DARKREADING
RESEARCH

MARCH 2025

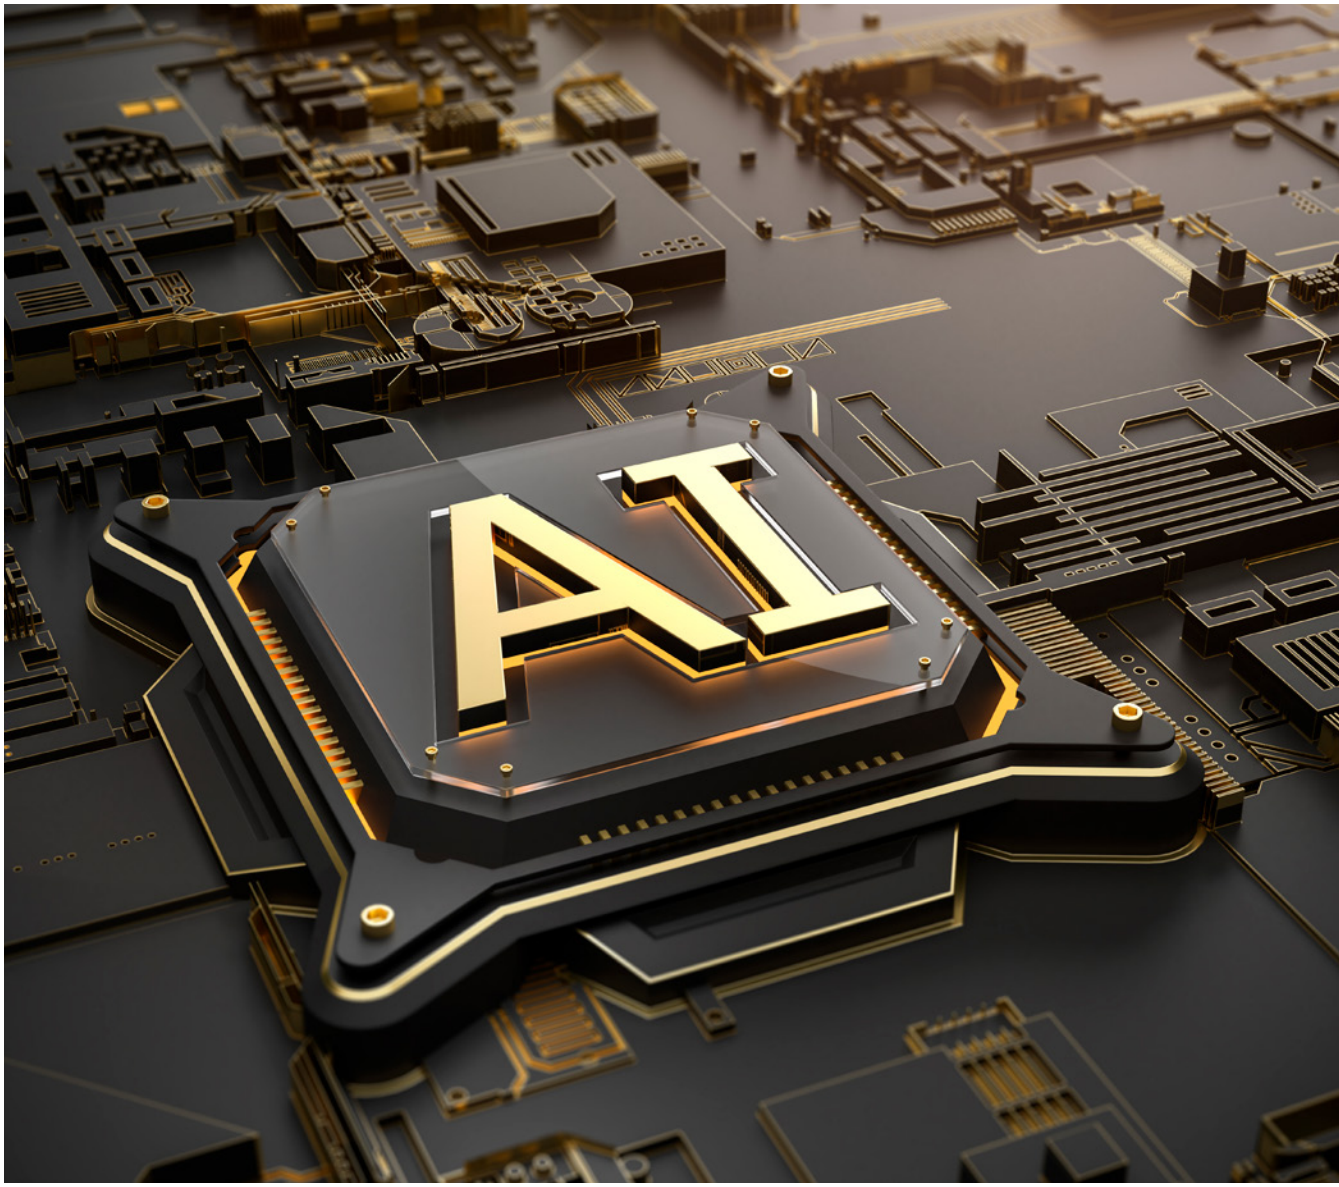

The Rise of AI-Powered Vulnerability Management

State of AI Adoption and Attitudes in Cybersecurity

seemplicity

Executive Summary

For many years, the hype of AI has overshadowed the actual benefits due to understaffed security teams. They've historically struggled to effectively apply AI-driven automation that demanded a lot of tuning and delivered spotty results with too many false positives and negatives. These problems haven't evaporated. But as AI technology grows more sophisticated, more reliable, and easier to use, the AI tide is shifting for security. The pros are increasingly outweighing the cons not just in threat detection, but also key security practice areas like vulnerability management, incident response, data protection, and IAM.

In this State of AI Adoption Survey, Dark Reading explored the prevalence and perceived benefits of AI in cybersecurity today.

Some of the highlights from the report show that:

- 86% of security teams today utilize some type of AI within their security tool stack
- 56% of security teams say the use of AI has become crucial to their team's operations
- 46% of security teams primarily depend on AI that is embedded in their security tools and delivered by their vendors versus building their own
- The top three most common security use cases for AI are endpoint security, basic vulnerability scanning, and antivirus/anti-malware
- 46% of firms say that they're actively trying to use AI to solve false positive issues, and 39% are using it to solve data overload problems that stymie vulnerability and exposure management work
- The top use case where security leaders say AI will offer most value is vulnerability and risk management, named by 74% of respondents
- The No. 1 security issue respondents are most hopeful that AI will help fix is the prioritization of disparate results from scanning tools, for which 82% are hopeful for gains
- Maintaining skilled security workforces that understand both security and data science will be an increasing concern — a lack of skills is the No. 1 named obstacle to effective use of AI today

Clearly, organizations see vulnerability and exposure management as the most promising use case for the next few years to advance applied AI in security work. They're especially hopeful that AI can help them solve the problem of data overload that has for many years weighed heavily on the security world.

AI Prevalence in Today's Security Environments

One of the clearest trends that surfaced from this study is that AI has reached critical mass in the security world today. The vast majority of security teams today — 86% — utilize some type of AI within their security tool stacks. These teams use it embedded in commercial security tools, apply it through custom tooling that's powered by internal data science work, or use some combination of homegrown and commercially developed AI to power their security work.

Breaking down those usage patterns, the most common way to leverage AI is to let security vendors handle the heavy lifting. Some 46% of organizations say that's primarily the model they depend on (**Figure 1**). Meanwhile, approximately 19% say they primarily apply AI to security through their own internal data science work. And another 21% say they depend on an even mix of both. Among those that use both internally developed and externally developed AI security tooling, the use cases are more likely to skew toward vendor-led AI rather than the other way around. This is largely because the people using in-house developed AI tend to do so to fill in gaps within their security stack or to experiment. Approximately 42% of these teams say they use both equally. Another 21% say that internally developed AI-driven security tools just fill in the gaps left by their security vendors. And 16% say that their usage of internally developed AI is strictly experimental.

Applying AI to security use cases usually requires a unique combination of skills that tightly marries data science with security expertise. Given the already dire shortage of security rock stars in the security employment pool, this kind of resume will remain rare for

Figure 1

AI USE WITHIN SECURITY TOOL STACK

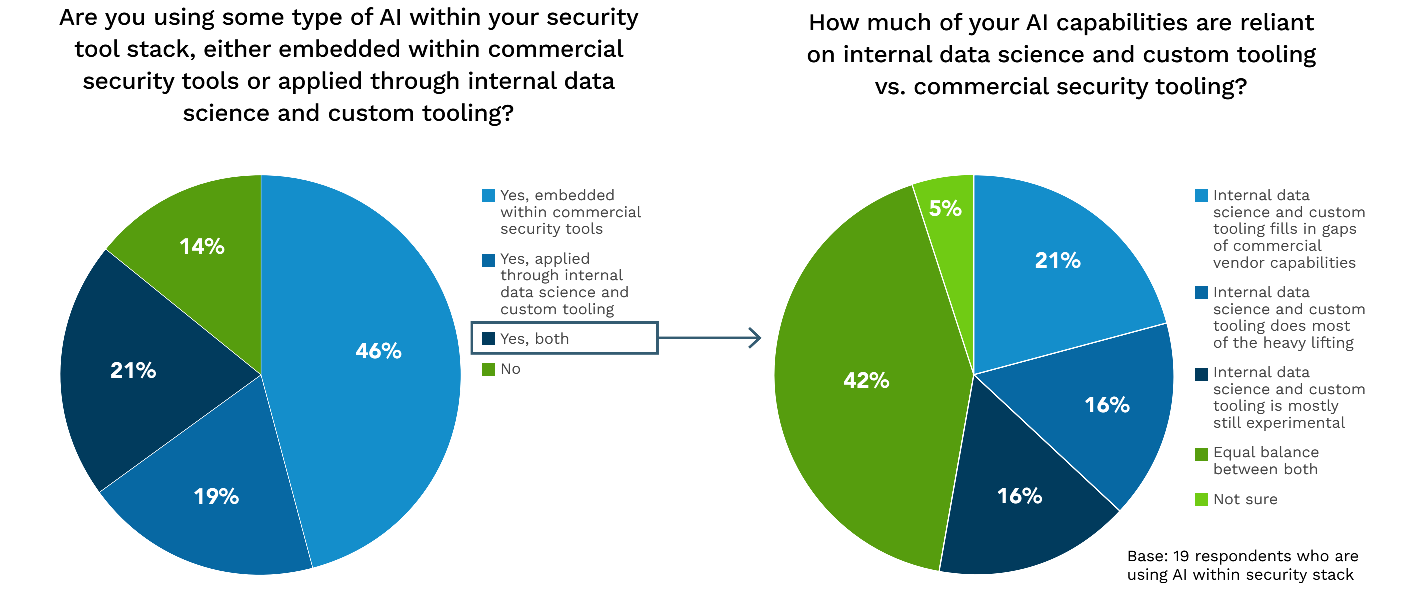

Data: Dark Reading survey of 94 Cybersecurity and IT professionals at large companies, November 2024

some time. The majority of companies that apply internal AI mechanisms to security problems are able to do this because they have bolstered their security workforce with data science expertise. Among those security departments that do their own custom AI work, 66% report that they have hired their own internal data science staff within their security teams. This comes as no surprise considering that homegrown AI takes significant investment in data science and model development to get right. This talent requirement will likely be one of the major reasons why security teams will continue to emphasize AI modeled and governed by security vendors rather than their internal teams.

In a minute, we'll dig into the mixed perceptions about the benefits and challenges that go along with increased AI use, but one thing is certain. The industry has reached an inflection point where — warts and all — over half of organizations say that at least some use of AI has become crucial to their security team's operations.

Applying AI to Security Pain Points

Security teams turn toward AI because they're overwhelmed by threats and vulnerabilities. When asked about the biggest security pain points they face today, far and away the sophisticated threat landscape was the most commonly cited point of friction, named by 60% of respondents (**Figure 2**). Inherently implied within this are issues around prioritizing the vulnerabilities that are most likely to be attacked by the bad guys — a difficult task given not only the sophistication of threats, but the sheer volume of automated attacks tuned to new and existing system flaws.

Interestingly, AI stands to both help enhance these points and exacerbate them as well. Number two on the pain points named by our respondents is AI and generative AI concerns. Approximately 1 in 4 organizations said they're concerned about how AI use in the enterprise will make them more attackable. As more

Figure 2

SECURITY PAIN POINTS

What are your biggest security pain points today?

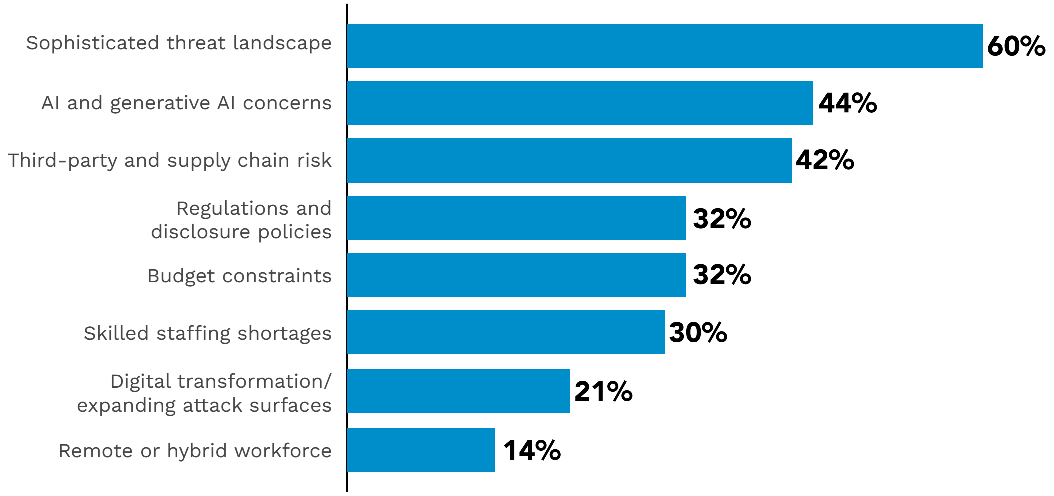

Note: Maximum of three responses allowed
Data: Dark Reading survey of 94 Cybersecurity and IT professionals at large companies, November 2024

AI-driven systems are used in broader IT applications this expands the attack surface of all software. New AI flaws will arise and further aggravate existing vulnerability management woes. Third on the biggest pain points was third-party and supply chain risk (42%). The complicated mesh of software dependencies in the enterprise has made vulnerability and threat prioritization orders of magnitude more complicated compared to even four or five years ago.

When asked how respondents are currently applying AI in their security tech stack today however, the use cases still remain fairly simplistic. The top three use cases were endpoint security (52%), basic vulnerability scanning (47%), and antivirus/anti-malware (40%) (**Figure 3**). These are all use cases for which simple machine learning has been applied for many years now, so the promises of more sophisticated AI-driven security automation have not yet been fulfilled within most environments. This isn't terribly surprising given the need for skepticism and caution in cybersecurity.

More sophisticated uses of AI for things like streamlining workflows, summarizing overwhelming heaps of security data, or prioritizing work are still relatively rare among security programs. For example, just 25% of teams use AI to power vulnerability prioritization, and only 18% use it to bolster vulnerability remediation workflows. Only 21% use it to automate away configuration management and system hardening and just 18% have utilized GenAI to speed up summarization and reporting work. As organizations develop greater comfort levels with AI, these use cases are all likely to see greater adoption.

Vulnerability Management Holds Most Promise

Given the relatively simplistic mix of AI use cases in place today, it should come as no surprise that only about 16% of security teams say their use of AI has been very beneficial and have made it a core part of their program. At the same time, there's very little antagonistic perception of AI. A scant 6% reported that it's

Figure 3

APPLYING AI IN SECURITY TECH STACK

Where are you currently applying AI in your security tech stack today?

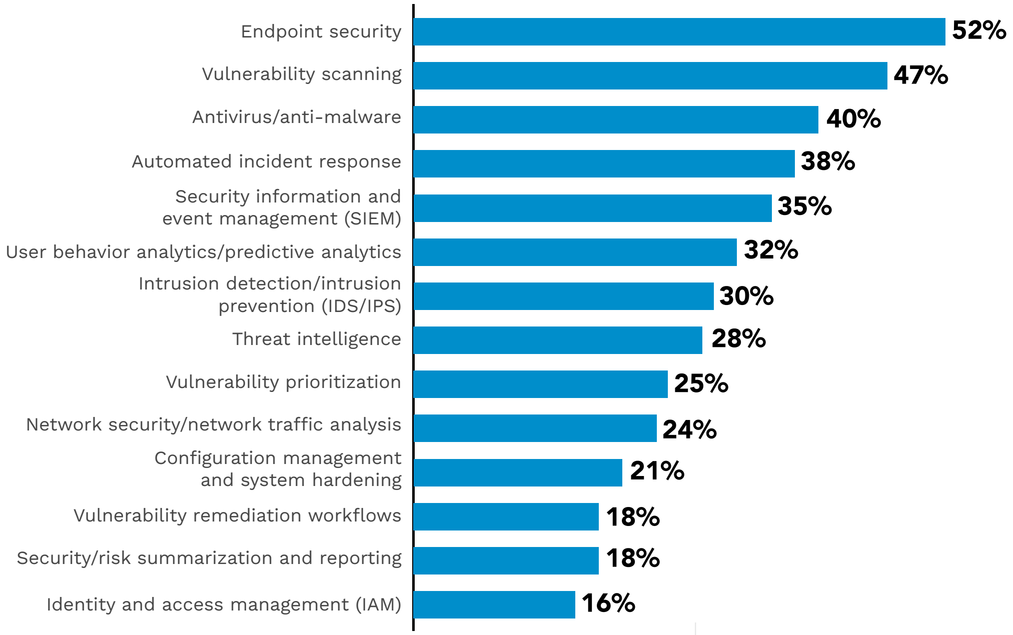

Note: Multiple responses allowed
Data: Dark Reading survey of 94 Cybersecurity and IT professionals at large companies, November 2024

detrimental to their security program.
For the most part, while AI is highly prevalent in security practices, most organizations are still realistic about the benefits they're deriving from its use. Around 45% say that it's moderately beneficial and they're starting to note the benefits. Another third are still evaluating its impact.

The perception among security pros is that the most effective uses of AI have been around threat detection, endpoint security, vulnerability assessment, and malware detection, around which at least 65% of respondents voted it moderately to very effective.

And the survey indicates a sea change coming where security teams are trying to enhance the sophistication with which they apply AI to their pain points. When asked about which security function they think AI will provide the most value in the next three years, vulnerability and risk management was head and shoulders above all of the rest, named by 74% of respondents. Another 59% cited incident response, which was No. 2 on that list. Data protection, identity and access management, and security reporting rounded out the top five on the list (**Figure 4**).

Clearly, organizations see vulnerability and exposure management as the most promising use case for the next few years to advance applied AI in security work. They're especially hopeful that AI can help them solve the problem of data overload that has for many years weighed heavily on the security world.

Figure 4

VALUE PROVIDED BY AI

Which security functions do you believe AI will provide the most value to in the next 3 years?

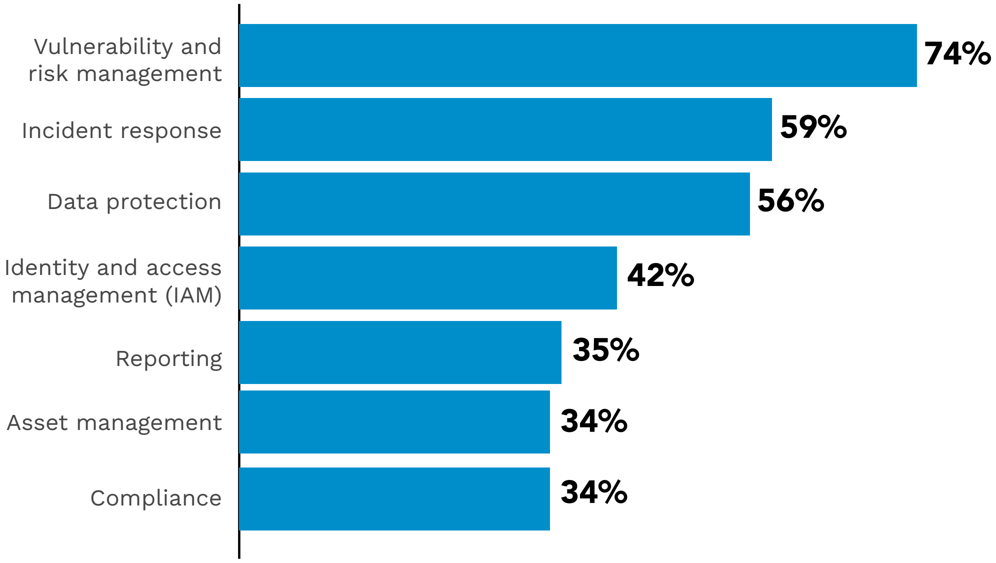

Note: Multiple responses allowed
Data: Dark Reading survey of 94 Cybersecurity and IT professionals at large companies, November 2024

Figure 5

USING AI TO SOLVE VULNERABILITY PROBLEMS

Which vulnerability and exposure management problems are you actively trying to solve with AI today?

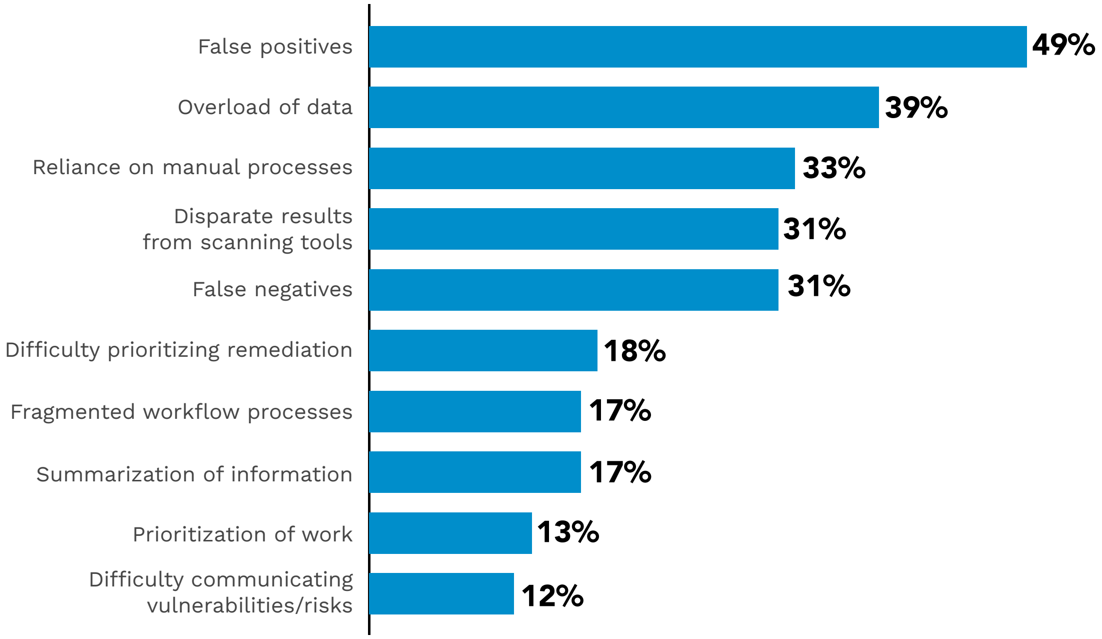

Note: Maximum of three responses allowed
Data: Dark Reading survey of 94 Cybersecurity and IT professionals at large companies, November 2024

The top five vulnerability management problems they're actively trying to solve with AI today were: false positives (49%), overload of data (39%), reliance on manual processes (33%), disparate results from scanning tools (31%), and false negatives (31%) (**Figure 5**).

The two problems that organizations are most hopeful for AI applicability were sifting through disparate results from scanning tools (82% optimistic) and dealing with overload of data (81% optimistic).

As organizations seek to ameliorate these issues and as AI tech advances within security tooling, the mix of use cases will likely change significantly. These results indicate the usage to be more heavily weighted toward vulnerability prioritization and workflow management in the coming years versus the traditional AI use in endpoint security and basic vulnerability scanning.

AI Hype and Efficacy

The survey shows that one of the big roadblocks hampering more widespread AI use and more meaningful value derived from its use is the fact that many security tools with embedded AI are underdelivering on the promise of their marketed tech advances. Approximately 56% of respondents reported that at least half of their security vendors tout their AI capabilities, with 1 in 5 reporting that 75% or more of their tool stack promotes AI capabilities. However, 77% of respondents reported that one or more of those vendors had overhyped their AI performance or are underdelivering on their promises.

The most overhyped capabilities are those that are currently also the most used: endpoint security, antivirus/anti-malware, and malware analysis (**Figure 6**). The big issue is that over half of organizations report that they're running into data quality problems and inaccuracy in their AI generated results (**Figure 7**). This could offer hints as to why AI deployments still remain fairly simplistic — if the most basic use of AI is overhyped some security teams may be turned off from further expanding their use of AI into more complex use cases.

One respondent was particularly detailed with their complaints, explaining that false positives are still the norm with a lot of AI-generated results:

> ***"One of the biggest challenges we face in applying AI to our security practices is managing false positives. While AI can be incredibly powerful in detecting potential threats, it often flags benign activities as suspicious. This leads to a lot of noise, which can overwhelm our security team and divert attention from genuine threats. The issue here is twofold: It not only increases the workload for our team, requiring them to sift through numerous alerts, but it also risks desensitizing them to alerts over time. If the team starts to see too many false alarms, there's a danger they might miss or underestimate a real threat. Balancing the sensitivity of AI systems to minimize false positives without missing actual threats is a constant struggle."***

Figure 6

PLACES WHERE AI IS OVERHYPED

In which security tech categories do you think AI is the most overhyped or underdelivering on promises?

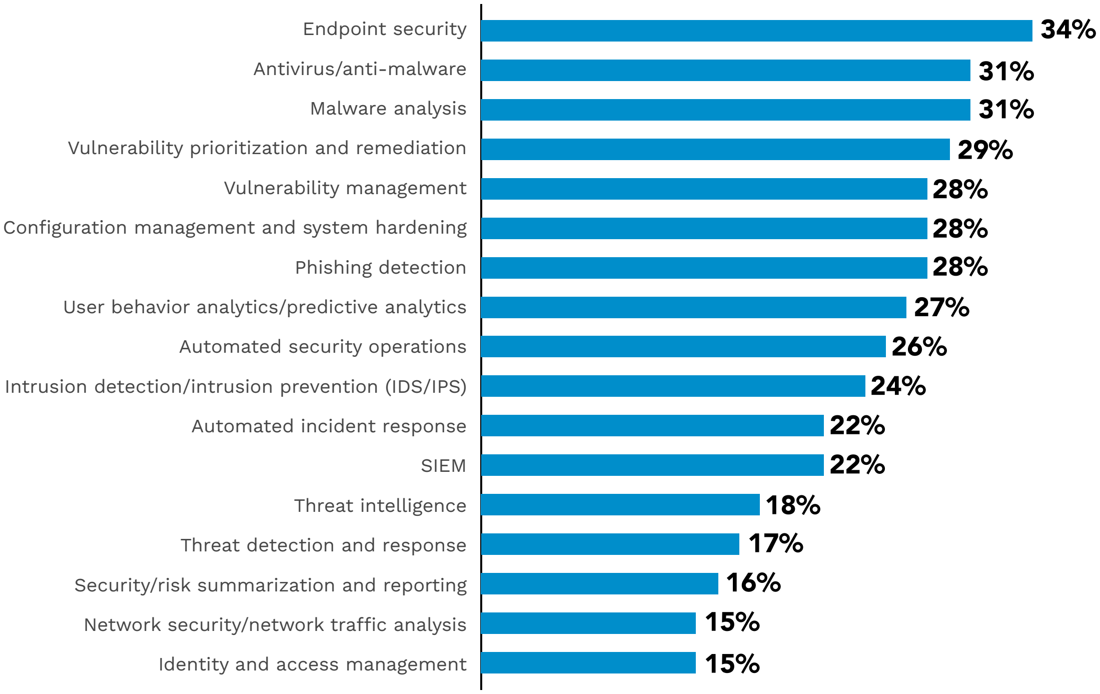

Note: Multiple responses allowed
Data: Dark Reading survey of 94 Cybersecurity and IT professionals at large companies, November 2024

Figure 7

OBSTACLES TO USING AI IN CYBERSECURITY

What are the biggest obstacles to the effective use of AI in cybersecurity today?

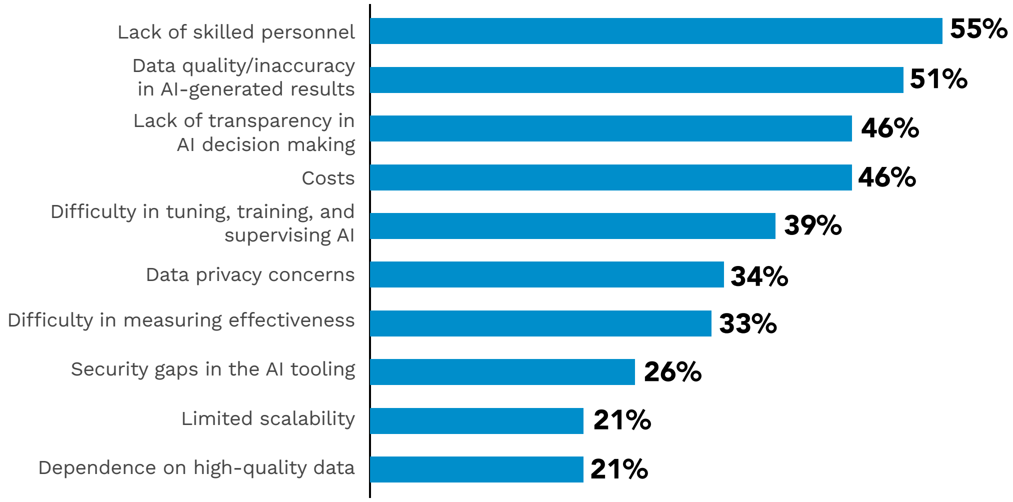

Note: Multiple responses allowed
Data: Dark Reading survey of 94 Cybersecurity and IT professionals at large companies, November 2024

Figure 8

EVALUATING AI EFFICACY

How do you evaluate the efficacy of AI?

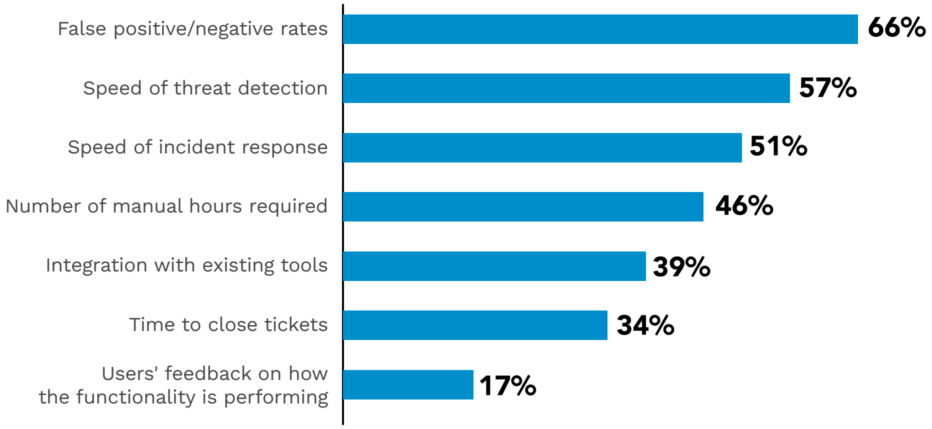

Note: Multiple responses allowed
Data: Dark Reading survey of 94 Cybersecurity and IT professionals at large companies, November 2024

This was a common complaint and likely why false positive and negative rates are the No. 1 way that organizations reported that they evaluate the efficacy of AI in security, named by 66% of respondents. That was closely followed by speed of threat detection (57%) and speed of incident response (51%) (**Figure 8**).

It's interesting to note the cognitive dissonance that false positives raise when it comes to AI value. On one hand, many organizations hope that the next iterations of security AI can help solve interwoven problems of security data overload, false positives, and decision fatigue. But on the other hand, it is clear that many AI-driven security tools in place are at the maturity stage where they still frequently exacerbate these issues with less-than-reliable results or a lack of transparency around how the results are derived, which was named by 51% and 46% respectively of respondents as major obstacle.

Addressing Resource Constraints

The biggest issue holding back security AI efficacy, though, is that many security teams don't have the wherewithal to effectively evaluate and appropriately manage the many resource-intensive tools that dominate the security tooling landscape today. Many AI tools developed by security vendors can still take significant care and feeding to work properly, including work on model training, tuning, and configuration.

Some 55% report that lack of skilled personnel is their biggest obstacle in effective use of AI in cybersecurity, which was the top response for named challenges. And in the free response question about their No. 1 AI challenge, many respondents cited issues around lack of training, knowledge, and resources to appropriately manage and tune their AI-enabled tools.

Some of the biggest resource requirements for effective AI usage have to do with model training and the meticulous data management that must scaffold this activity. As one respondent related:

"A significant challenge in applying AI to security practices is the need for vast amounts of high-quality, labeled data to train models effectively. This data is often hard to come by, especially for rare security incidents or emerging threats. Managing this data and keeping AI models up to date can be resource-intensive and complex."

Nearly a third of respondents reported that their team spends at least four hours per week training AI models within their own tools or within commercially available AI functionality. And just a fraction of respondents said that their tools come trained and/or tuned — 5%.

Those few organizations that do have the resources to become more sophisticated about their data management and data science are the ones that currently get the most benefit from AI in security.

It's a small sample group, but 100% of those who said their AI is very beneficial and a vital part of their security program have internal data science staff members.

That's great for organizations that can afford these investments. But realistically, only the largest organizations are going to be able to field an expert team of security data scientists.

All of this indicates the market opportunity and where buying evaluation should be focused on vendors selling security AI that works out of the box in the coming years. The average security teams need better support with more fully fleshed data science capabilities from their vendors to really start deriving value from their solutions around things like vulnerability prioritization, threat detection, and more. This means not only better support built into tools and contracts, but also perhaps more data science outsourcing opportunities. Currently, just 6% of respondents say that they fully outsource their AI training (**Figure 9**).

AI Governance in Security Tooling

One area that the field of vendors will also need to improve upon is better transparency and governance over how the AI works in their tools. The need for this clarity and control is twofold: both to improve the efficacy and also the safety and compliance of these features.

Just over half of respondents said that they regularly disable AI functionality in some or all security tooling due to a range of considerations. In fact, the majority of organizations (55%) say that they've enabled AI in under half the tools in their environments that have it available.

A lack of transparency and explainability was the top reason for turning off AI functionality, cited by 58%, and it was closely followed by security and privacy risks (55%), and vendor reliability and maturity (50%).

Figure 9

How much time does your team spend training AI models either within its own tools or commercially available AI functionality?

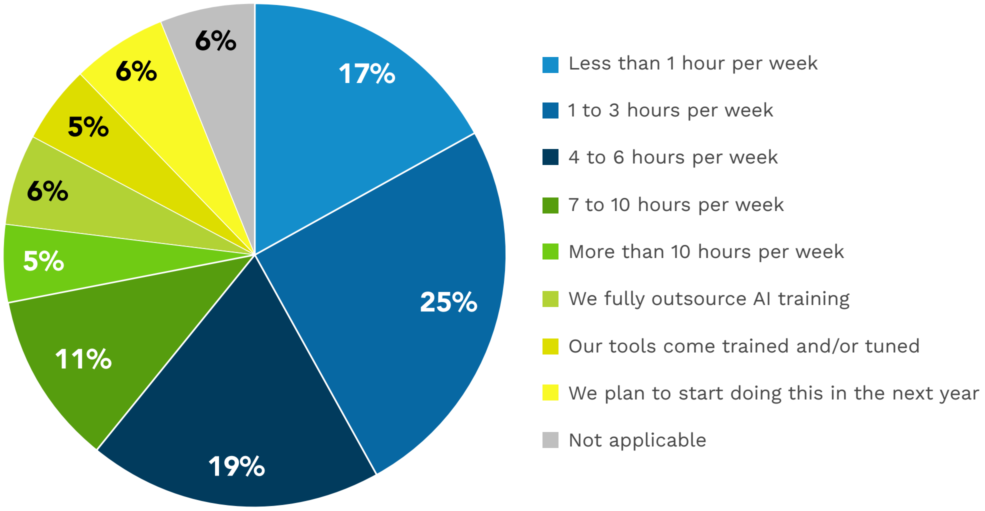

Data: Dark Reading survey of 94 Cybersecurity and IT professionals at large companies, November 2024

Chief information security officers (CISO's) are concerned that the AI capabilities in security tools could themselves become a threat vector if not properly controlled. And since a lot of security AI is still operating as black-box technology, risk professionals can't be sure whether those controls are in place. It's a maturity and trust problem that many security executives are following closely.

One respondent explained the dynamic in detail:

"I would say that our No. 1 challenge in applying AI to security practices is the lack of transparency and explainability in AI decision-making. Sure, while AI can significantly enhance our security operations, it often acts as a 'black box,' if you will, making it difficult to understand how it arrives at certain conclusions. This lack of transparency can lead to trust issues among our team and stakeholders, as we need to be confident in the accuracy and reliability of AI-generated results. Additionally, without clear explanations, it becomes challenging to identify and address any potential biases or errors in the AI models, which can ultimately impact the effectiveness of our security measures, in my professional opinion."

This inevitably dampens the widespread use of AI in daily security work. Vendors and security professionals should expect to see these governance issues continue to surface in the next few years.

3 Ways to Get More from AI in the Next Year

Clearly, AI use in security is maturing to the point where it truly moves beyond the hills and valleys of inflated expectations and fast-following disillusionment. Security teams that want to start reaping real benefit from AI embedded in their security tooling should consider the following tips for 2025.

Be discerning with vendor evaluation

Every vendor today can claim AI-powered something, but the hype is still very breathless. Buyers can dig into true benefits and capability by asking tough questions. This starts with those about false positives/negative rates and how much work the vendor expects the buyer to put into tuning and configuration. But buyers should also be turning the screws on vendors about their AI governance policies and asking for greater transparency in how their AI capabilities work under the hood.

Look for ways AI can boost vulnerability management

Many security practitioners hope that 2025 is the year that security teams and tooling finally make good on the promise of using AI to improve prioritization of vulnerability and exposure management. Not only can AI algorithms help prioritize flaws and exposures by risk, but AI could also be used to streamline and bolster the automation of remediation workflows. For practitioners, we encourage teams to start small, implementing AI enhancements incrementally. This is a smoother way to adopt new technology for both the staff and existing processes. Implementing AI in one part of the process at a time will make challenges more manageable and minimize the impact of errors that are inevitable in early deployment stages.

Invest in at least one internal expert

While most organizations with custom AI tooling have internal data science staff, there's still one-third of them that don't. Even for those that embed AI through commercial tools, it's still beneficial to have an internal AI guru who can help them get the most out of the AI functionalities. It may be a cost to begin with, but they're more likely to see better ROI from security AI in the long run. According to the data, for organizations that have a dedicated data science team, 100% see AI as "very beneficial," highlighting the importance of skilled expertise.

Finally, organizations should remember to have realistic expectations about benefits they can expect from their security stack's AI capabilities. As AI is still relatively new, it's not going to magically solve every single issue — especially in a domain as complex as cybersecurity. The benefits, as impactful as they might be, will be minimal to begin with and not omit of error. So, it's important that those who adopt/want to adopt AI understand this.

Survey Methodology

Seemplicity commissioned Dark Reading to conduct a survey about the benefits and challenges cybersecurity and IT professionals experience with AI as they navigate its use in cybersecurity.

The survey collected responses from 94 cybersecurity and IT professionals in North America all of whom were management-level titles or above and worked at large companies with 100 or more employees and $500 million or more in company revenue. The survey was conducted online in November 2024. Respondents were recruited via emailed invitations containing an embedded link to the survey. The emails were sent to a select group of Dark Reading's qualified database.

The final data set includes job titles from executives to manager level, predominantly located in North America. Nearly half of respondents (49%) held IT and cybersecurity executive job titles (CIO/CTO, CSO/CISO, VP of IT, or VP of cybersecurity). Other titles included cybersecurity or IT directors or heads of the department (18%), with the remaining titles primarily being cybersecurity or IT management (26%).

Respondents worked at large companies representing more than 20 vertical industries, including banking/financial services/accounting, consulting, healthcare, education, aerospace, and technology manufacturing, to name those cited by 7% or more. Thirty-one percent reported that they work at very large companies with 10,000 or more employees; 15% at companies with 5,000 to 9,999 employees; 40% at companies with 1,000 to 4,999 employees, and 14% from organizations with 500 to 999 employees.

Dark Reading was responsible for all survey design, administration, data collection, and data analysis. These procedures were carried out in strict accordance with standard market research practices and existing US privacy laws.