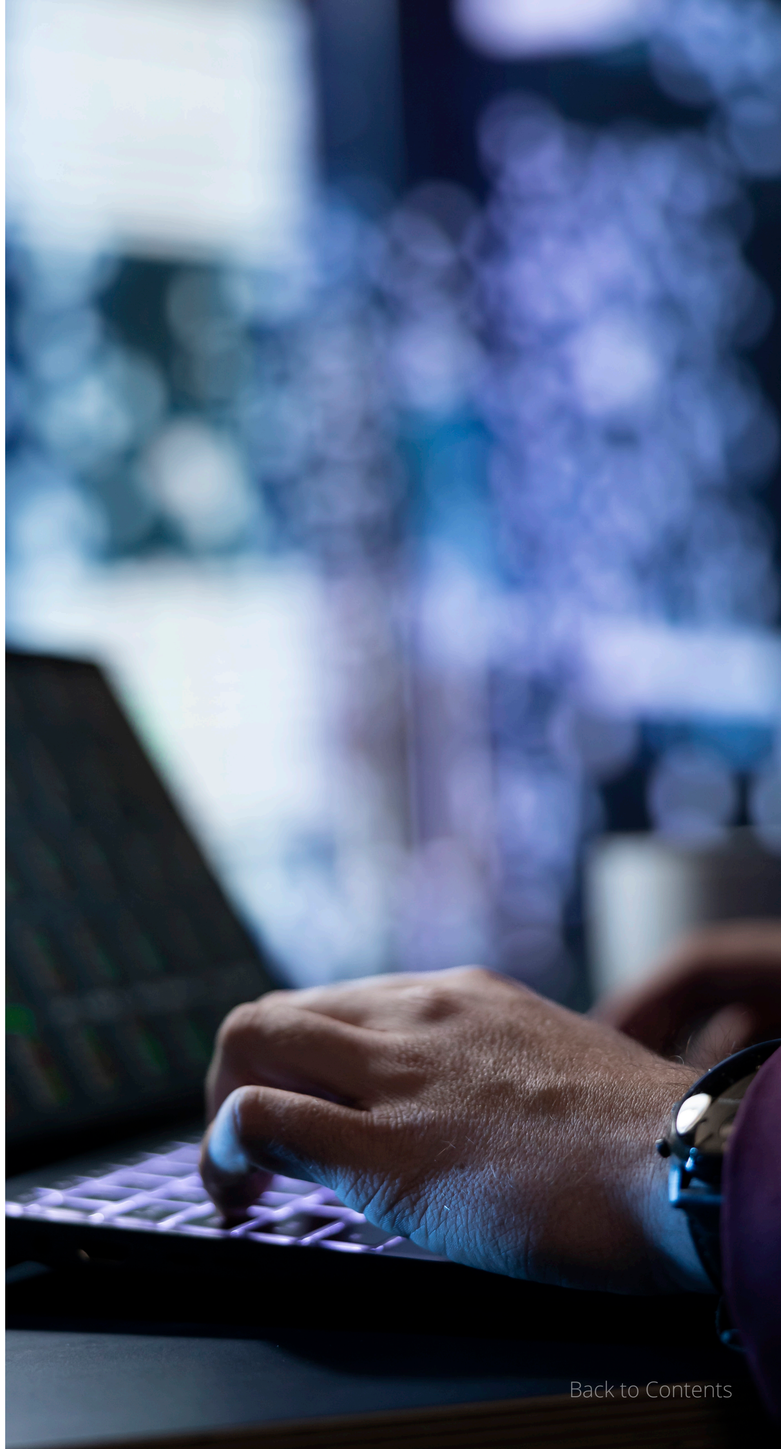

Back to Contents

RESEARCH METHODOLOGY AND DEMOGRAPHICS

To gather data for this report, Google commissioned Omdia to conduct a comprehensive online survey of security practitioners (i.e., SOC managers and staff) at enterprises with 1,000 or more employees. Organizations represented span both the private and public sectors and include respondents based in the U.S. (83%) and Canada (17%). The survey was fielded between October 21, 2025 and November 3, 2025.

All respondents were provided an incentive to complete the survey in the form of cash awards and/or cash equivalents. After filtering out unqualified respondents, removing duplicate responses, and screening the remaining completed responses (on a number of criteria) for data integrity, we were left with a final total sample of 300 qualified respondents. The margin of error for a sample of this size is + or – 6 percentage points. Totals in charts and graphs in this report may not add up to 100% due to rounding.

How many total employees does your organization have worldwide? (Percent of respondents, N=300)

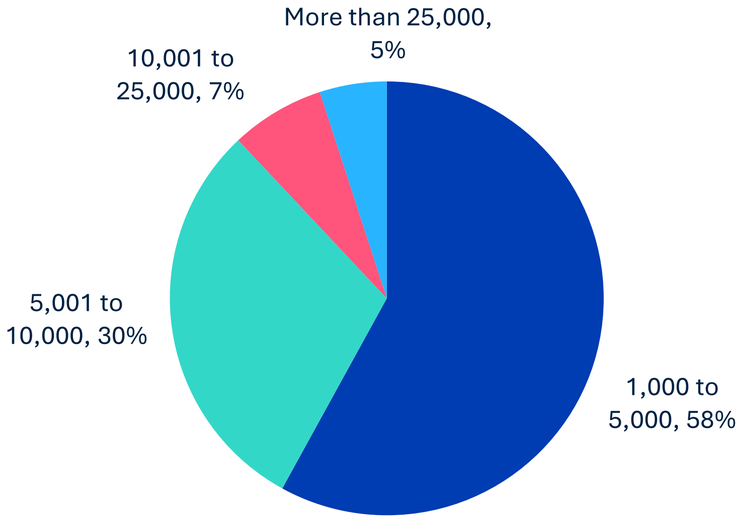

What is your company's primary industry? (Percent of respondents, N=300)

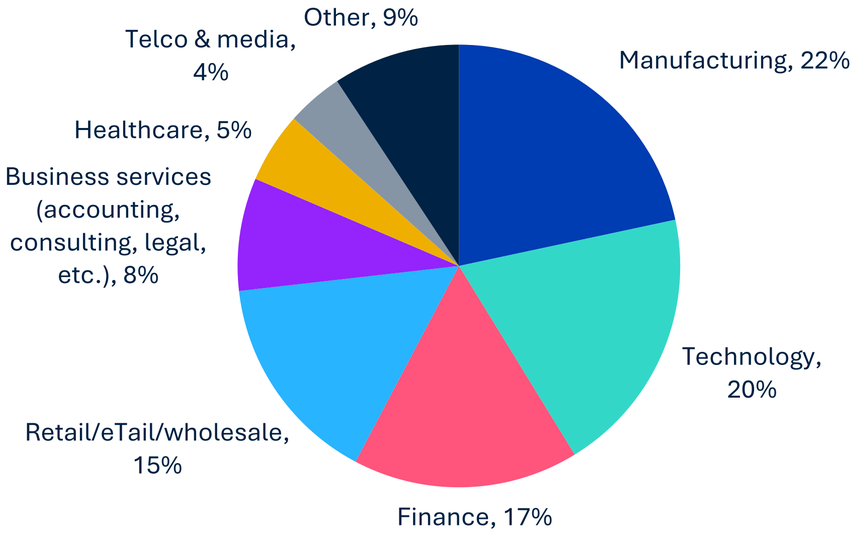

Which of the following best describes your current responsibility within your organization? If you don't see your exact role, select whichever is the best fit. (Percent of respondents, N=300)

Back to Contents

Omdia provides focused and actionable market intelligence, demand-side research, analyst advisory services, GTM strategy guidance, solution validations, and custom content supporting enterprise technology buying and selling.