VECTRA®

RESEARCH REPORT

2026 State of Threat Detection and Response:

Cyber Resilience in the AI Era

Executive summary

In this independent global research study, we surveyed 1,450 individuals involved in IT security with their organizations or who influence security decisions. The purpose of the research is to gain a deeper understanding of the challenges security teams face every day when detecting, investigating, and responding to cyber threats. To capture this perspective, responses were collected from security engineers, analysts, SOC leaders, CISOs, and other security team members working in organizations with at least 1,000 employees across North America, Europe, and APAC.

The data presented in this year's report provides insight into current practices for threat detection, investigation, and response. It examines where technology is helping defenders stop attacks, where it continues to fall short, and how teams are adapting to persistent challenges. This includes examining SOC workload, vendor trust, AI adoption, and the outcomes defenders hope to achieve. The findings also reference data from previous years of research, allowing us to track how challenges are evolving — whether improving, staying the same, or compounding further.

Introduction

What we conclude:

Three years of research into the state of threat detection and response raises a simple question: is cyberattack resilience getting better or worse? On one hand, 69% of practitioners now say they have the right number of analysts, and 80% believe their current tools provide adequate protection against attacks on hybrid, multi-cloud environments. On the other hand, defenders remain caught in the same cycle of noise, distraction, and doubt while the value outcomes of AI remain difficult to quantify.

Three challenges persist across defenders:

1. Hybrid, multi-cloud, AI environments are the norm, and defenders continue to struggle with too many siloed threat signals.
2. Alert volumes are declining, but threat detection latency isn't getting better. Teams spend the same number of hours triaging alerts and leave most unaddressed.
3. Fragmented observability and the inability to address threats at modern attack speed across modern environments overshadows progress.

Is this as good as it's going to get?

The evolution of AI-driven cybersecurity continues to progress towards meeting practitioners where challenges exist. Initially, AI tools were seen as supplemental in speeding up detection, however, today the SOC is at the threshold of an AI-led model, where systems and agents are trusted to take the first steps in detection, investigation, and response. While humans remain in control, defenders are actively engaged with AI agents and AI assistants as threat detection and response tools in their toolbox. Yet, with the promise and optimism around AI, it hasn't yet proven more resiliency.

Table of Contents

Key Findings

of defenders defend hybrid or cloud environments.

of defenders put aside important security tasks at least two days a week.

of security alerts received go unaddressed.

Defenders spend an average of 2.5 hours per day triaging alerts, while 41% of teams spend over 3 hours a day.

of defenders blame tools for ineffectiveness.

expect to use AI more next year, primarily to replace legacy detection and response tools.

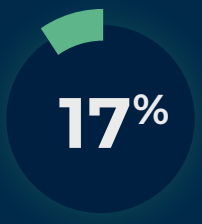

On average AI handles approximately 17% of defenders' workload.

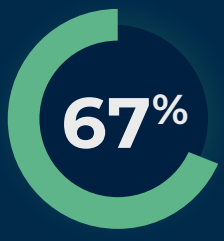

agree the implementation of AI-powered tools has had a positive impact on the ability to identify and deal with threats.

of defenders say they would like AI agents to handle alert triage and investigations.

Hybrid, multi-cloud is the norm, as is siloed visibility

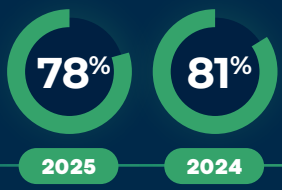

defend hybrid or multi-cloud environments

Visibility gaps remain in cloud networks

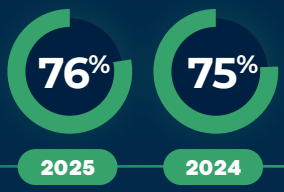

lack full visibility into cloud networks

SECTION 1

Hybrid environments, familiar challenges

Today's hybrid, multi-cloud environments are no longer considered "new" as nearly eight in ten defenders (78%) are tasked with defending them. Yet as defenders find themselves at the crux of defending these environments from attacks, they continue to find familiarity in many of the challenges that have persisted for years. Begging the question: Do defenders really understand who and what is on their network and if or when those behaviors introduce risk?

Visibility

Visibility across hybrid and multi-cloud environments continues to fall short and is uneven at best as roughly six in ten defenders say they have "full" or "almost full" visibility into core environments like endpoints (58%), on-premises networks (60%), cloud networks (60%), and identities (59%). The numbers go down from there depending on the surface:

How would you rate your visibility into the following environments?

ENVIRONMENT	% THAT CLAIM "FULL" OR ALMOST "FULLVISIBILITY"
Endpoints	**58%**
On-premises networks	**60%**
Cloud networks	**60%**
Identities	**59%**
SaaS	**56%**
IoT	**54%**
OT	**56%**
Vulnerabilities introduced by gen AI adoption	**49%**
Public cloud environments	**58%**

Even as defenders report having tools in place for visibility across each of their hybrid and multi-cloud surfaces, coverage doesn't guarantee a usable signal to effectively detect and prioritize malicious activity especially if visibility is fragmented. Additionally, the visibility they have into their hybrid networks largely remains the same year after year regardless of having tools to cover each environment.

YoY visibility comparison: Visibility into hybrid environments

ENVIRONMENT	2023	2024	2025
Endpoints	61%	61%	58%
On-premises networks	57%	62%	60%
Cloud networks	57%	60%	60%
Public cloud environments	54%	59%	58%
SaaS environments	55%	59%	56%
Identities (user + machine)	56%	59%	59%
IoT environments	54%	55%	54%
OT environments	54%	55%	56%
GenAI-related vulnerabilities	–	51%	49%

Compromise

It's a mixed bag in terms of gaining an attack signal that defenders can use to prioritize real threats. While defenders seem more confident that an attacker hasn't already compromised their environment than in the past, 46% agree that there's no point in constantly assessing security posture if they're already breached and just don't know about it yet.

2025		2024
37%	believe an attacker may already have compromised their organization without detection.	51%
44%	admit they are losing the battle when it comes to prioritizing real threats.	48%
47%	say it is not possible to deal with the increasing number of security threats.	52%

Noise

69% of defenders are still concerned about missing a true positive hidden in a flood of alerts at least once a week, which means not much has changed from a year ago when 71% of practitioners reported the same feeling. Are defenders just accepting alert noise as part of their environment or are there other indicators that signal change?

2025	2024
54% agree there's still too much noise, and it's only a matter of time before something slips through.	56%

Tools

Tool sprawl also shows little progress as stacks remain bloated even though 80% of defenders say their current suite of tools provides adequate protection against today's hybrid attack landscape — a figure that's up from 76% a year ago.

2025	2024
69% report using more than ten tools for detection and response.	71% had more than ten tools to detect, prioritize, and respond to threats.
39% juggle over 20 tools.	45% had over 20 tools.
56% lose hours every week jumping between tools.	57% lost hours jumping between tools.

Workload

71% of defenders say they put aside important security tasks at least two days per week to manage alerts and maintain existing tools — a statistic that shows little improvement from a year ago when 75% of defenders were having the same experience. In addition, more than half (55%) of defenders say that "more effective security tools" — more so than "hiring additional analysts" — would best ease their workload.

2025

71% of teams put aside important security tasks at least two days a week, while

45% of teams admit to putting aside important tasks at least 3 days per week.

2024

77% of teams put aside important security alerts at least two days per week.

43% did so at least three days a week.

What scenarios would best help ease the workload in your SOC?

	2024	2025
More effective security tools	55%	55%
More training and education	46%	47%
More time to respond to real threats	41%	43%
Outsourcing security operations	40%	40%
Hiring more analysts	36%	43%

Investments

Purchasing decisions aren't necessarily aligned with defender pain points. For example, 61% of defenders say a lot of their security tools are bought as a "box ticking" exercise for compliance with the percentage of defenders who say this is increasing each year. And while a "compliance" is cited as a top priority when evaluating solutions, defenders also want the ability to measure operational effectiveness. Are there other potential factors limiting the implementation of innovative tools such as those driven by AI from entering the SOC?

2025

61% say a lot of their security tools are purchased as a "box ticking" exercise for compliance.

2024

60%

of defenders say they wish other teams would consult with them before deciding to invest in a tool / technology.

say what really matters when evaluating solutions is risk reduction, compliance alignment, and measurable operational effectiveness.

What matters most when evaluating IT/security solutions?

Risk reduction and compliance alignment	**36%**
Operational efficiency and effectiveness	**36%**
Cost optimization and budget fit	**16%**
Vendor reputation and innovation	**11%**

Takeaway

Fragmented telemetry is split across tools and teams, resulting in defenders' inability to see the full picture. Defenders still face too much noise and not enough signal to prioritize attacks while visibility into identities shows little to no improvement. Teams put aside critical tasks to chase alerts, while "visibility" without actionable signal continues to create more noise and work. The question remains: have defenders accepted this as the current state of threat detection and response?

SECTION 2

Threat challenges

Whether it's putting aside important security tasks, concerns about missing critical alerts, or current tools not delivering the right signal for hybrid, multi-cloud environments — the data shows similar threat detection and response challenges as in past years with no clear core problem.

To what extent, if any, have the following negatively impacted your ability to identify and deal with threats?

	2025	2024
Size of attack surface	**30%**	**29%**
Number of attacker exploits	**34%**	**32%**
Number of unpatched vulnerabilities	**36%**	**35%**
Number of signatures and rules	**23%**	**22%**
Number of detection tools in place	**19%**	**20%**
Number of alerts	**33%**	**32%**

When looking at what defenders say is negatively impacting their ability to identify and deal with threats, there are slight differences from past years. However, the "number of detection tools" used are now viewed as less of a negative over the last two years compared to the first year of this research when 35% of respondents said their tools had a negative impact when it came to identifying and dealing with threats.

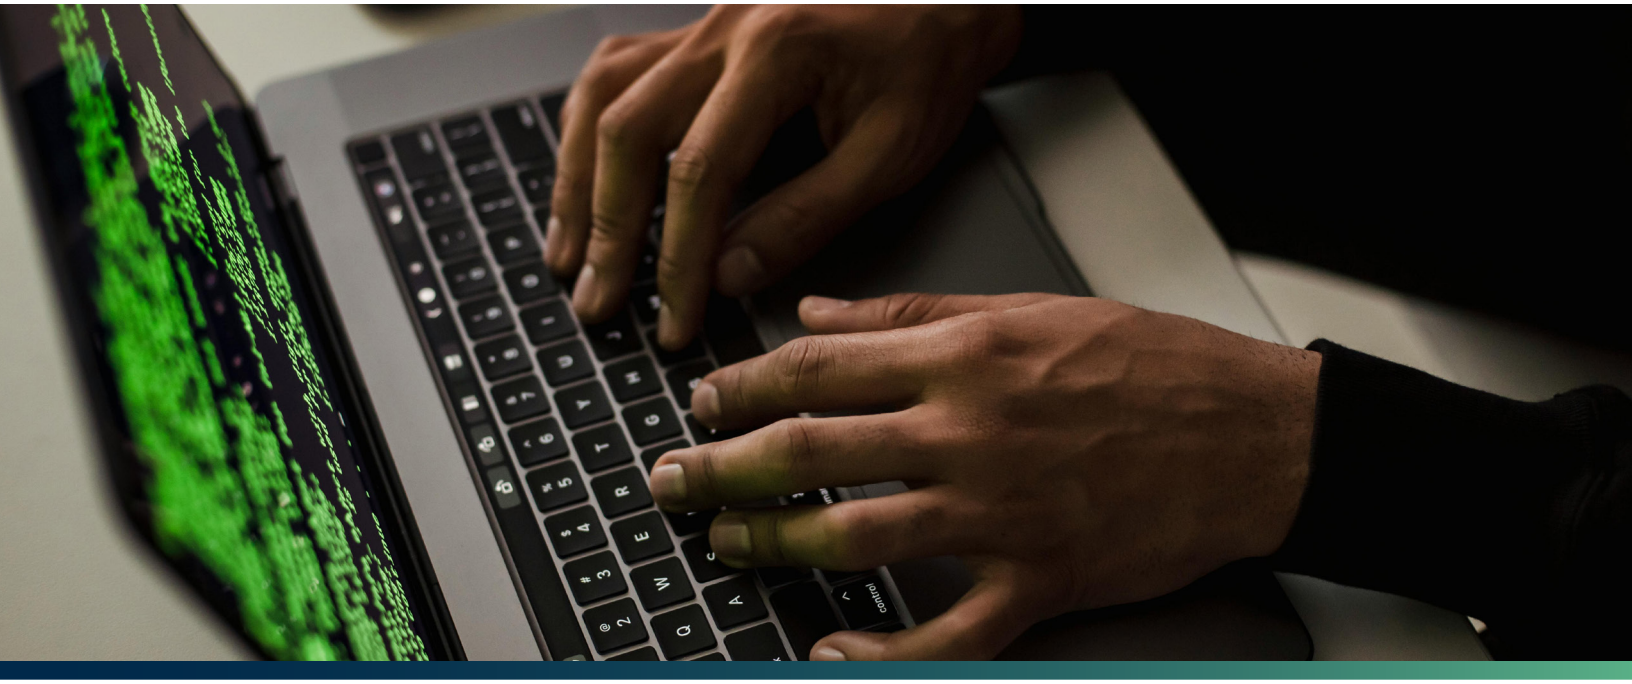

Threat detection

Defenders also report less overall alert volume from a year ago, however, they still feel they can't keep pace with alerts as the percentage of alerts they are able to deal with each day shows little improvement. In fact, **63%** of alerts go unaddressed — a figure largely unchanged — and practitioners continue to spend an average of **2.5 hours per day** triaging alerts, a figure that also lacks improvement.

What threat detection tools do you have in place that generate the most amount of alert noise?

Network threat detection	**43%**
Email threat detection	**40%**
Cloud threat detection	**39%**
Security incident and event management (SIEM)	**37%**
Identity threat detection	**35%**
Endpoint threat detection	**29%**
Intrusion detection	**29%**

	2025	2024
Typically, how many security alerts per day does your team receive?	2,992	3,832
Realistically, what percentage of these security alerts can you deal with per day?	36%	38%
What percentage of the security alerts you receive are "real attacks"?	14%	16%
How many hours per day do you spend digging through / triaging security events and alerts?	2.5 hours on average per day	2.5 hours on average per day

AI adoption

In terms of improvement, **67%** agree that the implementation of AI-powered tools are making a positive impact on the ability to identify and deal with threats (65% in 2024). This particular response was the most popular reason that defenders cited related to having a positive impact on their ability to identify and deal with threats. Overall, defenders aren't shying away from AI in the SOC as the technology continues to take a bigger role in threat detection, investigation, and response. In fact, 87% say they expect to use more AI-powered tools to replace legacy detection and response tools.

85% say their use of AI security tools has increased over the past year.

87% expect to use more next year, primarily to replace legacy detection and response tools.

75% report that the number of AI-powered tools they use has grown significantly over the past year.

67% of defenders agree that the implementation of AI-powered tools has made a positive impact on their ability to identify and deal with threats.

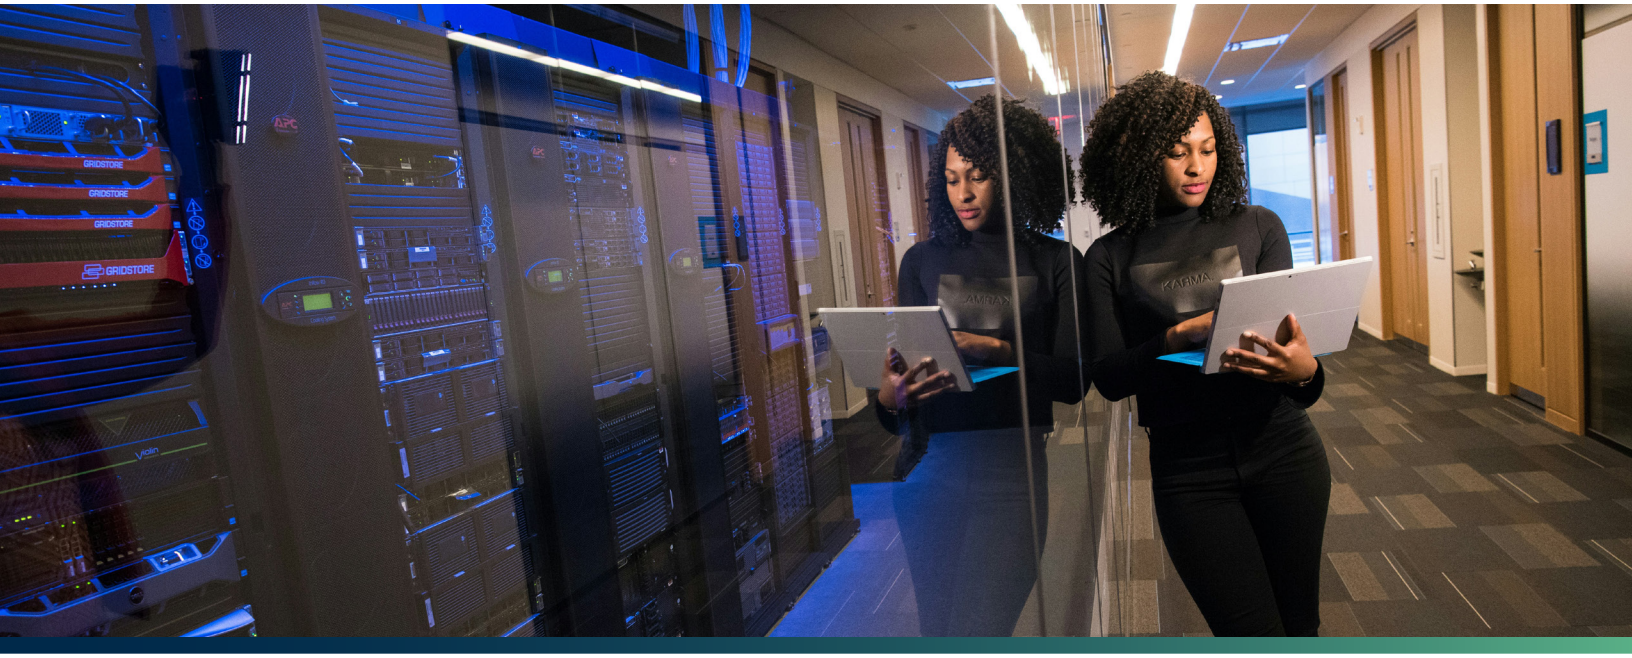

AI usage

With defenders now heavily adopting AI, and expressing its positive impact, how embedded is AI in today's SOC workflow? Last year's report found that **97%** of defenders had already adopted AI, and **85%** said their investment and use of AI had grown in the prior 12 months. That trend continues again this year as **85%** say their usage of AI tools has increased, and there are now more details about how much of the SOC workload and day-to-day tasks are being covered by AI.

What percentage of your workload is handled by AI (AI assistants and/or AI agents)?

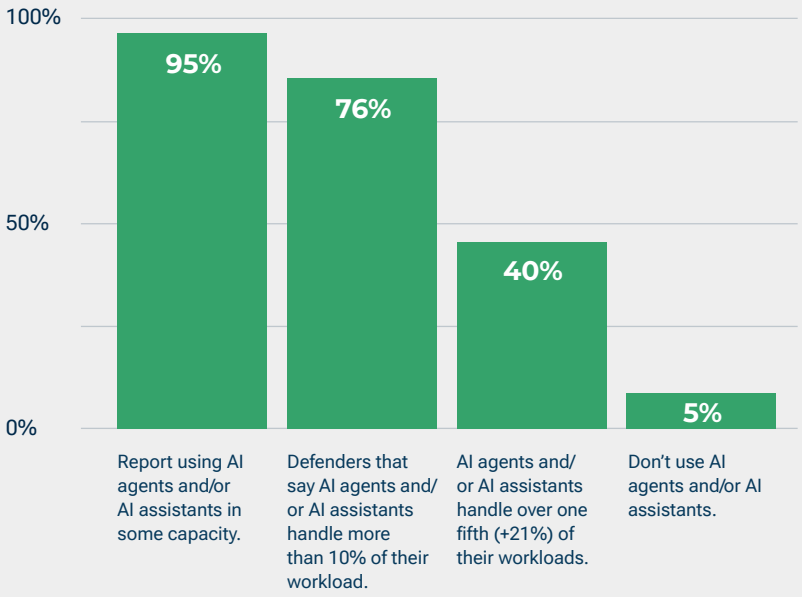

AI impact

Defenders are also more clear about what they want from AI, which falls in line with how it's helping them "identify and deal with threats." When asked which benefits matter most, they point directly to the areas that would help acquire a more accurate attack signal, such as improvements in "accelerating detection and response," and "uncover unknown threats."

What are the most important benefits AI-powered security tools offer to SOC teams?

AI value

The themes of identifying, dealing with, and prioritizing threats continues when defenders reported the parts of their job they would like outsourced to AI agents the most? "Alert triage" recorded the highest percentage of responses followed by "investigations," and then "reporting." Why do defenders want to turn over triage duties to AI? Defenders have a lot of tools: 69% use more than 10. Many lose hours jumping between them: 56% claim that to be the case. And they want help with threat prioritization as 69% of teams are concerned about missing a true positive hidden in a flood of alerts.

What part of your job do you most want AI agents to do?

Alert triage	**32%**
Investigations	**23%**
Reporting	**22%**
Response	**19%**
None of the above "I don't want AI agents to do any part of my job"	**3%**

Takeaway

Alert volumes are trending down, but defenders aren't feeling the difference. Defenders still leave most alerts unaddressed, spend hours every day on triage, and remain stuck in a reactive cycle where detection latency persists. The good news is that defenders recognize that AI tools can help them move beyond manual triage duties and they are optimistic about its ability to identify and deal with threats. The question is whether this optimism — combined with heavy AI adoption — signals the future SOC: one where latency is reduced as defenders focus on improvements that increase resilience?

SECTION 3

Is more trust in AI improving trust in vendors?

Defenders see the potential for AI to improve signal and are ready to lean on AI agents to handle triage duties and help prioritize the alerts that represent malicious behavior. What are the defenders who already use AI saying about its current impact in the SOC?

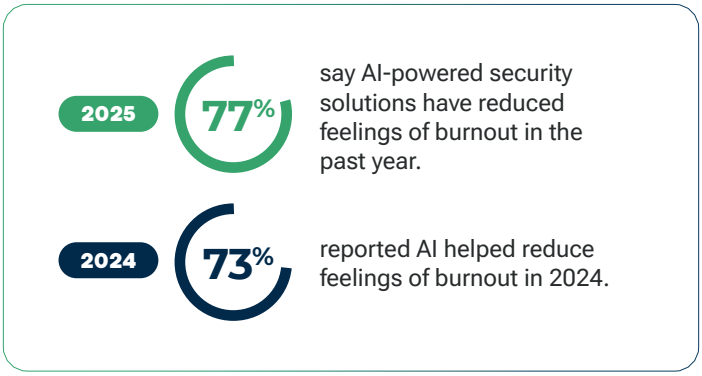

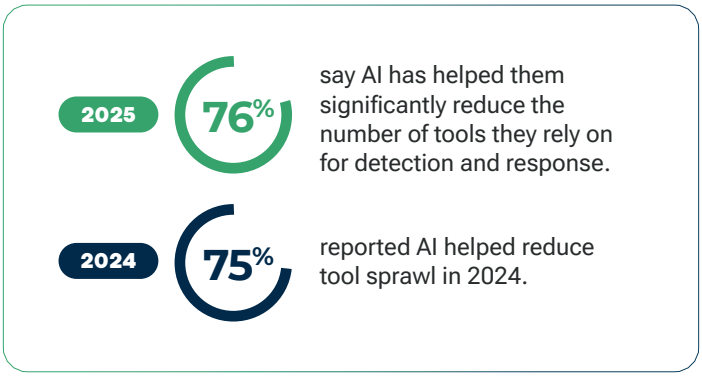

While over three quarters of defenders say AI is helping reduce the number of tools for threat detection and response, it's a slow progression with only a 6% decline in teams using over 20 tools and a 2% decline in teams using over ten tools.

AI agent needs

Defenders also acknowledge the potential that AI agents offer to help free up time for more fulfilling tasks. Instead of spending hours on alert triage, security professionals express that having more time to explore new tech, develop their skills, or threat hunt would be more fulfilling than constantly managing tools and alerts.

If AI agents freed up your time for more meaningful, enjoyable work, what is the top activity you would choose to spend it on?

28%	would spend time exploring emerging technologies.
23%	would invest in personal growth and development.
20%	would dedicate more effort to threat hunting.
18%	would spend time on deeper threat research.
11%	would focus on mentoring others.

AI agent concerns

While defenders still cite a crowded market, concerns about cost and uncertainty around which AI agents will add value as reasons for why AI adoption isn't happening — a fewer percentage of defenders expressed these concerns than they did a year ago.

Why is deployment and implementation of AI agents not happening in your organization?

2025	
22%	say there are too many AI agents on the market to tell which will add real value.
19%	worry AI agents could create more work instead of reduce it.
18%	say AI agents are too expensive.
15%	fear AI could render their jobs obsolete.

2024	
46%	were concerned AI agents will create more work opposed to reduce it.
38%	say there are too many AI agents on the market, making it difficult to determine which will add real value.
22%	said AI agents are too expensive.
16%	were concerned AI agents will render their jobs obsolete.

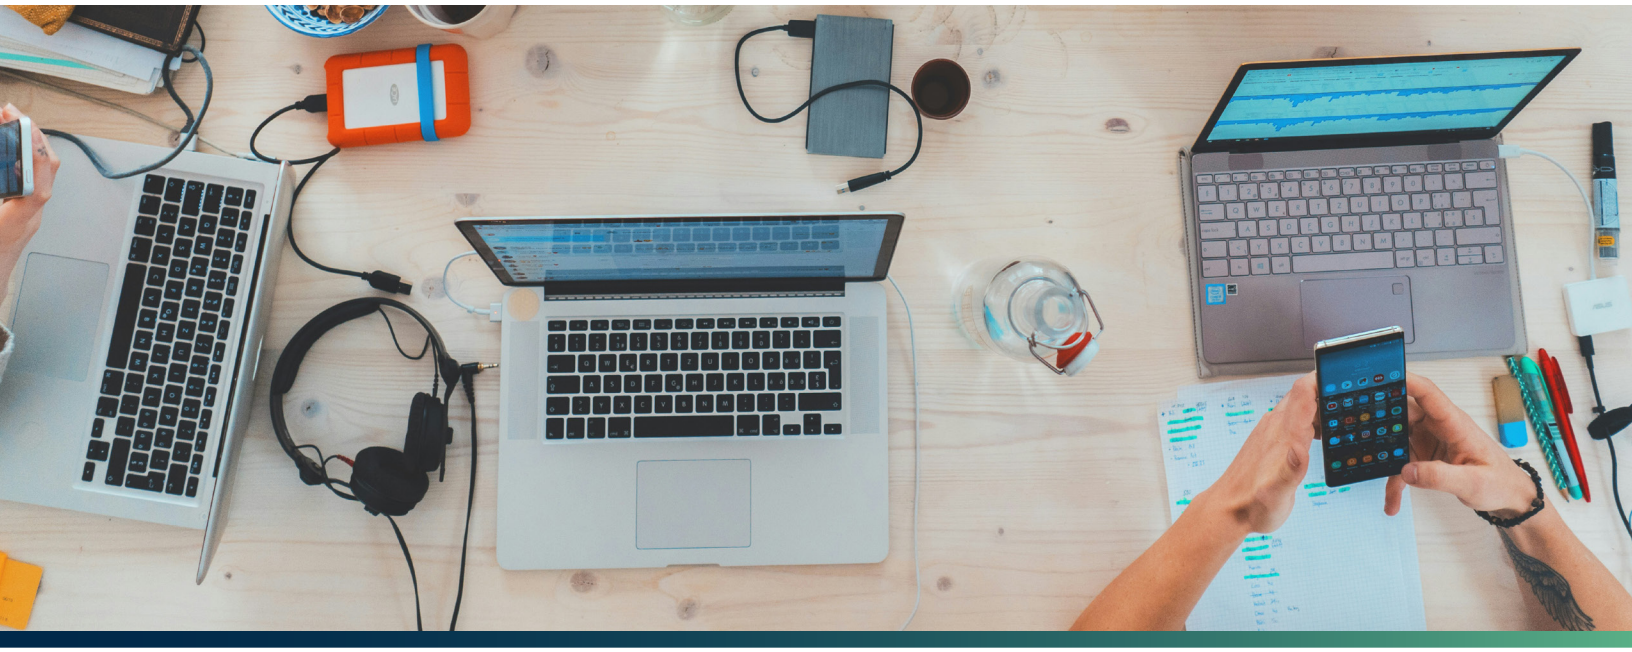

VECTRA

AI vendor trust

Despite an increase in confidence compared to a year ago that AI will add value and reduce workloads, defenders still express a similar sentiment towards the overall trust they have in vendors to deliver on their promises. For example, 48% of practitioners say that current tools are more "hindrance than help" in spotting real attacks (down 1% from 2024). Defenders still feel security vendors avoid responsibility, and well over half of the respondents say they remain frustrated with empty promises.

2025		2024
59%	feel vendors flood them with pointless alerts to avoid responsibility for breaches.	62%
58%	say vendors sell threat detection tools that create too much noise and too many alerts.	60%
74%	want vendors to take more responsibility when breaches occur.	71%
51%	say their tools increase workload rather than reduce it.	54%
58%	are frustrated with vendor's "empty promises" and products requiring constant tuning.	59%
48%	admit they don't trust security tools to work the way they need them to work.	46%

On one hand defenders see the potential of AI — especially as their usage grows. On the other, security vendors haven't improved their perception due to pointless alerts and increasing workloads related to tools not doing what's needed in the SOC. And as the data shows, defenders aren't just purchasing based on innovation alone — they want to see demonstrated impact and need to make sure compliance is met. Whether resiliency is being limited by the attack signal from current tools or if defenders are deciding on a new path to acquire better telemetry, it's the defenders who bear the burden of sorting out how best to move forward.

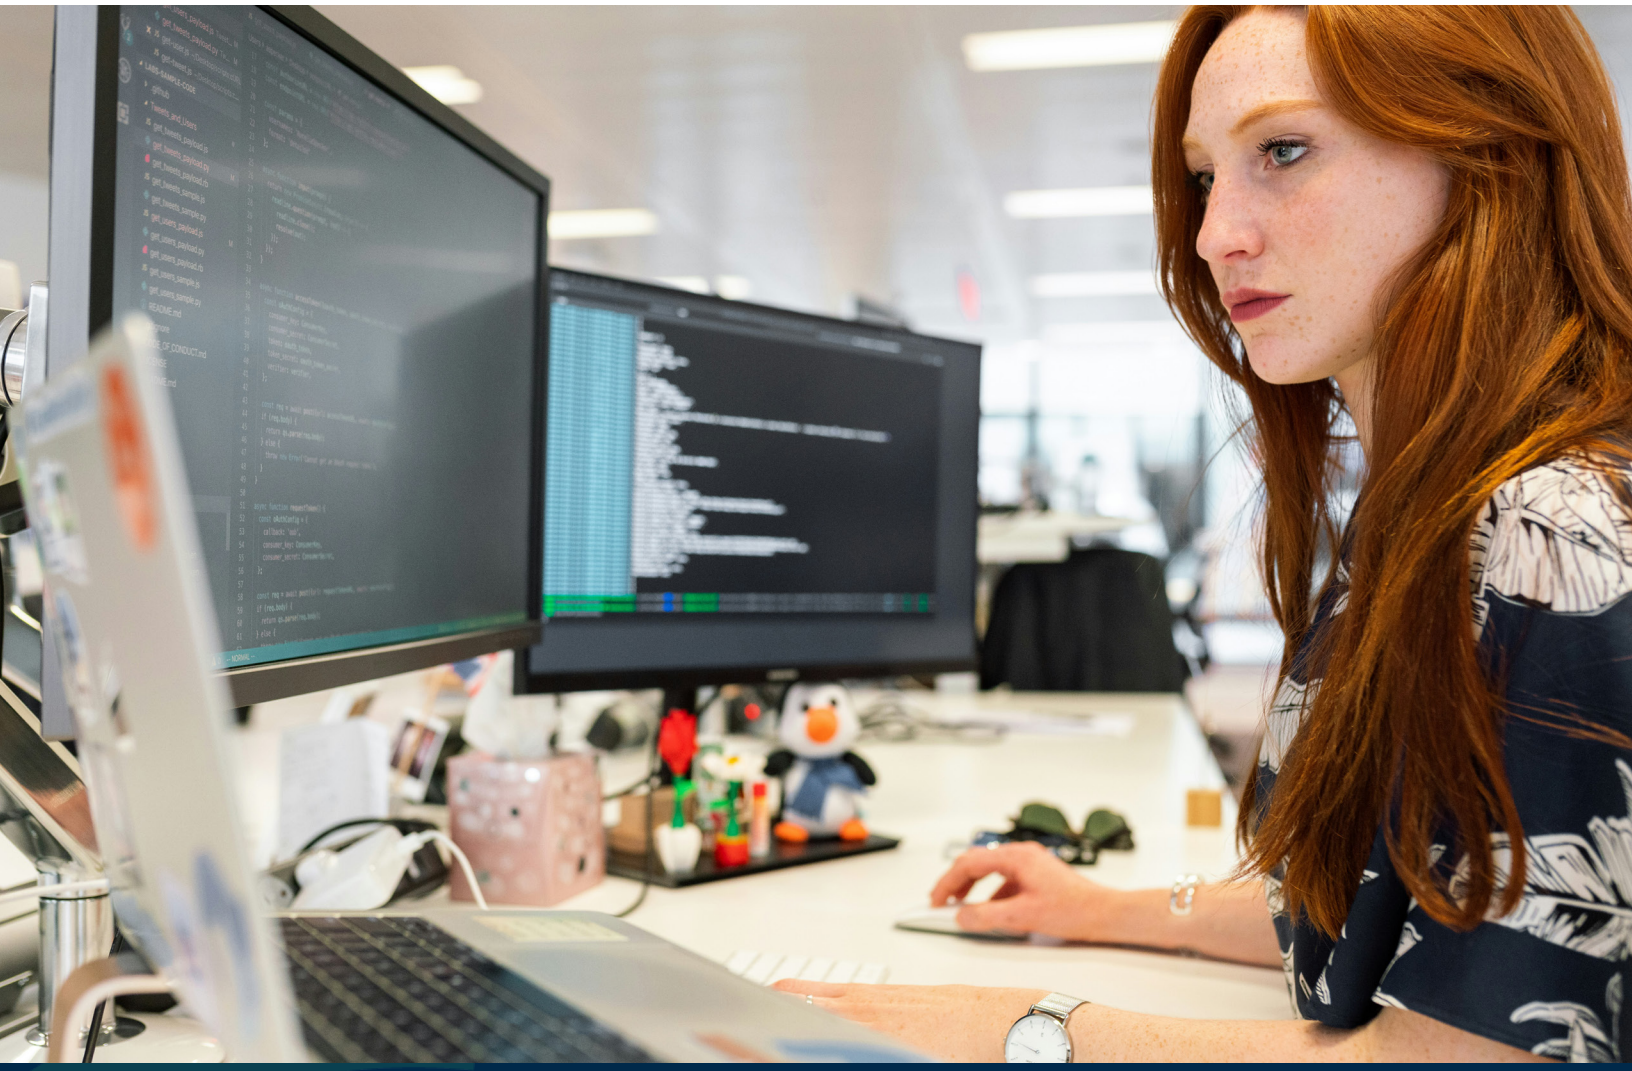

Takeaway

Defenders continue to see the potential of AI to improve signal and reduce workload and burnout, and their adoption and usage grows stronger each year. This momentum points to a possible shift — one where defenders rely more on AI to understand what's happening across the hybrid enterprise, gain deeper insight into identity behaviors, and act before impact occurs. Yet even with the potential of AI tools to help teams add resilience, trust in vendors remains stalled. Current tools still increase workload, operate in siloes, flood teams with alerts, and fall short of promised outcomes, with year-over-year data showing little meaningful improvement where it matters most.

Conclusion

Looking across the last three years of research, one theme is clear: despite incremental improvements in visibility and staffing, defenders remain stuck in the same cycle of noise and uncertainty. Many of the tools defenders use across hybrid, multi-cloud environments struggle to translate coverage and visibility into a usable signal. Alert volumes are trending down, yet defenders spend the same hours triaging while leaving most alerts unaddressed. And while confidence has ticked up — with more defenders saying they have enough people and that their current tools provide "adequate protection" — the core challenges haven't gone away.

This year's findings reinforce what defenders already know: visibility without signal is just noise, and vendor trust has not improved in this area. At the heart of this is a defender's mission: to keep their people and brand safe and stop breaches. However, true cyber resilience can't be achieved if uncertainty continues to exist. Defenders need to see all activity across their environment and understand which activity matters so they can do want they want to do — protect the organization from attacks.

As for AI, adoption and use continues to grow, with defenders clear about the benefits they want most: faster detection and response, sharper prioritization of real threats, and relief from the endless churn of alert triage. This signals a path forward, not as a supplemental tool, but as the foundation for the attack signal defenders have been missing.

The question is: what will it take for organizations to be resilient in the AI era? Endpoint, identity, and cloud controls are still a necessity, but to move at the speed of modern AI-driven attacks and be resilient to them, AI will need to play bigger role in helping defenders connect the dots about what's happening across their entire hybrid enviornment. Defenders know where AI can help and are eager to lean on it, but they need clarity and confidence that these solutions will deliver measurable, defensible outcomes.

Recommendations

Vectra AI recommendations based on 2026 findings

1

Build resilience through continuous observability

Building resilience at AI speed and scale requires modern network observability that continuously understands the identities on a network and how they are behaving. In today's modern hybrid environments, identities connect everything: people, services, workloads, APIs, and AI systems. When defenders have continuous observability to identity behaviors across the modern network, risk turns into insight, complexity turns into context, and resilience stops being an abstract goal and starts becoming an operating reality.

2

Redefine hybrid, multi-cloud defense in the AI era

Prevention remains essential, but it is no longer sufficient once attackers are inside the enterprise. It's easy to think in terms of separate environments and tools, but attackers operate across a single, interconnected attack surface. In the AI era, hybrid defense must reflect that reality. Tools built for static environments struggle to provide consistent visibility across modern enterprises. They introduce latency through manual processes, fragment context, and leave defenders unable to confidently assess whether resilience is improving. When security operates in silos, defenders lose the ability to track attacker movement as it unfolds. Effective hybrid defense requires a unified, AI-driven approach that brings together activity across the enterprise, prioritizes what matters in real time, and exposes malicious behavior wherever it occurs.

3

Use AI to transform, not just accelerate security operations

AI is no longer just a time-saving add-on; it is becoming the foundation for building resilience in hybrid enterprises. Defenders are clear about what they want from AI and want to be able to focus on higher-value work. Defenders who stay ahead will be the ones who continue to embrace AI as the path forward. Modern attackers are actively using AI to remove latency from their own operations and defenders must do the same. When AI is applied to both network and identity signals, complexity becomes clearer and exposure more measurable, enabling the possibility for earlier detection of risky behavior, faster and more confident response, and meaningful reduction in noise and investigation delays.

About Vectra AI

Vectra AI is the cybersecurity AI leader in protecting modern networks from modern attacks. From on-premises data centers to multi-cloud, identity, SaaS, IoT/OT, edge, and AI infrastructure, the Vectra AI Platform empowers security teams with the modern network observability, signal, and actions needed to preemptively reduce attack exposure, proactively contain active attacks in progress, and automate security operations to prove resilience in an always-on, AI-powered world. As the leader in Network Detection and Response and with 35 patents in cybersecurity AI, modern enterprises across the world trust Vectra AI to protect their modern network from modern attacks.

Methodology

This report is based on a 2025 study commissioned by Vectra AI and carried out by Sapio Research. The study was conducted among 1,450 individuals involved in IT security with their organizations or who influence decisions on IT security, working in organizations with at least 1,000 employees and based in North America (500), Europe (750), and APAC (200).

For more information please contact us:
Email: info@vectra.ai | vectra.ai