SANS | Research Program

Survey

SANS 2024 AI Survey

AI and Its Growing Role in Cybersecurity: Lessons Learned and Path Forward

Written by **Matt Edmonson** with **Matt Bromiley**

September 2024

hackerone

©2024 SANS™ Institute

Introduction

Artificial intelligence (AI) is rapidly becoming an integral part of many organizations, transforming business processes, enhancing operational efficiency, and improving security. From healthcare to finance to cybersecurity, AI can analyze vast amounts of data and automate decision making. Users can interact with AI systems to improve their efficiency, leading to sectors all over the world embracing AI technologies as soon as they can.

However, the rapid adoption of AI also introduces significant challenges. Incorrect or misleading results can undermine the trust in AI systems, and lead to a decline in adoption. As organizations rely on AI for more and more tasks, accuracy becomes paramount. Despite the impressive capabilities of AI, it is necessary to have strategies in place to mitigate inherent technology risks.

This paper reflects results from the 2024 SANS AI Survey, conducted in March and April 2024. The survey explored the current state of AI adoption for cybersecurity and provided insights into how various organizations manage and minimize the risks of AI shortfalls. Additionally, it helped examine the impact of AI on security operations and the workforce, looking for trends and how AI may be reshaping these areas. By understanding these dynamics, organizations can better prepare for the integration of AI technologies in their workflows.

A few of our key takeaways include:

- **High concern for AI-powered threats—**Most organizations (82%), whether or not they have currently implemented AI, are concerned about AI's impact(s) on offensive cybersecurity tactics, such as automated vulnerability exploitation or more advanced phishing campaigns.
- **Widespread AI adoption—**Approximately 43% of organizations are currently using AI as part of their cybersecurity strategy, while another 38% plan to adopt it. For those organizations that have implemented AI, key findings include:
  - **Impact on workforce and training—**Roughly 40% of respondents have an initiative to prepare their workforce for the evolving AI-driven cybersecurity landscape. Approximately 75% of these organizations are preparing their workforce for AI with ongoing training in AI fundamentals and applications.
  - **Improving morale—**Slightly less than half (48%) agree that AI has influenced job satisfaction among their cybersecurity professionals. Of these, 71% of organizations report higher satisfaction due to AI automating tedious tasks, allowing focus on rewarding work.

Finally, this paper also offers insights into future developments and trends in AI technology. As AI continues to evolve, staying ahead of these trends will be crucial for maintaining a competitive edge and ensuring that the use of AI in cybersecurity is well-implemented. As you work your way through this paper, we encourage you to compare the results to your organization, looking for areas where you may be succeeding or have questions for future implementations. You can see how *your* organization stacks up to others in your demographic (see Figure 1).

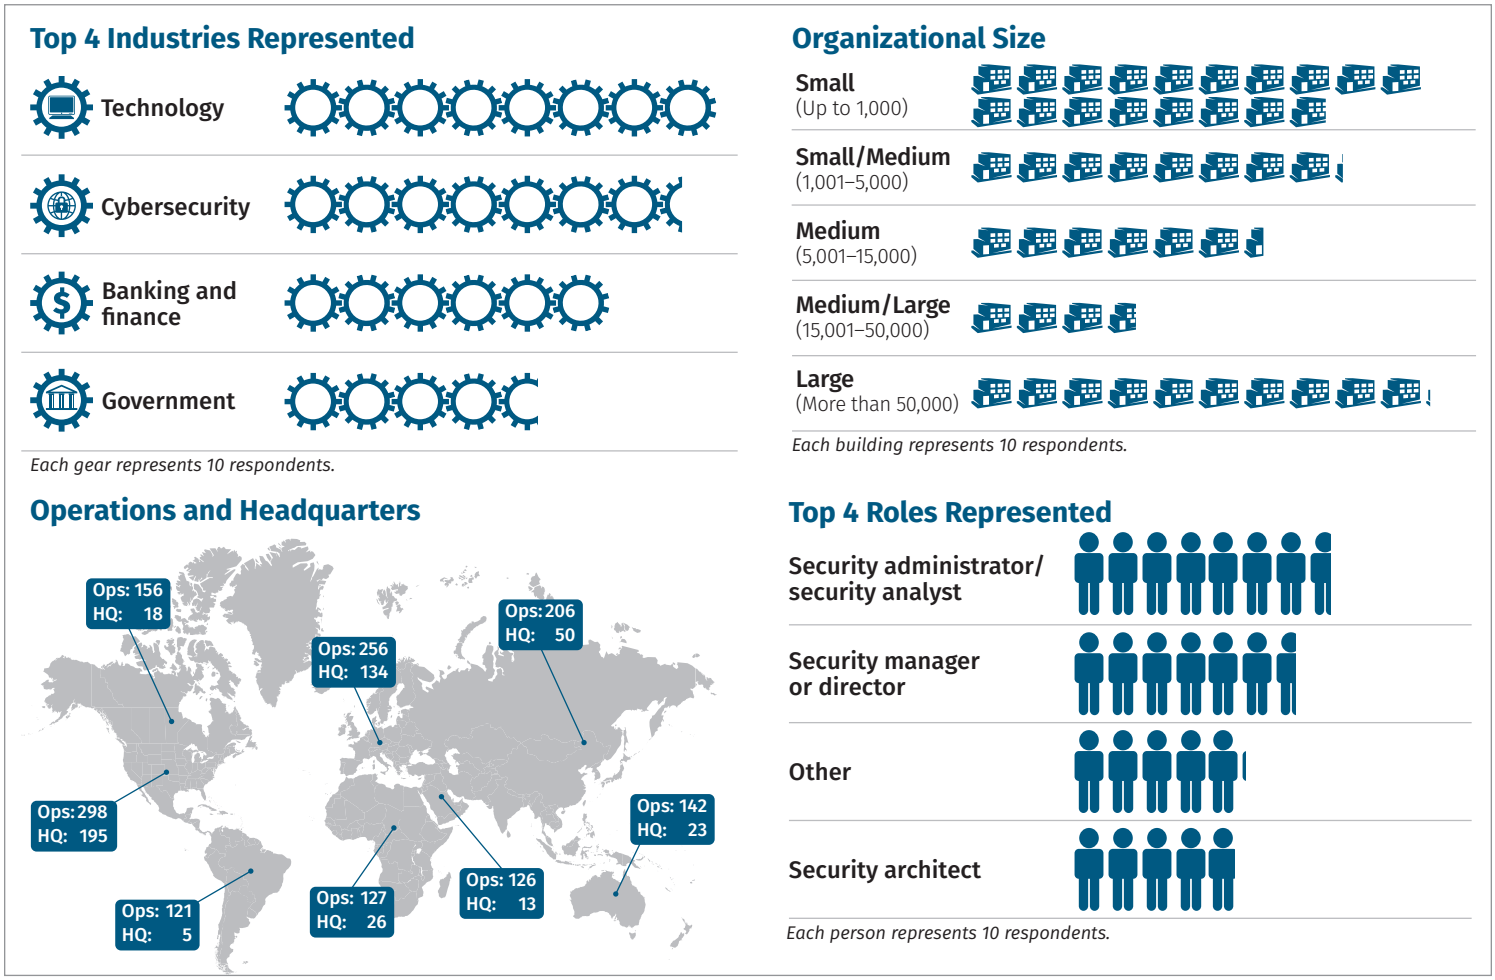

Figure 1. Survey Demographic Data

Current State and Adoption of AI

The adoption of AI within organizations has already seen enormous growth. Our survey showed approximately 43% of organizations are currently using AI within their cybersecurity strategies, and an additional 38% are planning to implement AI technologies (see Figure 2). This rate suggests that despite being around for just a few years, organizations already recognize the technology's potential to enhance efficiency and improve decision making. Digging deeper, we found that approximately 66% of respondents *currently utilizing* AI are doing so within their security operations centers.

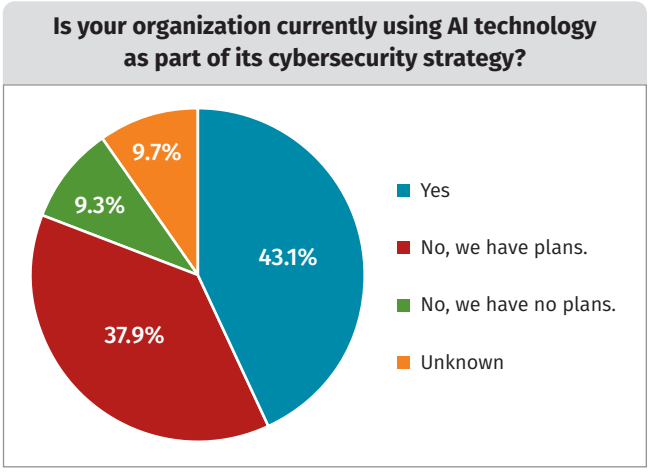

Figure 2. AI Usage

Although we did not collect specific reasons as to *why* AI has yet to be implemented, one potential reason could be a lack of awareness among users about AI tools and systems. We'll explore this later.

For those utilizing AI, our survey showed that AI is being applied in various areas of cybersecurity operations, including anomaly and malware detection (57% and 51%, respectively). Rounding out the top three is automated incident response (49%). See Figure 3.

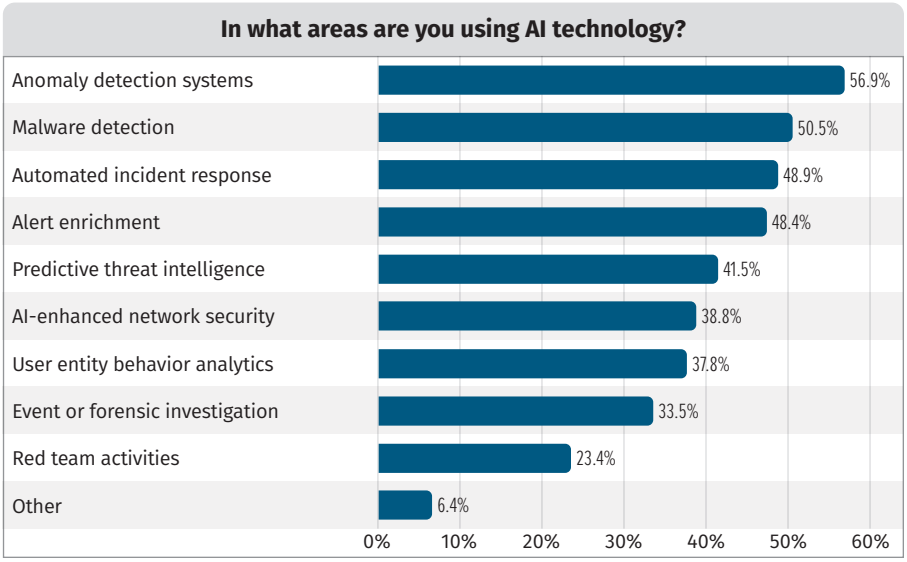

Figure 3. Key Areas of AI Usage

Other key areas of usage include alert enrichment, predictive threat intelligence, and network security. These use cases highlight AI's ability to analyze large datasets and quickly identify patterns that may indicate threats or anomalies to the environment. By automating these processes, organizations are responding to threats quicker and more efficiently.

The use of AI isn't limited to defensive operations. Our survey also found AI usage in red team operations, albeit not widespread. Approximately 23% of respondents indicated they were using AI in red team operations, with the majority (approximately 75%) utilizing AI to simulate more advanced cyberattacks.

Of course, these implementations do not come without risks or concerns. We observed that despite its many uses and implementations, there were no 100% implementations. Although organizations might like the idea of implementing AI in all areas of cybersecurity, many still recognize the need for human analysts to either keep the systems "in check" or focus on more rewarding tasks, allowing AI to handle automated, repetitive tasks.

Challenges and Shortcomings of AI

The adoption of AI can offer numerous benefits, but also presents challenges and shortcomings that should be top of mind. We must be aware of potential issues *prior* to widespread adoption to ensure we apply technology and train users correctly. With any security implementation, there is never a "one-size-fits-all" solution. AI is not immune to the need for a customized and well-integrated approach, with risks appropriately mitigated.

Skills and Legal Considerations

The integration of AI into organizational workflows requires a specialized skill set that many organizations currently lack. Identifying the risk skill set and balancing risks is a difficult task, especially given the age of some AI-driven technologies. Approximately 44% of respondents who have implemented AI identified a skills gap as a significant challenge in adopting AI technologies. Additionally, 42% responded they had difficulty trusting AI decisions due to a lack of transparency. These are valid concerns that AI technologies must address. See Figure 4.

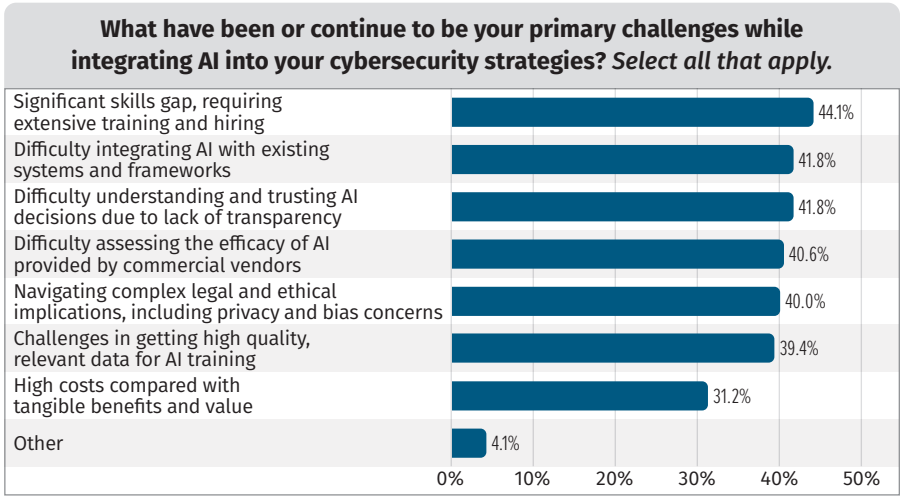

Figure 4. AI Integration Challenges

Interestingly, we also found that 40% of respondents cited legal and ethical concerns regarding AI privacy and bias. The diversity of AI use cases contributes to a diverse array of legal considerations for AI implementations. Only 31% of organizations reported cost as an issue.

AI implementations don't come without their own shortcomings in security operations. Thirty-nine percent of those respondents who are using AI indicated they have faced *significant* shortcomings of AI in detecting or responding to cyber threats. See Figure 5.

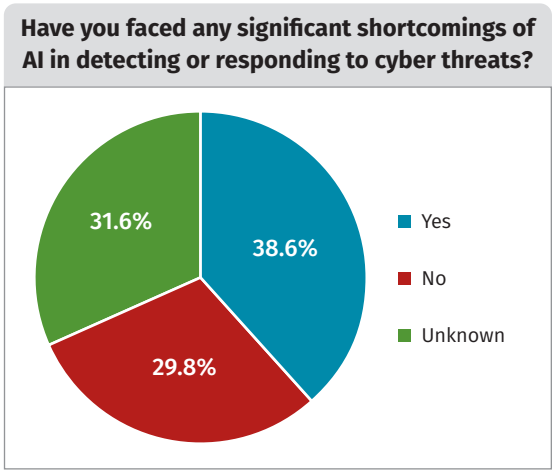

Figure 5. AI Shortcomings

Some of the top shortcomings include:

- Approximately 71% reported AI systems generating false positives, leading to alert fatigue.
- Nearly 58% of respondents recognized heavy dependencies on quality and relevance of training data.
- Approximately 56% found that their AI systems struggled to identify new threats or other outlier indicators, again due to training data.

Large language model (LLM) fine-tuning is an iterative process, and the necessary trial and error testing can be intensive and demanding on already-busy technical schedules. Another significant issue in AI is the phenomenon of "AI hallucinations," where the system generates incorrect or misleading information. These hallucinations can undermine the effectiveness of AI systems, leading to incorrect conclusions and problematic results for systems and users *dependent* on these decisions. The accuracy and reliability of AI outputs remain critical for cybersecurity purposes, and lead to an enhanced requirement for transparency in AI technologies.

Transparency in AI Decision Making

Building trust in AI systems requires making their decision-making processes more transparent. This survey indicated a growing need for transparency in AI, helping users understand how decisions are made and ensuring that outputs are reliable. Sixty-one percent of respondents indicated that AI decisions needed more transparency *and* a refinement in AI algorithms to reduce false positives and alert fatigue. See Figure 6.

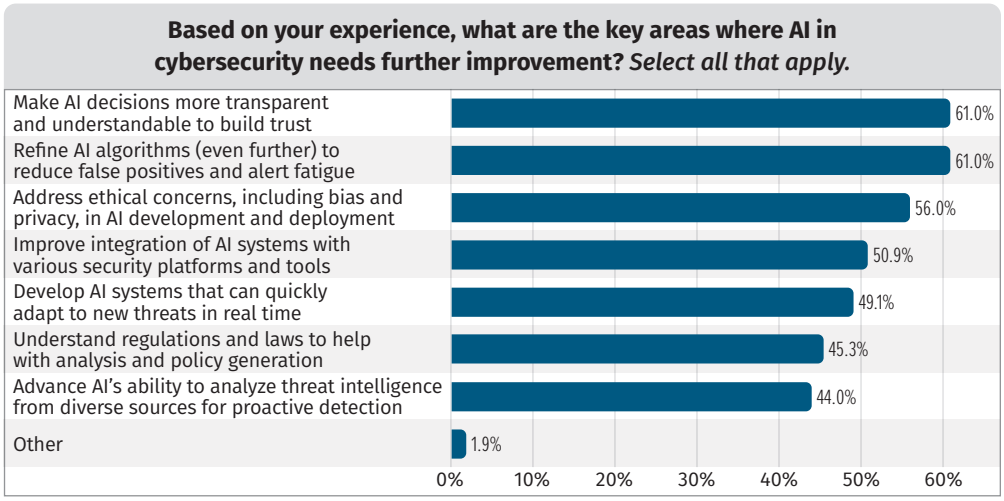

Figure 6. AI improvements Needed

A close third, at 56%, was the need to address ethical concerns, including bias and privacy in AI development and deployment. We also found concerns on AI bias stem down into specific functions, such as red team operations. Approximately 60% of respondents indicated they were concerned with AI respecting privacy or creating skewed results from bias.

The integration of AI into organizational workflows presents several challenges and shortcomings that must be addressed to fully leverage the power of AI. Between skills gaps, legal considerations, and AI hallucinations, transparency and lack of bias are critical to future adoption and implementations.

AI in SecOps and Workforce

The integration of AI into security operations centers (SOCs) and its impact on the workforce are pivotal aspects of successful AI adoption and trust building. According to the survey data, AI is significantly influencing security operations and reshaping roles within those organizations. Approximately 66% of applicable respondents indicated they are using AI in their SOCs, underscoring the growth AI has experienced in this area of security.

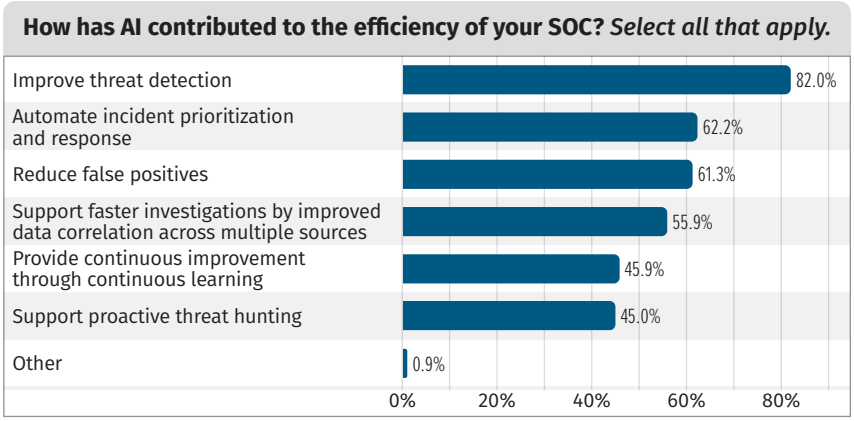

Figure 7. AI Efficiencies

AI's effectiveness in the SOC is further demonstrated by the ability to automate various tasks that might otherwise consume an inordinate amount of time. A whopping 82% found AI useful for improving threat detection—an expected result because AI can easily assist in the analysis of adversary tactics, techniques, and procedures (TTPs) and crafting of associated detections. See Figure 7.

Approximately 62% of organizations are using AI to automate incident prioritization and response, minimizing potential downsides and tedious, time-wasting tasks better suited to automated systems. Another excellent use of the technology, found in 56% of respondents, is supporting faster investigations with improved data correlation across multiple sources.

AI for Red and Blue Team Operations

Our survey found that AI is making significant inroads in both red and blue team operations. Of the 30% who use AI in their red team activities, 74% are leveraging AI to simulate more sophisticated cyber-attacks in their red team training.

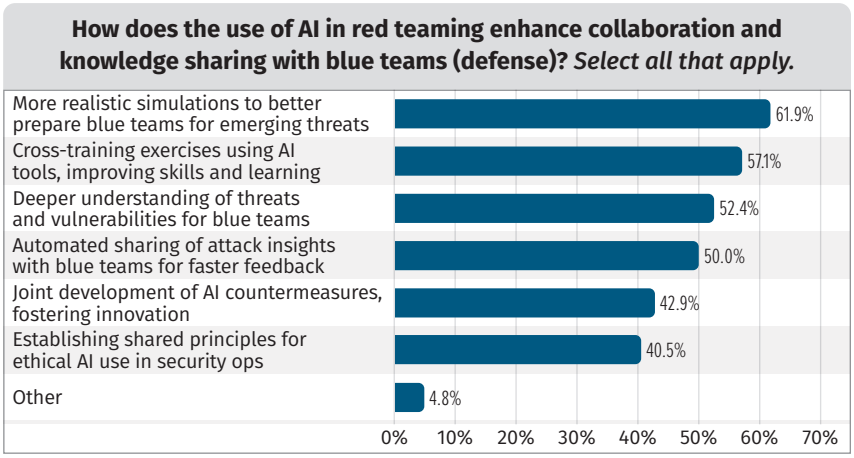

Figure 8. Collaboration Enhancements: Red Team and Blue Team

Approximately 62% of our respondents indicated that AI is used to create more realistic attack simulations, better preparing blue teams for emerging threats. A little over 57% of respondents found that cross-training exercises using AI tools provided better skills and learning opportunities for red/blue activities. See Figure 8.

Other notable areas include a deeper understanding of threats and vulnerabilities (52%) and automated sharing of attack insights with blue teams for faster feedback (50%). We cannot overstate this: Red teams exist to make blue teams stronger. AI-positive integrations between red and blue team activities only help strengthen the organization's overall security posture and encourage adoption of AI technologies. However, as we noted earlier, respondents are concerned with the highly complex and ethical issues of using AI in offensive security operations. Furthermore, approximately 36% of respondents indicated that red teams might have an issue keeping up with rapidly evolving AI defenses deployed by blue teams.

Impact on Training and Job Roles

The integration of AI into security operations will have a profound impact on job roles and training requirements. According to our survey, approximately 60% of organizations report changes in training needs for the security team due to AI adoption. Top areas of education and training concern include:

- More specialized AI and cybersecurity courses
- Emphasis on continuous learning with a focus on AI
- More hands-on experience with AI-based tools through simulations and experience
- Modules on ethics, privacy, and responsibility

Our survey also found that AI is having an influence on observed changes in job positions. More organizations (44% of those respondents that are actively using AI) are creating AI-focused security positions and emphasizing continuous education. Approximately 64% of these applicable respondents are seeing changes in continuous education (with a focus on AI advancements), and 63% are seeing increased collaboration between cybersecurity, IT, and AI experts. Rounding out the top three, approximately 55% are seeing new AI-focused positions being created, such as AI security analysts and data scientists. See Figure 9.

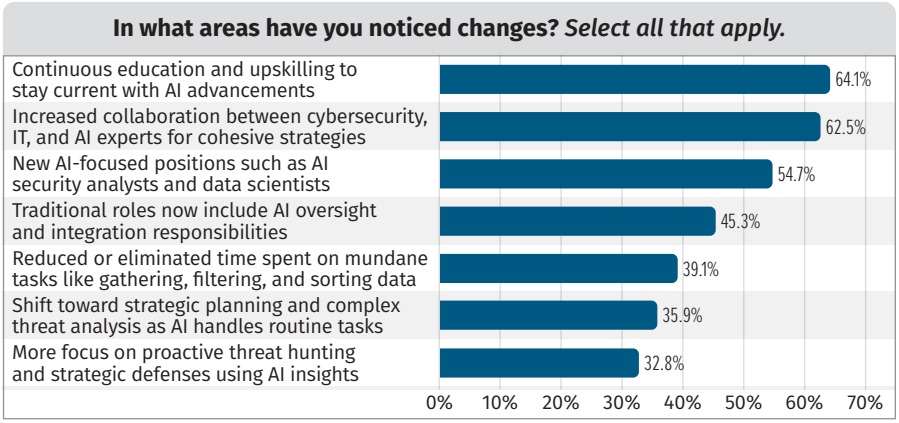

Figure 9. Changes in Job Positions

Another result we were happy to observe from the survey: Approximately 48% of respondents who have implemented AI are seeing an influence in cybersecurity job satisfaction from AI. First, anything that improves quality of life or job satisfaction for cybersecurity professionals is an instant plus for us! Key areas of improvement include:

- Higher satisfaction due to AI automating tedious processes (71%)
- A greater sense of accomplishment after AI integration (61%)
- Opportunities for growth by acquiring new skills and advancing careers (54%)

AI can introduce significant benefits for security teams and cybersecurity professionals. Increasing morale and automating tedious tasks is only half the battle. One area we celebrated was the 44% who indicated that employees developed appreciation for a better work/life balance through AI-driven efficiency. We'll take those results any day!

The survey also indicated that approximately 58% of organizations implementing AI have ongoing training initiatives to prepare their workforce for AI. This underscores the need for reliable, effective, and trustworthy AI implementations within security teams. Workforce preparation includes:

- Ongoing training on AI fundamentals (75%)
- Training on ethical considerations, including privacy and responsible use (47%)
- Organizing workshops, seminars, and hackathons for knowledge sharing and hands-on experience (42%)
- Partnering with universities and online platforms for cutting-edge courses and certifications (33%)

These initiatives, which we expect to grow in future years, ensure that employees are equipped with the necessary skills to leverage AI effectively and stay ahead of risks and emerging threats.

Best Practices and Future Developments

The successful integration of AI in organizations requires adherence to best practices and an understanding of future developments in the technology. We authored this survey in 2024, knowing that the landscape of AI technologies is likely to change drastically in the coming months and years.

Best Practices for Integrating AI into Cybersecurity

To effectively integrate AI into cybersecurity, organizations must align their AI initiatives with an overall cybersecurity strategy. According to our survey respondents, approximately 75% of respondents believe in a strategic implementation that aligns with overall cybersecurity strategy and business goals. See Figure 10.

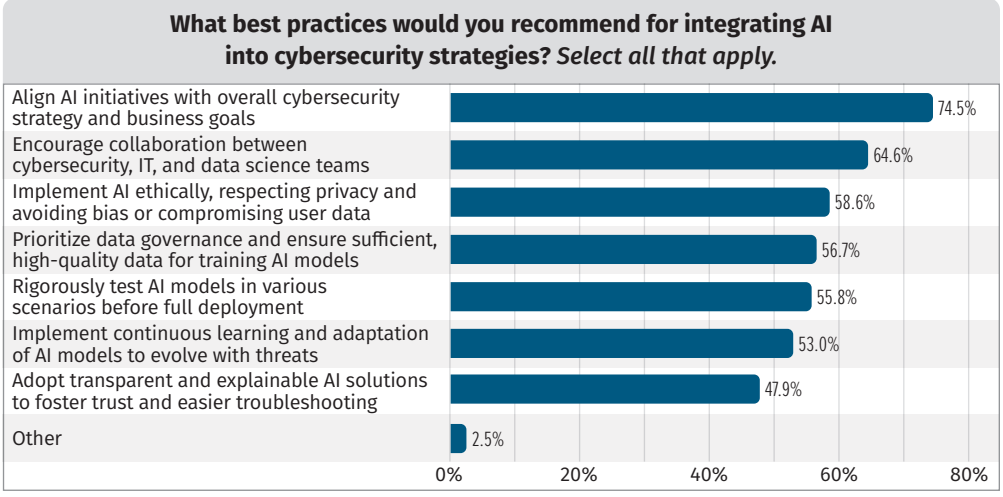

Figure 10. Integration Best Practices

In second place, approximately 65% of respondents encourage collaboration between cybersecurity, IT, and data science teams. In third place at 59%—and perhaps no surprise—is the best practice to implement AI ethically, respecting privacy and avoiding bias. Aligning AI with cybersecurity objectives ensures technologies are used to enhance, rather than undermine, a security posture. Our survey respondents indicated as such, with 58% responding that the cybersecurity team should contribute to mitigating risks associated with AI adoption. A close second, at 54%, indicated that the cybersecurity team should be performing regular audits of AI systems. See Figure 11.

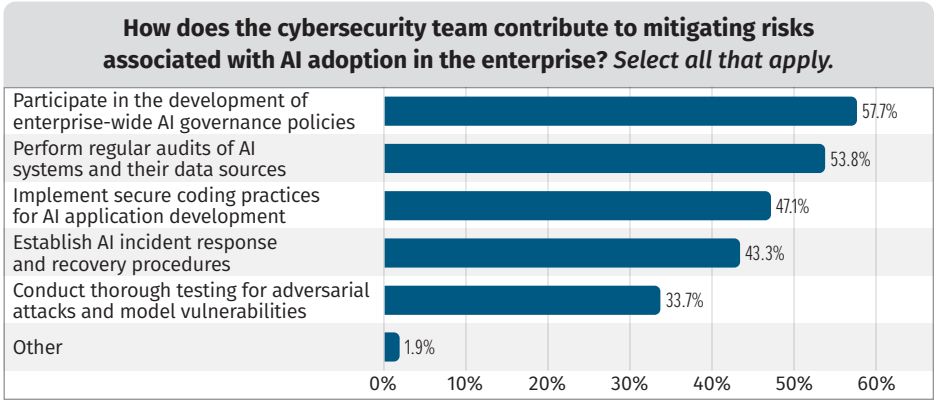

Figure 11. AI Risk Mitigation Tactics

As seen in Figure 11, our survey respondents find that the cybersecurity team has an integral role in mitigating AI-associated risks. Furthermore, this role is not just in auditing and development—approximately 43% found that the cybersecurity team should be establishing AI incident response and recovery procedures.

Another unique takeaway from the above question is the approximately 34% of respondents who see the cybersecurity team as conducting thorough testing for adversarial attacks and model vulnerabilities. This is perhaps one of the most future-looking findings in this survey. Already a third of our respondents are recognizing that part of AI risk mitigation includes testing for attacks and modeling vulnerabilities.

Developing frameworks and methodologies to test, confirm, and assess the safety of AI technologies is quickly becoming a must for organizations who are looking to embrace AI safely. Furthermore—as our paper has focused on the ideas of bias and privacy concerns—it is only through repetitive, structured testing processes that we'll be able to alleviate these concerns.

Additionally, organizations should focus on continuous learning and adaptation to keep up with AI advancements. Promoting a culture of continuous learning ensures that employees are equipped with the latest skills and aligns with previously discussed training objectives. This also helps organizations stay ahead of emerging threats and maintain a competitive edge.

AI Enablement and Governance

Not unlike previous concerns of bias and privacy issues, our survey revealed that organizations adopting AI had concerns about AI-powered threats. As previously discussed, threats such as deep fakes, advancements in phishing campaigns, and other automated attacks pose significant risks to organizations. These threats leverage AI and put security teams in the juxtaposition of using AI to battle AI.

Addressing these concerns requires robust AI governance and security strategies. Approximately 69% of respondents indicated that the cybersecurity team had a role in enabling and governing the use of AI across the enterprise. See Figure 12.

Key roles for the cybersecurity team, with respect to AI enablement, include:

- Developing AI risk management frameworks and policies (33%)
- Collaborating with IT and data teams to secure AI deployments (27%)
- Conducting AI security assessments and penetration testing (17%)
- Monitoring AI deployments for potential security breaches and anomalies (17%)

Cybersecurity-led governance can be a force multiplier for successful AI integrations, because the security team is likely in tune with and aware of risks and concerns and can effectively assist in designing mitigation strategies.

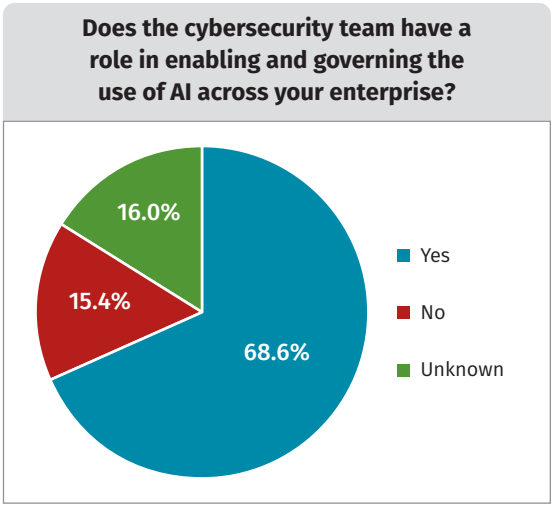

Figure 12. AI Governance and Enablement

Anticipated Developments and Key Trends in AI

Our survey indicated that the demand for AI professionals is expected to grow, with 66% of respondents anticipating an increase in this need. This only underscores the growing importance of AI, as this survey clearly reflected. AI is also becoming an integral part of cybersecurity education and training, according to 60% of respondents. See Figure 13.

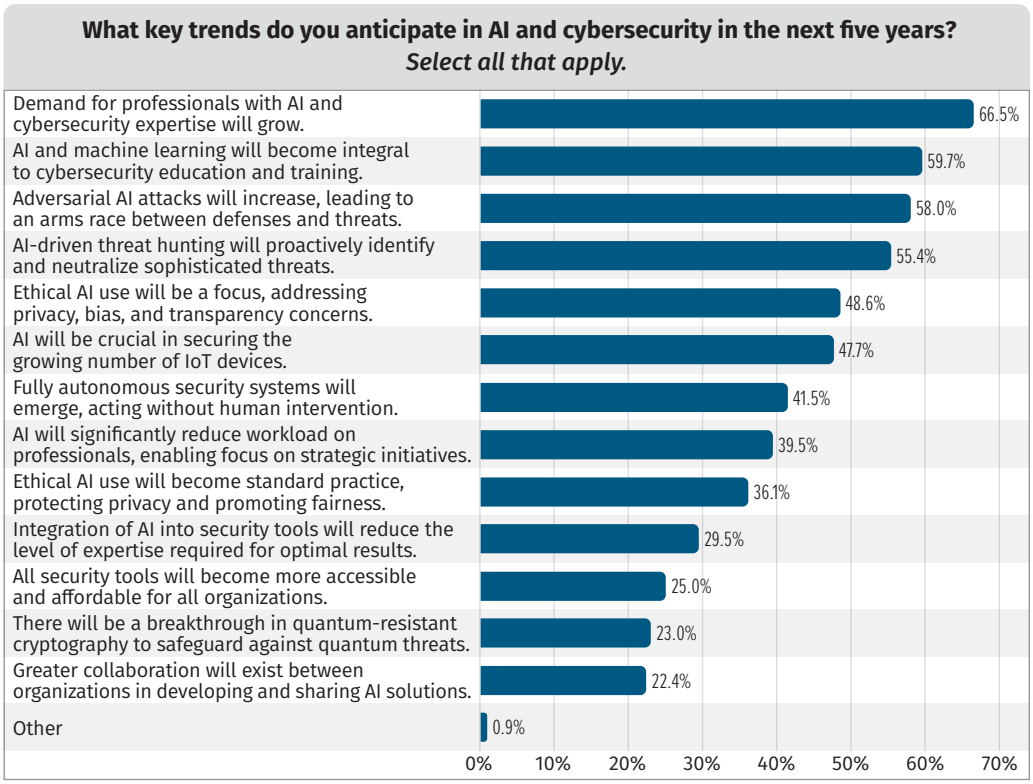

Figure 13. Key Trends

Interestingly, a close third place (58%) feel that an "arms race between defenses and threats" is also on the horizon. Other posture-related predictions include:

- Ethical AI use to become a focus, addressing privacy, bias, and transparency concerns (49%), with ethical AI usage becoming standard practice (36%)
- Criticality of AI in securing IoT devices (48%)
- Continued reduction on the workload of professionals (40%)
- Security tools becoming more affordable and accessible (25%) while reducing the expertise required to achieve optimal results (30%)

One of the most consistent findings in this survey, especially with respect to future trends, is that transparency, privacy, and bias concerns will get better. This signals a wholehearted wish to embrace AI technologies, but also a moment to pause and ensure that the technology is in the "right place," protecting organizations and their data. We can lessen these concerns by utilizing the AI technologies available, training them to be free of bias and privacy concerns. In short, we're nowhere near "done" when it comes to AI technology. It will take years of training and tuning in order to establish a baseline confidence level that bias and privacy concerns are minimal.

We view building trust as a paramount step forward for AI technologies and developers.

Promoting a culture of continuous, AI-focused learning is another key trend identified throughout the survey. Organizations that prioritize ongoing education and skill development will be better positioned to adapt to AI advancements and maintain their competitive edge. This will also contribute to raising their confidence level in AI technologies. Similar to other data concepts, privacy concerns will need to be addressed *as much as possible* ahead of time, but also as they come up. Do we expect to see legislation in the future addressing AI concerns? Absolutely.

Addressing Misconceptions and the Importance of Transparency

One of the biggest misconceptions is that AI is a job killer. This survey indicates that AI should be viewed as an enabler, rather than a threat to jobs. By automating routine tasks, AI empowers employees to focus on more strategic and value-added activities, enhancing (as we've seen) overall productivity and job satisfaction. See Figure 14.

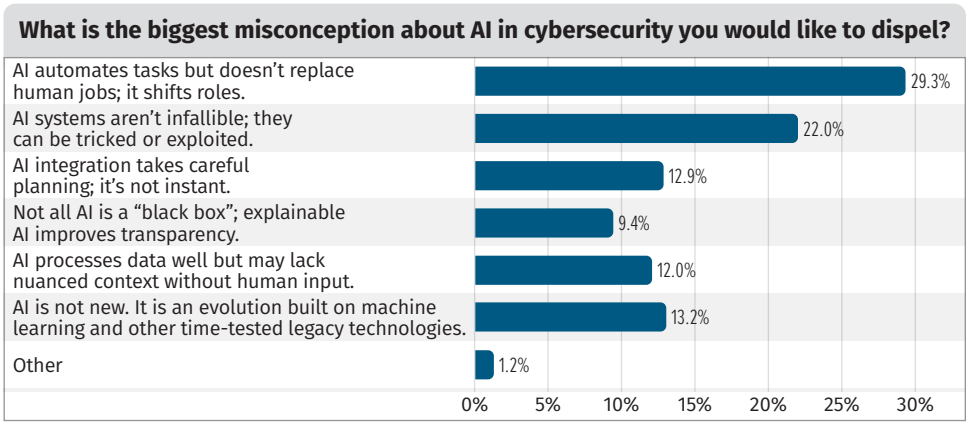

Figure 14. Misconceptions About AI

Transparency in AI algorithms is crucial for building trust and addressing bias concerns. According to the survey, compliance is a major driver of transparency initiatives. Ensuring that systems are transparent and comply with relevant regulations helps build trust and facilitate future AI technology use and growth.

Conclusion

AI is transforming various industries by enhancing efficiency, improving decision-making processes, and providing valuable insights. However, the successful integration of AI into organizations requires addressing several challenges, including trust, transparency, skills gaps, legal considerations, and the growth of AI-powered threats. By adhering to best practices and aligning with a cybersecurity-led strategy, organizations can effectively leverage AI while mitigating associated risks.

In this paper, we explored the current state of AI adoption, highlighting key challenges and shortcomings. We also looked at the impact of AI on security operations and looked ahead to best practices and future trends. As mentioned in the introduction, this survey can serve as an excellent barometer to help you address your current implementations and/or concerns.

When adopting AI, organizations are encouraged to prioritize transparency, invest in training and education, and align initiatives with an overall cybersecurity strategy. Build a trustworthy environment your users can depend on, address emerging threats, help automate tedious tasks, and find an increase in work/life balance and employee morale. As AI continues to evolve, staying proactive and forward-thinking will be crucial for maintaining an edge with AI-enabled technologies.

Sponsor

SANS would like to thank this survey's sponsor:

hackerone

hackerone

SANS | Research Program

Product Briefing

AI with HackerOne:

Insights from the 2024 SANS Institute Survey

September 2024

AI is everywhere, and yet it requires thoughtful guidance and monitoring from humans to realize the promised benefits. The most effective approach is to deploy AI for the task at which it excels, while relying on human expertise to keep the door shut on malicious actors.

HackerOne

As powerful as AI is, it's not a cure-all for security risks. The best protection in a constantly changing threat landscape is a combination of AI processing power and human assessment skills. That starts, as the SANS AI Survey indicates, with training your people in AI basics, but it doesn't end there.

HackerOne backs up its AI threat detection with a defense force of ethical hackers who can fight back against malicious actors—even against attacks that count as "unknown unknowns," the ones AI can't spot.

We all want to stop vulnerabilities as early as possible in the software development process—on the left side of the flowchart, the way most people depict it. With HackerOne, you can start by integrating real-time testing for code changes and live applications, ensuring that any vulnerabilities are caught early and that code remains secure throughout the development cycle.

HackerOne also provides human-led audits of an organization's codebase and, as you might expect, penetration testing to make sure nothing has been missed. Unlike many pentesting organizations, a HackerOne test lets you, the client, monitor what the testers are discovering in real time and, if needed, shut down the test before the end date.

Key Findings

Most organizations are concerned about AI's impact on offensive cybersecurity tactics, with 79% worried about AI-powered phishing campaigns and 74% about automated vulnerability exploitation.

AI has limitations when it comes to detecting threats—70.5% of respondents indicated issues with AI generating false positives while missing outlier indicators or new threats.

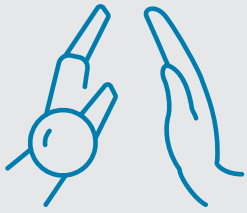

Approximately 71% of organizations report higher satisfaction due to AI automating tedious tasks, allowing focus on rewarding work.

You can work with HackerOne to set up a limited-time challenge or a continuous bug bounty program, both of which deploy a global network of registered security researchers against your applications. In fact, some organizations use these challenges in the form of AI red teaming to test the robustness of their own AI systems, including their models and software components. If application security is mission-critical to your organization, this is the level of assurance you need (Figure 1).

Figure 1. HackerOne Platform Capabilities and Comprehensive Offensive Security Testing Portfolio

Of course, HackerOne uses AI, for use cases like these:

- Tailored advice and personalized remediation guidance
- On-demand assistance with intricate reports, proofs of concept, and technical details
- Actionable insights from visuals and videos
- Custom apps for repetitive tasks
- Accelerated internal workflows and security tools with the HackerOne API

The latest product, Hai, is an AI co-pilot that adds context to vulnerability reports, transforms natural language into filtering queries, and uses vulnerability data from across the HackerOne platform to provide recommendations. The goal is to improve detection rates, prevent regressions, and speed up triage so analysts can spend more time fixing issues. Hai can read screenshots and the organization plans to introduce the ability to analyze videos that automatically detect key moments in videos, such as when payloads are sent to a system and executed. Integrating Hai into the workflows of people handling incoming trouble tickets helps front-line analysts call on the power of HackerOne's community through security-focused translation capabilities and auto-routing so that tickets quickly reach the team member best prepared to handle them.

For more information, visit **www.hackerone.com**

The HackerOne team is clear that at least for now, it's not time to trust AI with making final decisions about security issues and mitigations. There's always a human pulling the trigger, and their goal is to make sure that humans are as well prepared as possible, no matter what comes in the door. That lines up with where a lot of the SANS AI Survey respondents fall on the subject of AI trustworthiness, with more than 40% saying they have difficulty understanding and trusting AI-based decisions due to a lack of transparency.

If your digital resources need to be protected no matter what, and if you believe the best way to stop a hacker is with a better hacker, HackerOne may be the solution you need.

Note that SANS Product Briefings do not represent a SANS endorsement of a sponsor or its products, but rather an overview of its offerings and their capabilities.