ARCTIC WOLF

AI & Cybersecurity Trends Report

2026

TABLE OF CONTENTS

Foreword

Cybersecurity has always been defined by periods of inflection. Times of steady progress are followed by breakthroughs that fundamentally reshape how defenders operate and how technology enables them to succeed.

Today we find ourselves in one of those periods.

Over the last decade, organizations have invested heavily in security tools, analytics, and automation. These innovations have strengthened visibility and improved detection and response, yet the core challenge has remained the same: security teams continue to face overwhelming volumes of alerts, increasingly sophisticated adversaries, and environments that grow more complex every year. And all the while **the frequency and cost of breaches continue to rise.**

Defenders need a step change in capability, not incremental improvement.

Artificial intelligence has the potential to deliver that change. Across the industry, the promise of AI, and especially agentic AI for security operations, has generated excitement about a future where intelligent systems can plan, investigate, and take action alongside human experts, bringing speed, scale, and consistency to an increasingly complex problem.

At the same time, security leaders have been rightfully cautious. Concerns around hallucinations, model drift, and brittle reasoning have limited adoption. Many solutions have **struggled to demonstrate the reliability** required for security operations, where trust is paramount and mistakes carry real consequences.

It is in this environment that we conducted the survey powering this report. Speaking with those leading today's organizations, including owners, C-level executives, vice presidents, and directors ensures the findings within truly capture the perspectives of an authoritative and informed population.

The survey revealed that:

- The large majority of organizations suffered a "significant" cybersecurity incident in the preceding 12 months, with many suffering long-term disruptions as a result.
- Despite the prevalence of significant incidents, leaders have overwhelming confidence that their teams can keep up with the volume and complexity of today's threats.
- When it comes to applying AI to security operations, most organizations remain in the exploratory phase. One reason for stalled rollouts is that leaders don't trust AI-powered utilities enough to allow autonomous intervention.

Foreword

Two things can be true at once. Organizations taking a DIY approach to integrating AI into security operations have experienced limited returns (and encountered many unpleasant surprises). At the same time, using mature machine learning and capable AI agents to augment the broader set of people, processes, and technologies within an existing SOC can fundamentally change outcomes. The goal is not to replace human expertise, but to enable security operations to move at machine speed, with AI handling the scale and repetition while humans provide judgment, context, and oversight.

Cybersecurity is no longer the sole responsibility of IT and security teams. The findings in this report highlight a broader shift: cyber risk is now business risk, with direct implications for financial performance, operational continuity, regulatory exposure, and long-term competitiveness. For executive leaders, this means that cybersecurity outcomes must be evaluated not only in technical terms, but in their ability to preserve revenue, protect brand equity, and sustain business operations in the face of disruption.

Across the C-suite, priorities are beginning to converge. CEOs are increasingly focused on resilience and the ability to withstand disruptive events; CFOs are scrutinizing cybersecurity investments in terms of cost, risk transfer, and measurable return; and boards are demanding clearer visibility into both risk exposure and incident readiness.

In this context, the challenge is not simply adopting new technologies like AI, but ensuring that investments translate into measurable outcomes.

The organizations making the most progress are those embracing a new operating model for security: one built on a true partnership between humans and AI. In this model, AI brings the speed, scale, and consistency required to process vast amounts of data and execute workflows in near real time, while human expertise ensures accuracy, adaptability, and trust. Together, they enable security teams to move faster, respond more effectively, and reduce risk in ways that neither could achieve alone.

Ultimately, the path forward is not about choosing between human-driven or AI-driven security. It is about combining the strengths of both: aligning people, processes, and technology to deliver outcomes that map directly to business priorities. Organizations that succeed will be those that make this shift, transforming cybersecurity from a reactive function into a driver of resilience and competitive advantage.

ADAM MARRÈ
Chief Information Security Officer
Arctic Wolf

Methodology

This report draws upon a survey commissioned by Arctic Wolf and conducted by Sapio Research in April 2026.

To ensure a representative sample, Sapio surveyed 1,350 IT and security decision makers at director level or above, from organizations across the United States, United Kingdom, Canada, ANZ (Australia, New Zealand), DACH (Germany, Austria, Switzerland), the Nordics (Norway, Sweden, Denmark, Finland), Benelux (Belgium, the Netherlands, Luxembourg), Singapore, Japan, and South Africa.

Organization Size

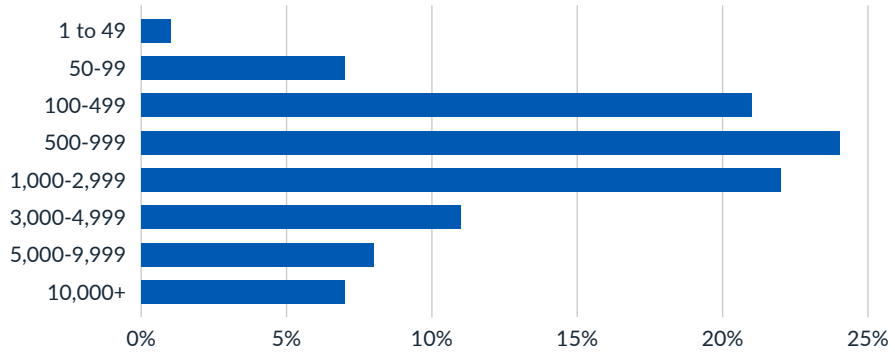

Organization Industry

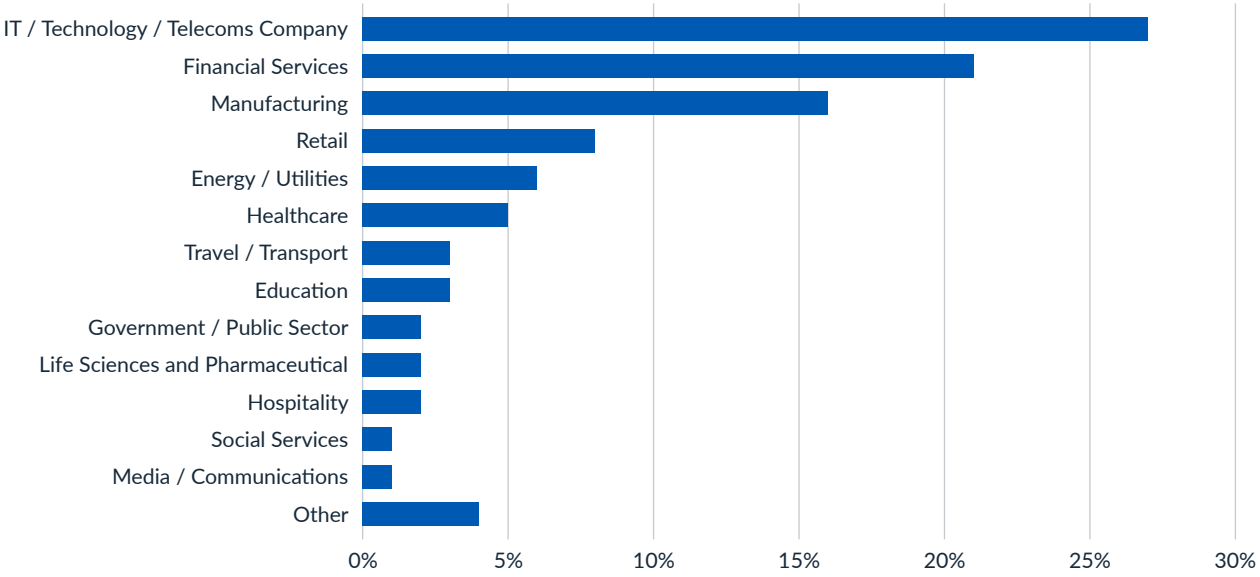

Respondent Role

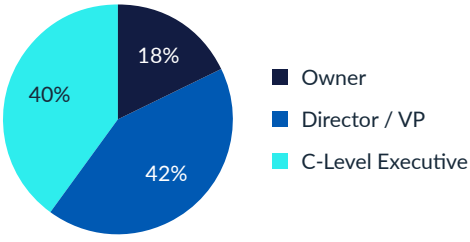

In this particular study, the chances are 95 in 100 that a survey result does not vary, plus or minus, by more than 2.7 percentage points from the result that would be obtained if interviews had been conducted with all persons in the same universal pool represented by the sample.

01

SECTION 1:

Threat Trends

TREND 1

The AI Boom: Adoption Is Outpacing Security, and Leaders Know It

Artificial intelligence isn't new. Machine learning (ML), in particular, has been a foundational element of data processing algorithms and applications for several decades, enabling companies to make sense of, and to find opportunities within, large volumes of information.

But since the debut of large language models (LLMs), including ChatGPT and Claude, two shifts have occurred:

1. **In pursuit of cost savings,** accelerated workflows, greater agility, and other business benefits, executives have prioritized the introduction of artificial intelligence tooling into a broad array of their organizations' everyday operations.
2. **Seemingly every industry** is exploring the potential of LLM-powered utilities to deliver returns; in the past, AI was primarily only used by tech companies and those working with vast datasets.

In just a few short years, LLMs and their related technologies have become ubiquitous in organizations of all sizes and in every industry. Our survey shows that 94% of organizations use LLMs in some capacity, with nearly half of this group (39% overall) having 2-3 approved tools (with the use of others being banned or prohibited).

Organizational approach to AI and LLMs

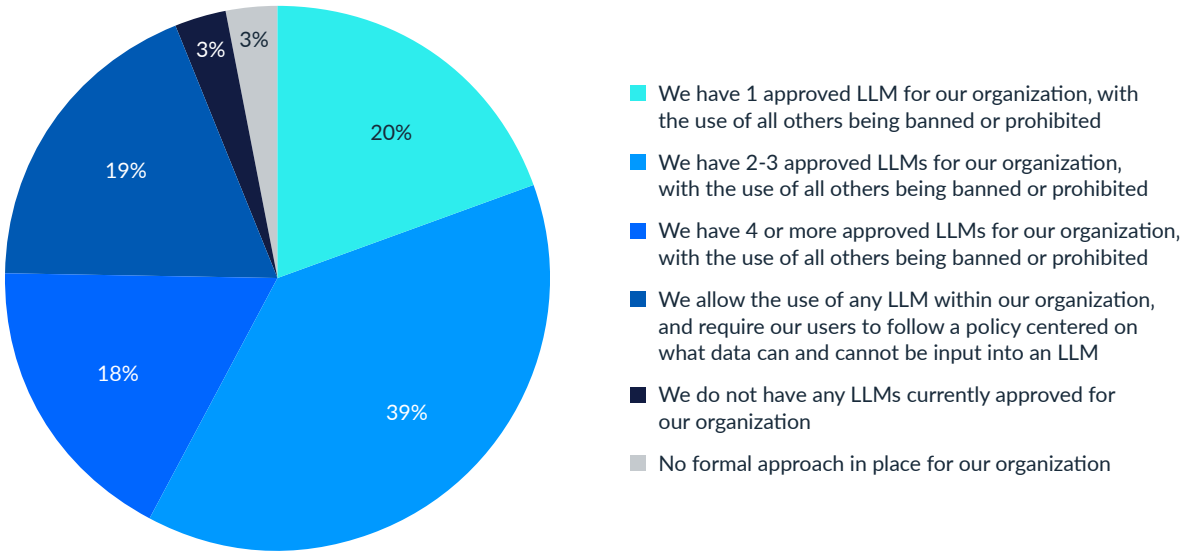

But LLMs collectively represent just one component of the contemporary AI revolution.

Next came the AI agents, which introduced greater autonomy, aided by more and deeper connections into other business systems. It wasn't long before "super agents" arrived, with the ability to connect to multiple LLMs and to spawn ever-more agents, with each potentially connecting across apps, APIs, SaaS tools, and data repositories in response to prompts.

Agentic AI represents a meaningful shift in how artificial intelligence operates within organizations. Generative AI can draft a report or answer a question, but it completes one task at a time and returns control to the user. In contrast, agentic AI can be given a broader objective and then pursue that objective through a series of coordinated, sequential actions; reasoning, planning, and executing without requiring a human to manage each individual step along the way.

As with the further expansion into a mobile, remote workforce, and adoption of cloud technology, any change to the IT environment has the potential to introduce new risks. Major changes that are introduced very quickly are therefore especially prone to doing so. Consider that last point, that cloud transformations and migrations have been going on for years, and organizations still struggle to secure these environments.

Therefore, it's no surprise that AI and its associated privacy risks are repeated as the top cybersecurity concern among survey participants (selected by 35%). And while we didn't explicitly ask about shadow AI (i.e., AI tools deployed within an organization outside of officially sanctioned channels), recent studies from **Verizon** and **Okta** have begun to shine a light on this often overlooked risk.

AI concerns first dethroned ransomware as the top area of concern in our 2025 Trends Report, and the gap between the two has expanded to 10 percentage points in the year since that report's publication, despite year-over-year growth in the share of respondents who regard ransomware and related threats as their primary area of concern (from 21% a year ago to 25% now).

Primary Cybersecurity Concerns

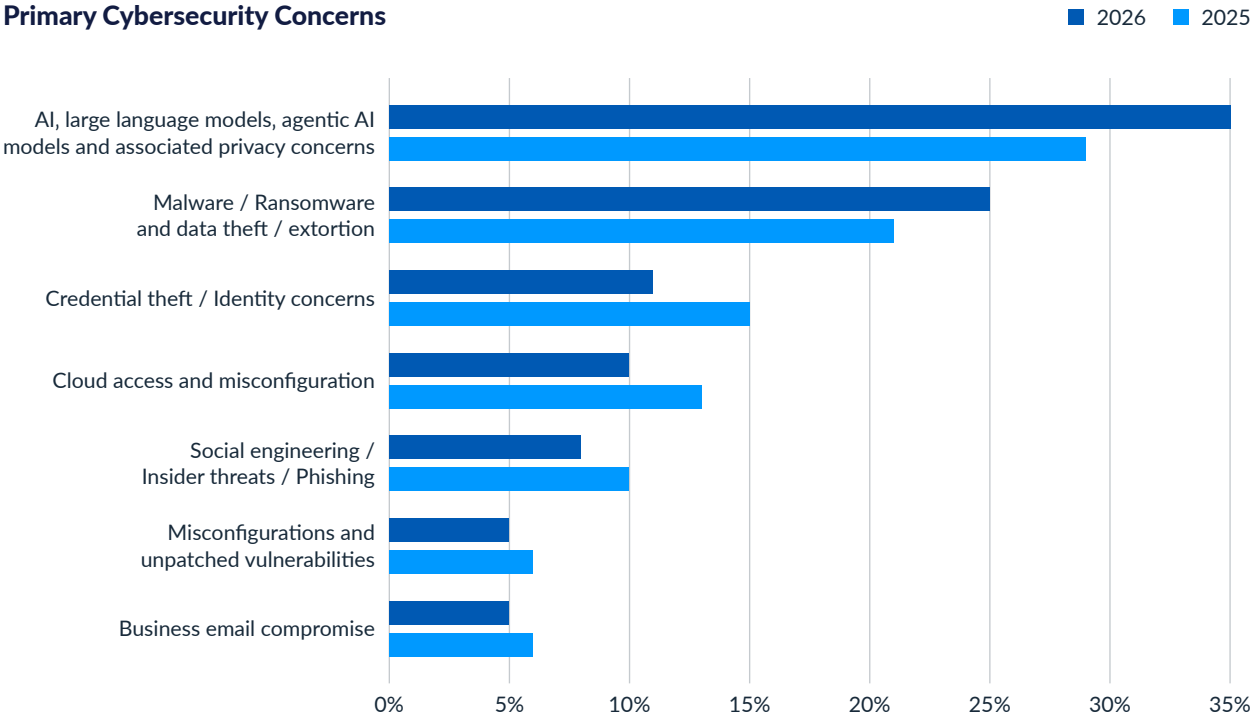

As we noted in last year's report, it's encouraging that security leaders take the threats associated with AI very seriously: from **attack campaigns powered by AI-discovered zero-days** to more mundane (but all-too-common) **accidental privacy breaches**, AI introduces a whole new set of risks. And AI agents have already shown the propensity **to go somewhat off the rails** when not diligently supervised.

However, an often-overlooked risk surrounding AI is that the attention it receives is distracting leaders from more familiar business risks. For example, while more leaders reported that ransomware is their primary cybersecurity concern compared to a year ago, fewer said the same of credential theft/identity concerns, cloud access and misconfiguration, and social engineering.

All three of these areas remain legitimate risks for today's organizations, and the challenge for security leaders is to simultaneously develop and implement plans to safeguard against emerging risks (like those associated with AI) without neglecting to safeguard against historical risks that remain just as real today.

Drilling Deeper

In general, smaller companies (fewer than 100 employees) are more concerned about traditional threats (33% malware and ransomware vs 24% AI and LLM), while medium to large organizations (100+ employees) are more concerned about AI and LLM-related threats.

One interpretation of these results is that smaller companies are prioritizing keeping the operational lights on, while larger organizations are sufficiently mature to view AI through a risk-based lens.

And here's an interesting sidenote: the travel and transportation industry appears in the data as a significant outlier, with only 7% of respondents identifying AI as their primary concern, while 32% pointed to credential theft and identity-related risks.

From a practical standpoint, these findings make sense. Organizations in this sector often manage customer accounts that contain both financial value and sensitive personal information. Rewards points can be stolen and resold if accounts are compromised, while many travel providers also store digital copies of passports and other identity documents, making these accounts particularly attractive targets for cybercriminals.

TREND 2

Significant Cybersecurity Incidents and Long-Lasting Disruptions Remain All Too Common

Focusing our survey on leaders (especially owners, C-level executives, vice presidents, and directors) allows us to tap into an authoritative and informed population. If something happened that threatened the organization's ability to operate, these people will know about it.

In this year's survey, 63% of respondents were certain that their organization had experienced at least one significant cybersecurity incident in the previous 12 months (where significant was defined as "any incident resulting in financial loss, time loss, data loss, operational disruption, reputational impact, regulatory or legal exposure, requirement for security policy/program revisions, or the need for outside assistance in recovery.")

Only 29% of respondents were confident their organization had avoided such a scenario, on par with the 27% result in last year's Trends Report. These results highlight the sobering reality that significant cybersecurity incidents represent the norm for today's organizations, rather than the exception.

Let's focus for a moment on the remaining 7%: those leaders who believe that an incident may have occurred, but that it went undetected. To make informed cybersecurity decisions, leaders need clear answers, and that requires unambiguous detection mechanisms and the investigative capacity to create certainty. Absent such measures, risk-based prioritization becomes a guessing game, with potentially severe consequences for those who make the wrong guess.

Cybersecurity incident occurrences

Critically, these incidents aren't mere hiccups and they don't come and go quickly:

For roughly five out of every six victimized organizations, each significant incident was considered disruptive to the organization, resulting in a loss of time or productivity. That means only one out of every eight escaped relatively unscathed — very poor odds for those who might be willing to take a chance.

For nearly half (48%) of victimized organizations, the loss of time or productivity lasted for two weeks or more, and for almost 10% of incidents, respondents reported disruptions lasting two quarters or longer.

These productivity losses underscore not only the importance of threat prevention, detection, and mitigation, but also of a sound incident response (IR) plan that enables quick recovery and restoration.

In a world of ever-increasing risk, one in which an organization is more likely than not to experience a significant cybersecurity incident, an IR plan isn't a luxury. It's a necessity.

Loss of Productivity

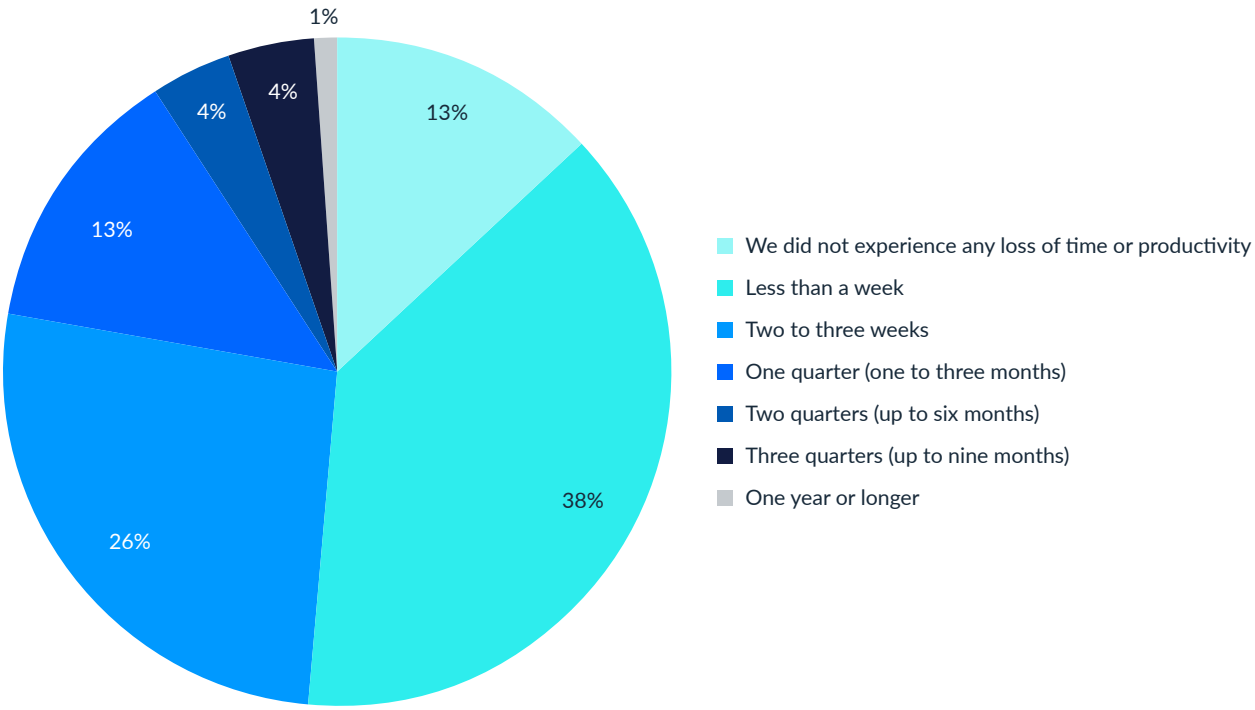

Drilling Deeper

69% of financial services respondents were certain their organization suffered a significant cybersecurity incident, the highest rate among the industries with significant representation in the survey. Another 8% of financial services respondents believed that an undetected incident may have occurred.

Travel/transport is once again an outlier, as only 34% of respondents were certain their organization suffered a significant cybersecurity incident. Worryingly, a further 22% believed that an undetected incident may have occurred, suggesting that more than one fifth of these organizations simply don't have the visibility or incident investigation capability to be certain.

Data Incidents are the Most Impactful, But Ransomware Remains a Costly Menace

In the Arctic Wolf 2026 Threat Report, we revealed that data incidents accounted for 22% of our IR cases, a remarkable 11x increase over the previous reporting period.

Unfortunately, not only are data incidents increasing in frequency, but our survey respondents also report that such incidents have a significant impact on the victimized organization.

First, 21% of those surveyed indicated that data exposure represented the most significant/impactful type of cybersecurity incident their organization experienced in 2025. It's important that readers recognize that this type of incident is distinct from ransomware-related data theft, and instead captures instances in which an organization unintentionally or inadvertently:

- Left data exposed, resulting in its acquisition by an external actor
- Shared or uploaded sensitive data to unauthorized parties (e.g., a broader-than-intended email distribution, mistakenly uploading sensitive data into an LLM)

Second, a further 8% of respondents reported that an attack-based data exfiltration (again, exclusive of ransomware) had the largest impact. In this scenario, a malicious actor deliberately broke into a protected IT environment and stole the data. While sometimes the work of an individual, recent years have seen some threat actors (e.g., Silent Ransom) abandoning ransomware and encryption altogether to focus purely on data exfiltration and extortion in hopes of better net returns.

Most significant cybersecurity incident experienced in 2025

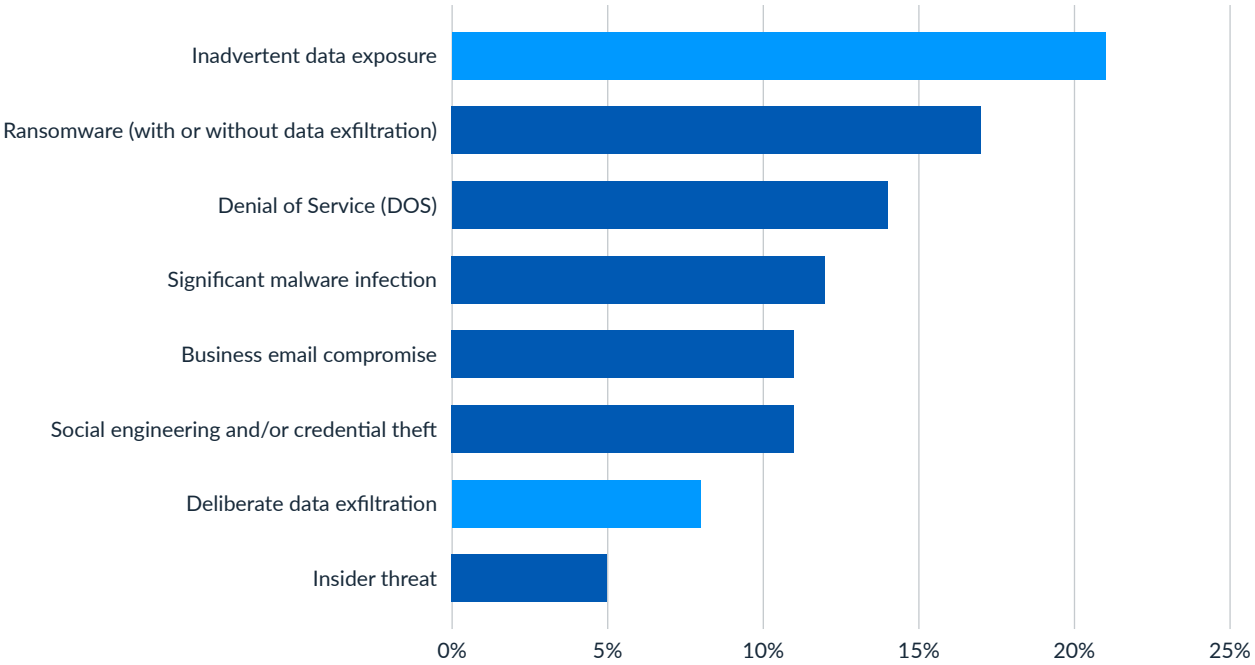

Unsurprisingly, ransomware remains a disruptive threat, with 17% of respondents citing a ransomware incident (with or without data exfiltration) as having the largest impact.

At Arctic Wolf, our position aligns with the general recommendations of the FBI, other law enforcement agencies, and governments: If possible, ransom demands should not be paid, as starving the perpetrators is the only way we can collectively hope to eliminate these attacks.

Nevertheless, the decision on whether to pay is one that must be made by stakeholders within the victim organization once presented with all possible information and options.

Of those respondents victimized by ransomware, more than half, at 56%, reported paying some portion of the ransom demand. While this is disheartening, as it runs counter to the cybersecurity community's effort to make ransomware a thing of the past, responsibility for that decision was not always one that was strictly internal. In half of these cases, it was the organization's own decision makers who chose to pay. Our results, however, show that in the other half of these cases it was a third party, such as a cyber insurer acting on the organizations' behalf that chose to make some form of payment.

The contrast with Arctic Wolf's own incident response data is striking. According to our 2026 Threat Report, just 23% of ransomware IR cases handled by Arctic Wolf's Incident Response Team involved any form of ransom payment, meaning 77% were resolved without a single dollar paid to the attacker. According to Kerry Shafer-Page, Arctic Wolf's Vice President of Incident Response, this gap reflects a troubling pattern: Organizations working without expert guidance too often rush to pay, viewing it as the fastest path to resolution. It rarely is. Many threat actors are seasoned negotiators who exploit urgency and fear, frequently making alternative resolution approaches both more cost-effective and more prudent.

Ransomware attacks

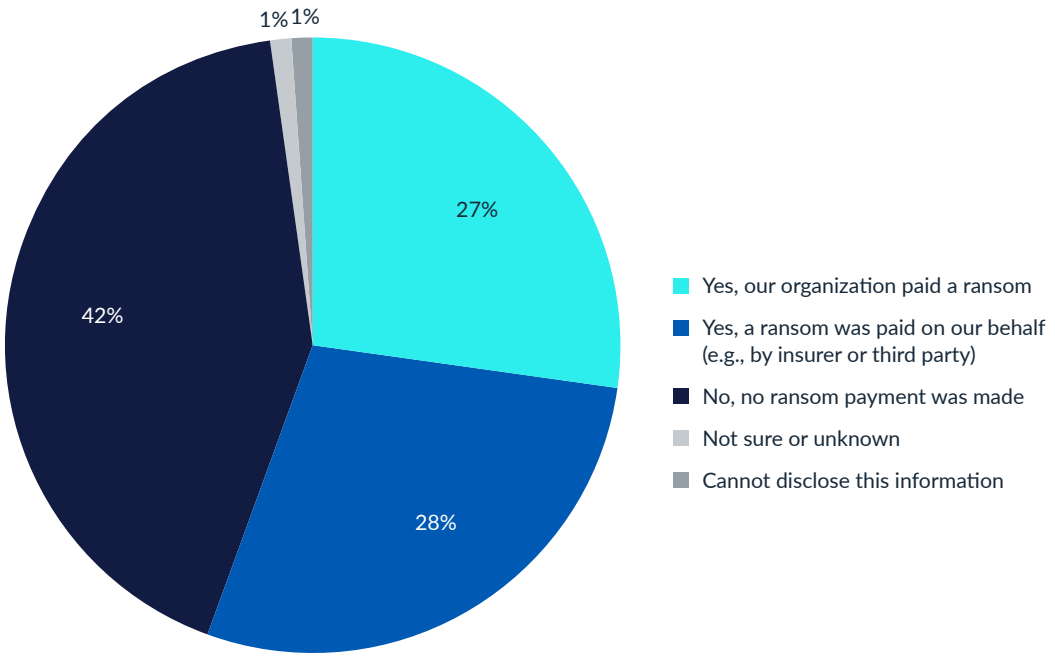

In general, we believe that engaging with a ransomware actor is best left to the experts, as they generally have a great deal more experience with handling these events than any in-house personnel.

However, the outcomes may vary, due to factors like the degree of disruption inflicted, the organization's IR capabilities, the skill of the negotiator, and the disposition of the attacker.

One More Thing...

Before we move on, it's worth considering that the first graph from this section shows that a meaningful proportion of the survey respondents experienced significant disruption from eight distinct types of cybersecurity incident.

This response distribution underscores a point we've long emphasized: that a strong security posture cannot result from focusing on one particular type of threat at the expense of the long list of others. Yes, security investments should be prioritized based upon factors including the general threat landscape, prominent threats against your industry, your risk tolerance, and other relevant factors. But care must also be taken not to lose sight of the fact that there are many different types of threat, each of which has the potential to disrupt your organization.

TREND 4

Regulators and Insurers Remain the Driving Forces Behind Incident Disclosure

In prior reports, we found significantly high rates of disclosures for data breaches (96% in 2025 and 97% in 2024).

This year, we expanded the scope of our questioning to include the disclosure of any/all significant cyber incidents, rather than those that met the stricter criteria of a breach. Yet even with this expanded scope, 94% of respondents still indicated that their organization disclosed their cybersecurity incident.

Consistent with our past surveys, this high disclosure rate is due mostly to obligations:

45% of respondents reported disclosing due to a legal requirement to do so (versus 51% last year)

36% cited a requirement from either an insurance provider or other outside entity (versus 34% last year)

However, 13% of organizations elected to disclose simply to share their findings with the broader security community. Such selfless transparency in the face of potential negative consequences is critical in the collective fight against cybercrime, and a collaborative approach is especially so, as AI-powered vulnerability discovery is simultaneously increasing the number of vulnerabilities and shortening the time to exploitation.

Cybersecurity incident disclosure

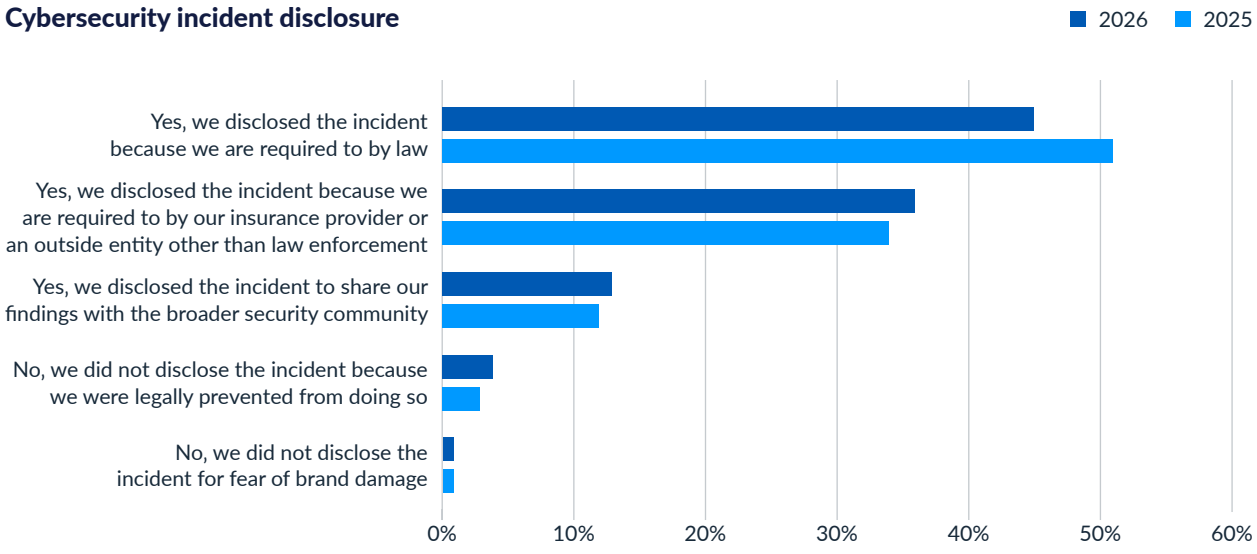

02

SECTION 2:

Readiness Trends

Data Transformation and AI Initiatives Are Driving 12-Month Cybersecurity Strategies

We already saw that AI broadly, and specifically LLMs, agentic AI models, and their associated privacy risks collectively represent the primary cybersecurity concern for our survey respondents.

Consistent with this view, two closely related AI-centric needs emerged as the primary drivers of cybersecurity strategies in the next 12 months.

Topping the list for the second year in a row, 41% of respondents selected **"Data transformation and secure AI adoption."** This survey option focuses on the implementation details of deploying AI utilities, including how data is handled (e.g., workflows pertaining to training data, information repositories, and data pipelines) and issues relating to Identity and Access Management (IAM), which itself encompasses authentication and authorization for employees and the AI agents themselves.

However, from listening to our customer base about their AI initiatives, we knew that operationalizing AI is only one piece of the bigger picture. Many leaders are under tremendous pressure to prevent their organization from falling behind competitors or the broader technological landscape. Accordingly, we added a new option to this year's survey: **"Keeping pace with rapid AI developments."** Reflecting how quickly the landscape changes, 35% of those surveyed selected this option.

Primary Drivers for Cybersecurity Strategy for the next 12 months

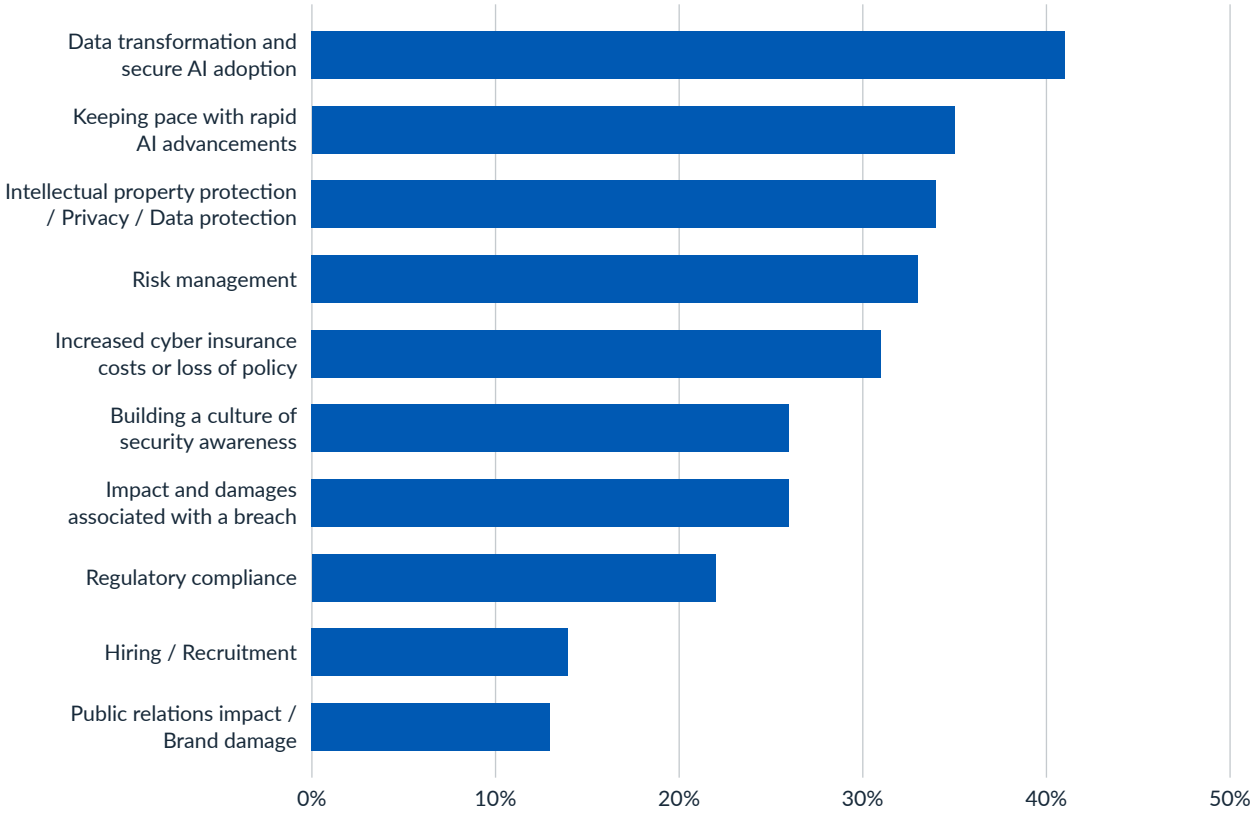

Of course, priorities may differ due to a wide range of factors, and when we look closer we see that the main drivers of 12-month cybersecurity strategies vary considerably by industry. For example:

Risk management topped the ranking for healthcare and travel organizations.

Governmental organizations are strongly influenced by cyber insurance and compliance, and are comparatively less driven by the need to keep pace with AI advancements.

Energy/utilities and travel/transport organizations are especially wary of disruptive breaches.

Two areas that are quite consistent across industries?
Hiring/recruitment and PR/brand issues have very little influence over cybersecurity strategies.

	Energy / Utilities	Financial Services	Government / Public sector / Education / Social services	Healthcare / Life sciences and pharmaceutical	IT / Technology / Telecoms company	Manufacturing	Retail	Travel / Transport
Data transformation and secure AI adoption	1	1	3	4	1	1	1	1
Keeping pace with rapid AI advancements	5	3	6	1	2	4	1	7
Intellectual property protection / Privacy / Data protection	2	2	5	3	3	2	4	3
Risk management	3	5	2	1	4	3	5	1
Increased cyber insurance costs or loss of policy	6	4	1	5	6	4	3	4
Building a culture of security awareness	7	7	6	7	5	7	8	4
Impact and damages associated with a breach	4	5	8	7	7	6	6	4
Regulatory compliance	8	8	3	5	8	8	7	8
Hiring / Recruitment	10	9	9	9	9	9	10	9
Public relations impact / Brand damage	9	9	10	10	10	10	9	10

Leaders Have Misplaced Confidence that Their Team Can Keep Up With Today's Threats

In an ever-evolving threat landscape, especially one being reshaped by AI-powered vulnerability discovery and automation, acquiring and operationalizing up-to-date knowledge of cyber risks are essential.

Slightly more than half (53%) of leaders surveyed for this report are very confident that their respective security teams can keep up with the volume and complexity of today's threats. A further 43% are somewhat confident, which brings the sum of those expressing confidence to 96%.

At this point, it's worth recalling that 70% of respondents believed that their organization had experienced at least one "significant" cybersecurity incident in the previous 12 months.

In fact, confidence levels are actually higher in organizations that experienced such an incident: 57% of leaders from victimized organizations were very confident, compared to 47% from leaders in organizations that did not suffer a significant incident. While it's possible that suffering a disruptive attack led these organizations to dramatically increase their proactive defensive measures, it seems at least as plausible that leaders are succumbing to a survivorship bias.

In a less-confusing finding, confidence is also much higher in those leaders whose organizations have an active IR retainer (57% of whom are very confident) versus those lacking this outside expertise (only 40% of whom are very confident).

Confidence in team capabilities for managing threats

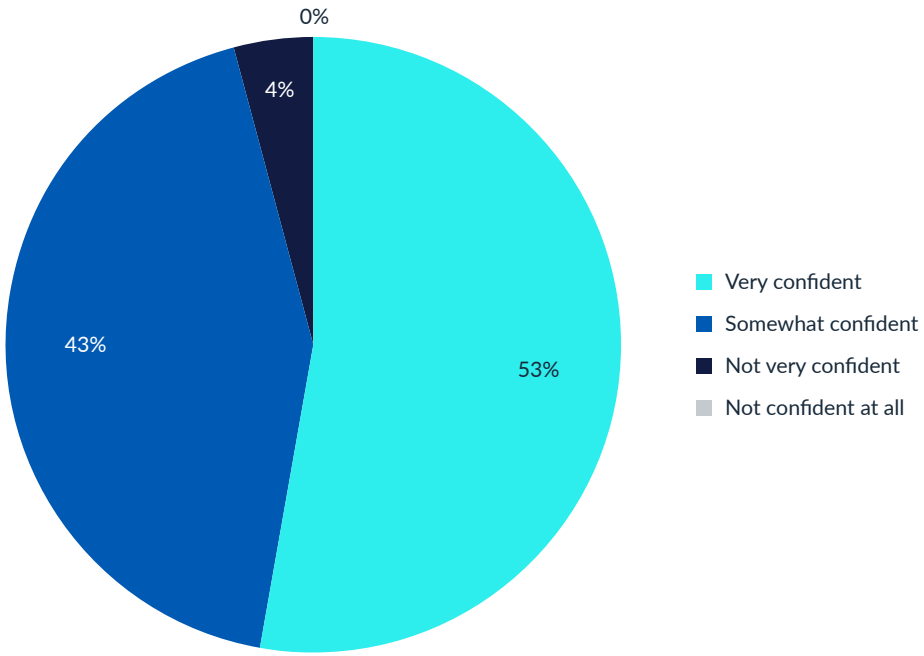

As cyber threats continue to grow in volume and sophistication, organizations face increasing pressure to detect malicious activity as early as possible. For important context, recall that the Arctic Wolf 2025 Security Operations Report drove home the reality that threat actors don't schedule their activities to coincide with their targets' working hours. During the period covered by that report:

51% of all alerts were generated outside of traditional working hours

1/6 or nearly that, of each week's total alert volume occurred on the weekends

Accordingly, achieving and maintaining full visibility across the IT environment are necessities, as even minor visibility gaps can result in delayed or missed detections.

And yet, we continue to see that organizations simply aren't doing enough to protect themselves.

A slight majority of respondents (57%) were confident that their organization had sufficient internal resources to monitor their IT, cloud, and network environments around the clock. An additional 32% report that they have the same coverage through reliance upon a third party.

Again, these results should be considered in the context of 70% of respondents also reporting experiencing at least one disruptive cyber attack, suggesting that whatever monitoring was in place may have been sufficient to detect the attack, but clearly did not lead to effective containment.

Moreover, fully 11% of respondents indicated that their organization lacked 24x7 monitoring capabilities.

Resources to monitor organizational IT, cloud and network 24×7

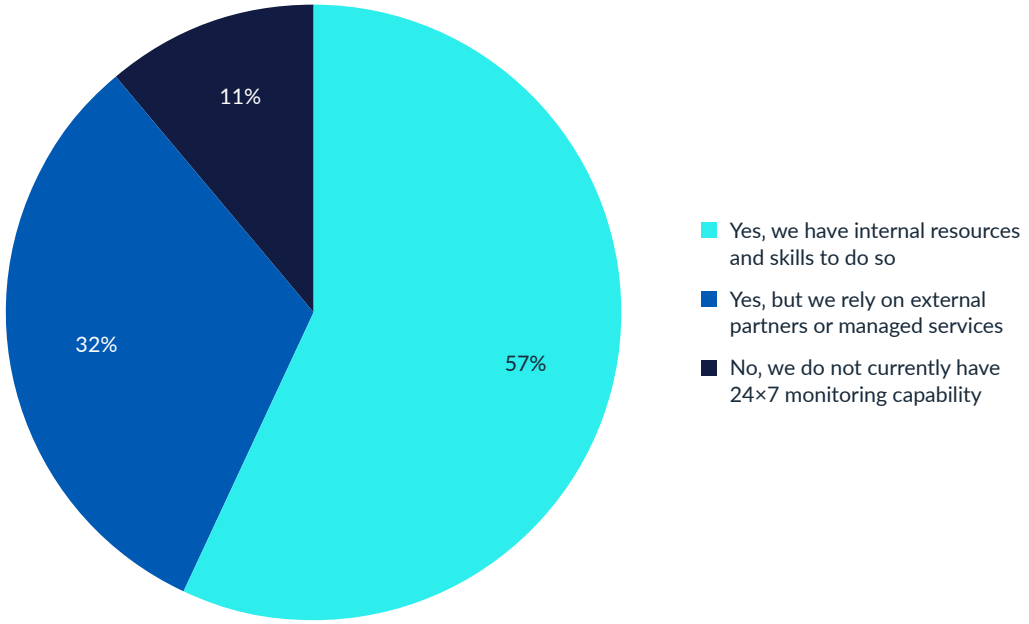

Looking deeper than the aggregate numbers, we see some significant variation across different industries, both in terms of professed coverage and in how that coverage is achieved.

Fully 95% of retail respondents reported having 24x7 coverage, topping the list by just edging out energy/utilities services and technology companies. Conversely, only 79% of respondents from healthcare organizations reported around-the-clock coverage, while 80% of leaders from public sector organizations said the same. In other words, one in five healthcare and governmental organizations lack 24x7 monitoring of their IT environments, which is alarming.

Healthcare leaders were also the only group to report a higher reliance upon third-party providers (41%) versus internal resources (38%). At the other end of the spectrum, financial services companies were the least reliant upon external parties, with 64% equipping and trusting their internal staff to perform around-the-clock monitoring.

Resources to monitor organizational IT, cloud and network 24×7

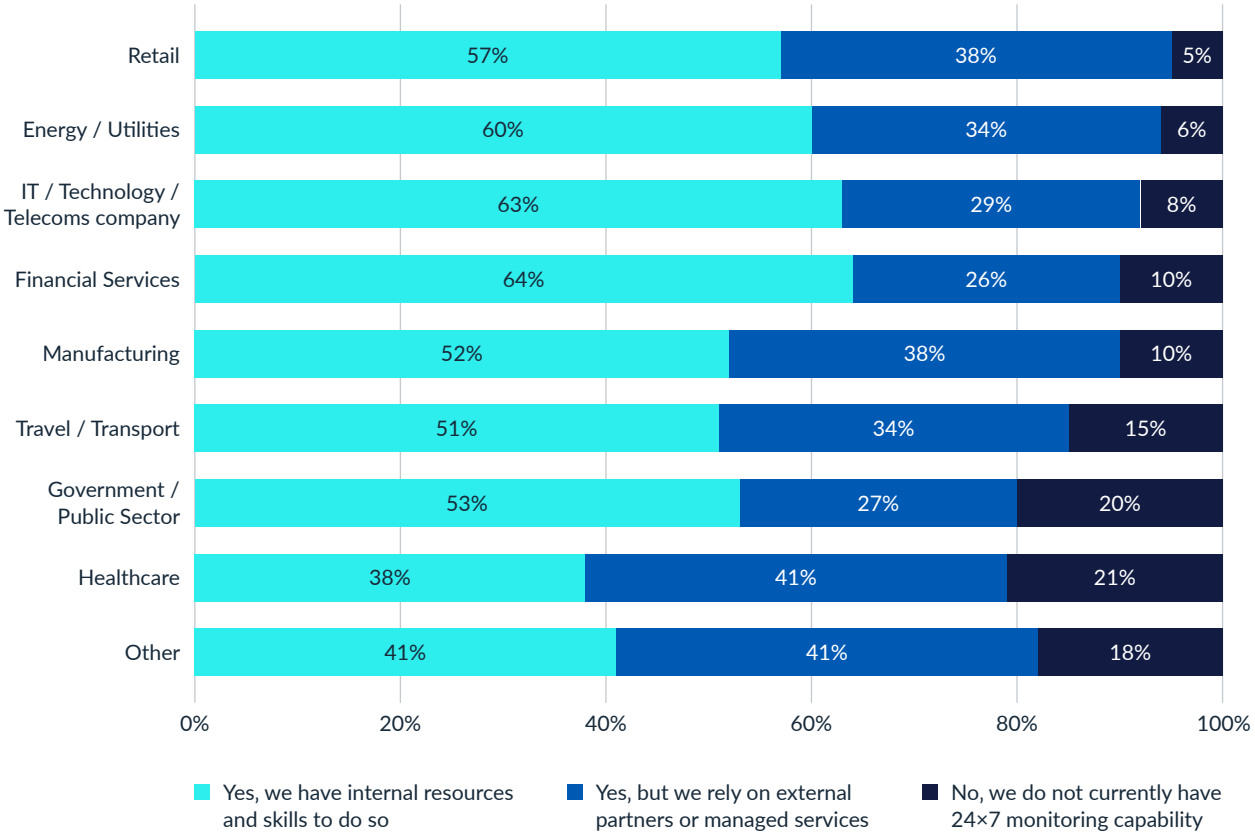

TREND 3

Cybersecurity Teams Are Stretched Thin, and Leaders See Automation and AI as Potential Solutions

IT and cybersecurity teams are under constant pressure both to safeguard the organization against today's attacks while also preparing the organization to withstand the attacks of tomorrow.

Practically, this requires striking a balance between meeting the operational and tactical demands of right now and devoting enough time to strategic research, planning, and solution sourcing for the demands of the future.

To better understand this division of attention, we asked leaders how much time their personnel spend each week on nine security-related tasks.

The results stand out not for highlighting a hierarchy of priorities, but for astonishing uniformity: Leaders report that their respective teams spend between 13 hours (for activities at the bottom of the figure) and 15 hours (for activities at the top of the figure) per week on each activity.

If anything, these results highlight the risk of competing priorities and strongly suggest that security teams are stretched thinly across a long list of important and essential activities.

Time spend on different tasks

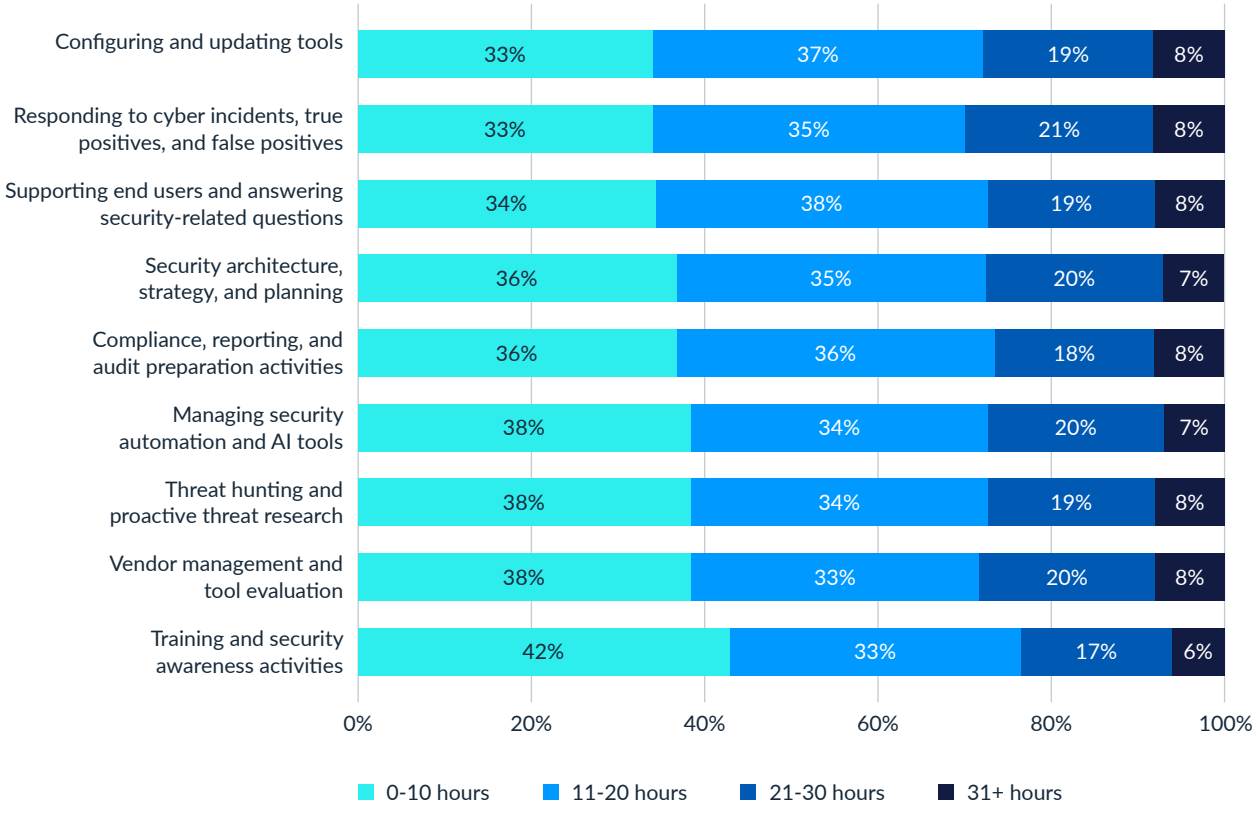

In general, the smaller the respondent's organization, the less time they spend on these activities, although the variation is quite small.

However, we do see some interesting variations when we divide responses by industry. For example, governmental organizations spend the most time on every one of these activities, while travel/transport organizations spend the least overall (and on eight of the nine, individually).

	Energy / Utilities	Financial Services	Government / Public Sector / Education / Social services	Healthcare / Life sciences and pharmaceutical	IT / Technology / Telecoms company	Manufacturing	Retail	Travel / Transport
Security architecture, strategy, and planning	15.6	16.5	16.7	14.3	16.3	13.8	13.7	13.5
Managing security automation and AI tools	16.7	16.3	17.9	14.9	16.5	14.1	15.3	13.6
Supporting end users and answering security-related questions	15.1	15.1	17.7	16.2	16.7	14.3	16.2	13.3
Training and security awareness activities	14.0	15.6	16.6	15.4	15.5	14.4	15.9	12.0
Threat hunting and proactive threat research	15.6	15.4	16.9	14.4	15.7	13.3	14.0	14.0
Configuring and updating tools	16.4	16.3	18.6	13.0	15.7	13.8	14.6	12.4
Responding to cyber incidents, true positives, and false positives	15.3	15.4	17.7	13.0	16.6	13.1	15.2	12.5
Compliance, reporting, and audit preparation activities	15.3	15.8	16.1	14.9	15.6	13.8	14.5	12.5
Vendor management and tool evaluation	12.6	14.9	15.0	13.7	14.7	13.6	13.9	11.0
Mean	15.2	15.7	17.0	14.4	15.9	13.8	14.8	12.7

For many years, organizations worldwide have grappled with a cybersecurity skills gap that manifests at both the macro level (i.e., a general shortage of qualified professionals) and the micro level (i.e., challenges building an internal team with the requisite skills).

To address this problem, leaders point to two primary strategies:

1. **Investing in automation or AI to reduce manual workload** (chosen by 59% of respondents)
2. **Investing in training and certifications** (chosen by 57%)

Executing on these two strategies should allow existing team members to better apply their skills and finite time while also expanding their capabilities.

However, these courses of action are not without their own unintended consequences.

For example, the introduction of AI, AI-powered tools, and automation is rarely straightforward. More often than not, doing so effectively requires redesigning workflows, strengthening IAM controls, modifying IT environments, updating data governance guidelines, and introducing safeguards (just to list a handful of considerations). The risk, then, is that AI isn't reducing the workload, per se, but is simply rearranging or redistributing it. Readers of a certain age may recall the early SIEM years. Yes, these tools were powerful and gradually became essential; however, the costs and operational headaches associated with deploying, configuring, and maintaining them often caught early adopters by surprise.

And while investing in training and certifications is an effective means of upskilling your employees, it also makes them more attractive to other organizations. Consequently, paying for training and certifications should be accompanied by increased compensation. This approach increases costs, but also increases the likelihood that your organization will be the one to benefit from your employees' new knowledge and capabilities.

In contrast, an outsourcing-centric approach allows your organization to tap into cutting-edge expertise without the risk that it could walk out the door the moment it receives a better offer

Actions addressing cybersecurity skill gaps

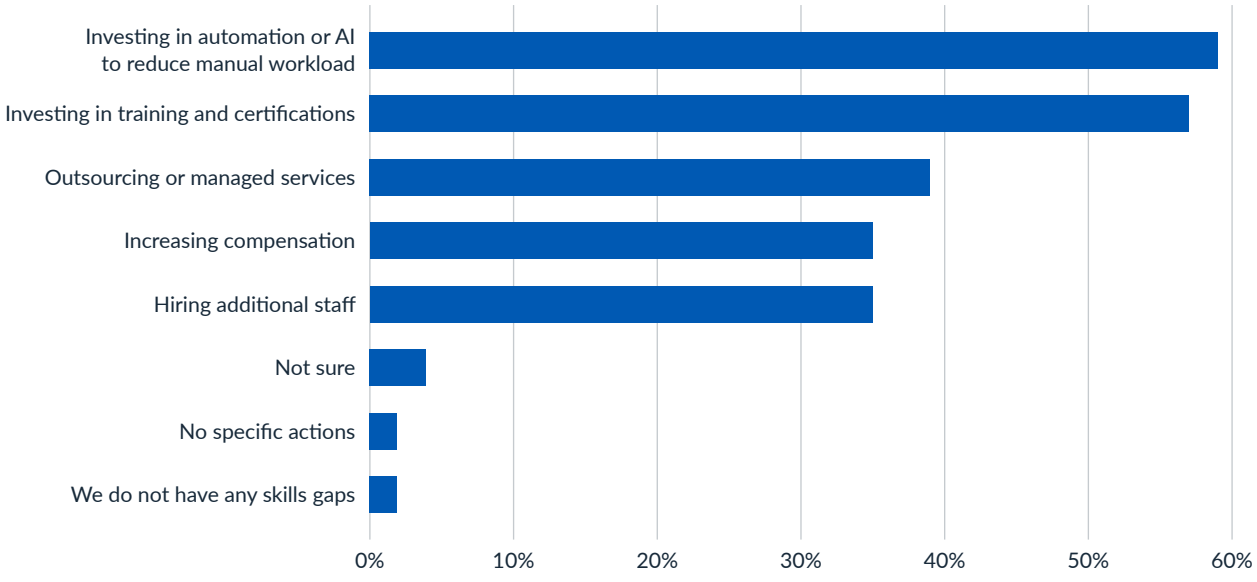

Most Organizations Have an IR Retainer, and Most Retainers Continue to Be Used

Incident response (IR) is a set of processes and tools used to identify, contain, and remediate cyber attacks, and to restore the organization to pre-incident operations.

Important functions include:

- Securing an environment by eliminating the threat actor's access
- Analyzing the cause and extent of the threat actor's activities while inside the network
- Restoring the network to its pre-incident condition (including ransom negotiation and payment, if required)

Ideally, each function is performed concurrently and is complemented by information, insights, and outcomes emerging from the others.

When we surveyed IT and security leaders two years ago, we found that 64% of organizations had invested in an IR retainer. In this year's survey, that figure has grown to 74%. This strongly suggests that the IR retainer model has been embraced by organizations seeking to transfer risk and ensure support during an incident.

Maintenance of IR retainers and working with IR firms

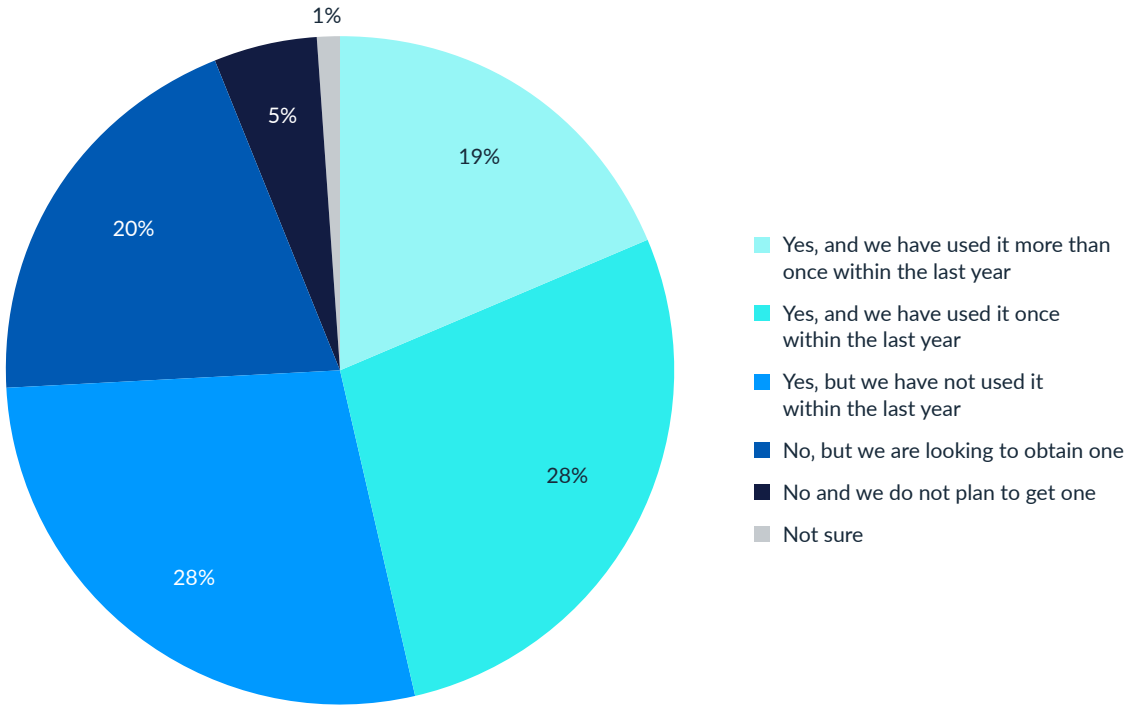

Moreover, the survey responses indicate that 62% of organizations with an IR retainer or contract put it to use at least once in the last year.

The most frequently cited reason for doing so (selected by 59% of respondents) was in response to an actual cyber incident, while 50% of respondents reported turning to the IR vendor for support with a non-incident-related scenario (e.g., training, tabletop exercises, etc.).

Reasons for utilization of IR retainer funds

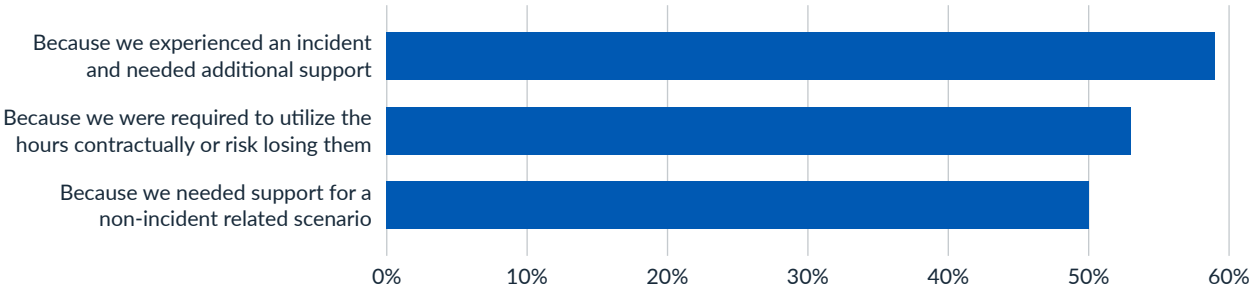

However, it's important that leaders looking for an IR provider understand that there are different varieties of retainer.

For instance, traditional IR retainers are built on outdated, hours-based models that leave many organizations scrambling to "use it or lose it" by year's end. Slightly more than half (53%) of leaders who participated in our survey found themselves in this position, and that's hardly the way to get the most out of your IR budget.

And what happens when a major incident strikes your organization but your hours are already exhausted? Be sure to understand exactly what it is you're getting (or not) for your money.

Cyber insurance carriers will usually approve the use of reputable IR vendors that aren't included on their preapproved list. To eliminate delays when IR services are needed, it's recommended that organizations obtain written confirmation of this approval from their insurance carrier once a retainer is established and before experiencing an incident.

IR costs are usually covered (less policy deductibles) by cyber insurance, but you'll want to review your specific policy and/or endorsements to determine how much coverage you may have.

Organizations that choose prepaid retainers are often looking for preferred access to IR teams or want to negotiate a lower hourly rate on a large block of hours.

Before purchasing a prepaid IR retainer, verify with your insurance carrier that the prepaid hours are covered by your cyber insurance policy. **Spoiler alert: they probably aren't.**

Reimagining incident response

Our assurance-based approach to incident response addresses limitations of traditional retainers, providing predictable coverage, proactive readiness, and integrated tools:

- **JumpStart Retainer:** Entry-level plan with a 4-hour SLA, access to IR planning tools, runbooks, and a cyber resilience assessment (CRA)
- **Incident360 Retainer:** Includes everything in JumpStart, plus a 3-hour SLA, full incident coverage for one incident, and threat intelligence access (restoration services such as reinstalling app and software and patching are capped at 30 hours)
- **Incident360 Retainer Plus:** Adds a tabletop exercise, IR plan review, and CRA Review; three readiness touchpoints which provide a comprehensive readiness and response package

"Arctic Wolf's Incident360 Retainer breaks the laws of traditional hours-based incident response retainer models. Cost certainty mixed with an outcome-based model will provide IT, cybersecurity, and C-suite leaders with some reduced anxiety when — not if — their incident response retainer gets utilized in a crisis situation."

CRAIG ROBINSON, RESEARCH VICE PRESIDENT, SECURITY SERVICES, IDC

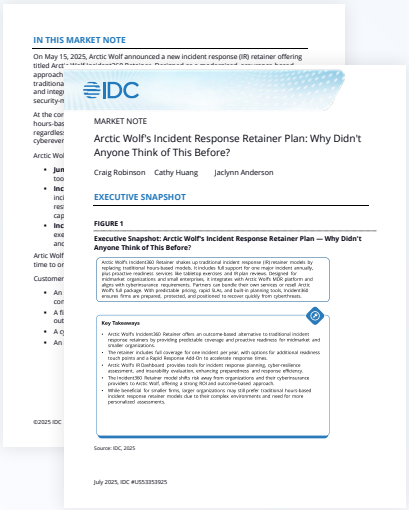
IN THIS MARKET NOTE

IDC

MARKET NOTE

Arctic Wolf's Incident Response Retainer Plan: Why Didn't Anyone Think of This Before?

Craig Robinson Cathy Huang Jaclynn Anderson

EXECUTIVE SNAPSHOT

FIGURE 1

Executive Snapshot: Arctic Wolf's Incident Response Retainer Plan — Why Didn't Anyone Think of This Before?

Key Takeaways

Source: IDC, 2025

July 2025, IDC #US53353925

03

SECTION 3:

AI Trends

As Organizations Aim to Supercharge Security Operations, AI Capabilities Are Becoming Purchase Prerequisites

A rapid, accurate, around-the-clock threat detection capability is a necessary condition for an effective defense, but to enable effective response, detection must be paired with equally rapid and decisive response actions.

However, threat signals are hidden within a tremendous volume of noise: as we detailed in our 2026 Threat Report, Arctic Wolf produces one alert for every 138 million raw data observations, on average. Similarly, threat containment and remediation may require a collection of actions, executed in parallel across many systems, with sufficient speed to outpace nimble attackers who may themselves be aided by automation and AI tooling.

Accordingly, leaders are intrigued by the potential security applications of AI, including using machine learning algorithms to identify threats, utilities to equip human analysts with crucial context and correlated signals, and AI agents to execute workflows and runbooks at machine speed.

In fact, almost every organization represented in our survey is adopting AI in their cybersecurity program:

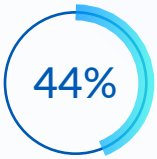

are in the early stages, either evaluating AI use cases (22%) or actively piloting AI-enabled tools (22%)

have deployed AI in limited areas

have comparatively mature deployments in which AI is either embedded within multiple security workflows (20%) or is central to the organization's security operations strategy (14%)

Stage of AI adoption in cybersecurity

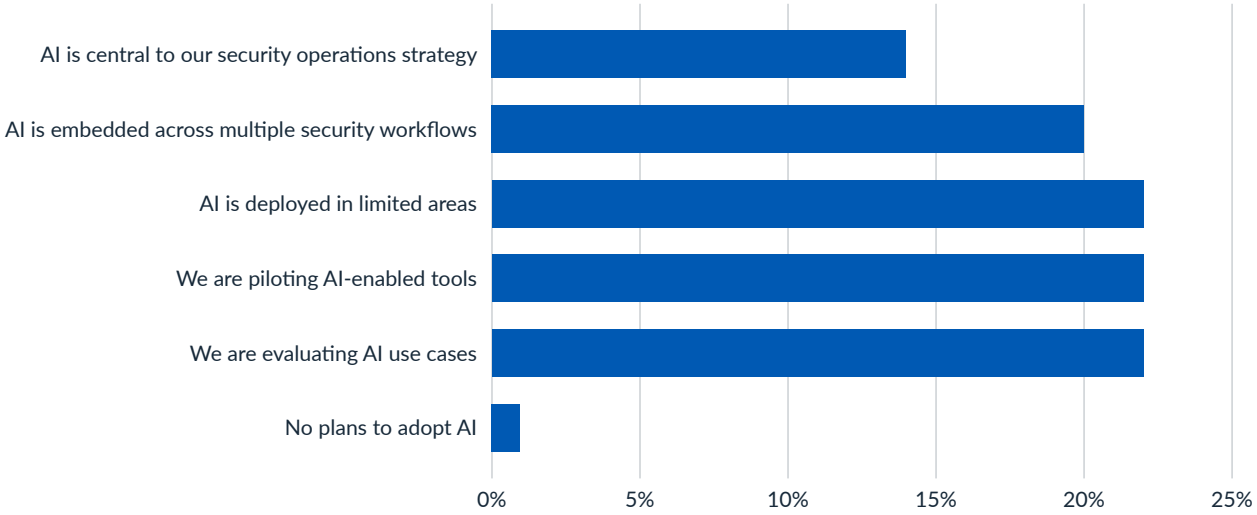

In this context, it's no surprise that decision-makers are carefully evaluating vendors' AI capabilities:

43% of respondents indicated that they strongly consider a vendor's AI capabilities when making a purchase or renewal decision

51% go even further, and regard AI capabilities as a requirement

5% suggest that AI capabilities have little influence on such decisions

However, leaders would do well to recognize that many vendors have simply bolted AI onto existing products, rather than architecting new solutions around AI's strengths. When such products are deployed, security teams are asked to prepare data for AI, redesign workflows for agents, build and test agentic systems, validate outcomes, govern model behavior, and manage performance over time — all of which can contribute to long pilots, uncertain returns, and AI initiatives that struggle to deliver value in production.

Predicted influence of AI-powered capabilities on purchasing and renewal decisions

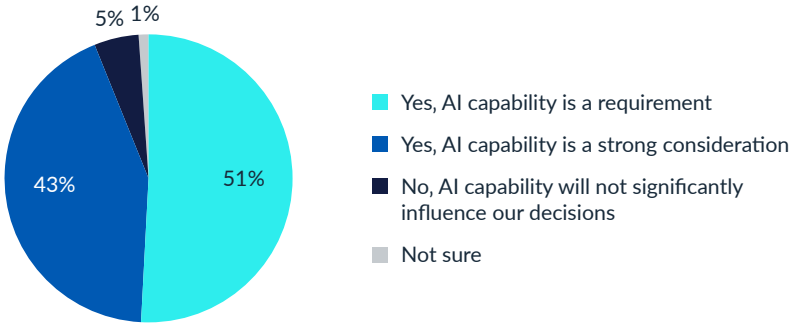

Diving a little deeper into the survey responses reveals that larger organizations, those that recently experienced a significant cyber incident, and those with an active IR retainer are even more likely to consider AI capabilities as a requirement.

The biggest differences of opinion appear when we look at individual industries. For example, financial services providers and the technology sector are much more likely to regard AI capabilities as requirements, whereas healthcare providers, travel companies, and governmental organizations place less importance on such features.

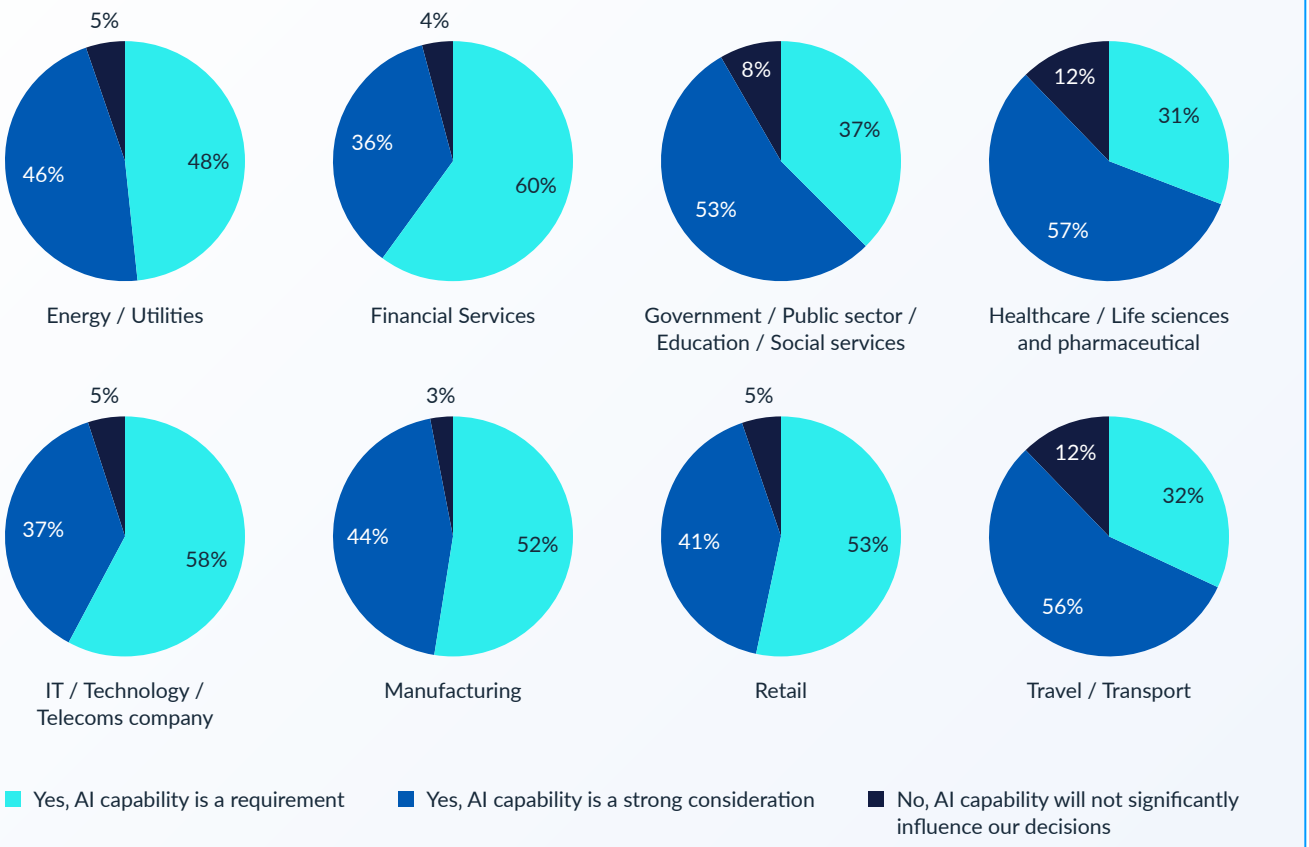

TREND 2

Leaders Expect AI to Enhance Security Operations and Already Consider AI More Capable Than Humans of Improving Their Organization's Cybersecurity

To evaluate the effectiveness of a solution, it helps to have clear expectations of what that solution will do, and the leaders who participated in our survey expect AI to contribute to a collection of security outcomes, with three in particular standing out:

85% of IT and security leaders expect AI to improve their organization's ability to detect new or elusive threats

83% expect AI to correlate and analyze their security data

82% expect that AI will provide a consolidated view of security alerts, remediation steps, and resolution processes

Perceptions of AI's future role in security operations

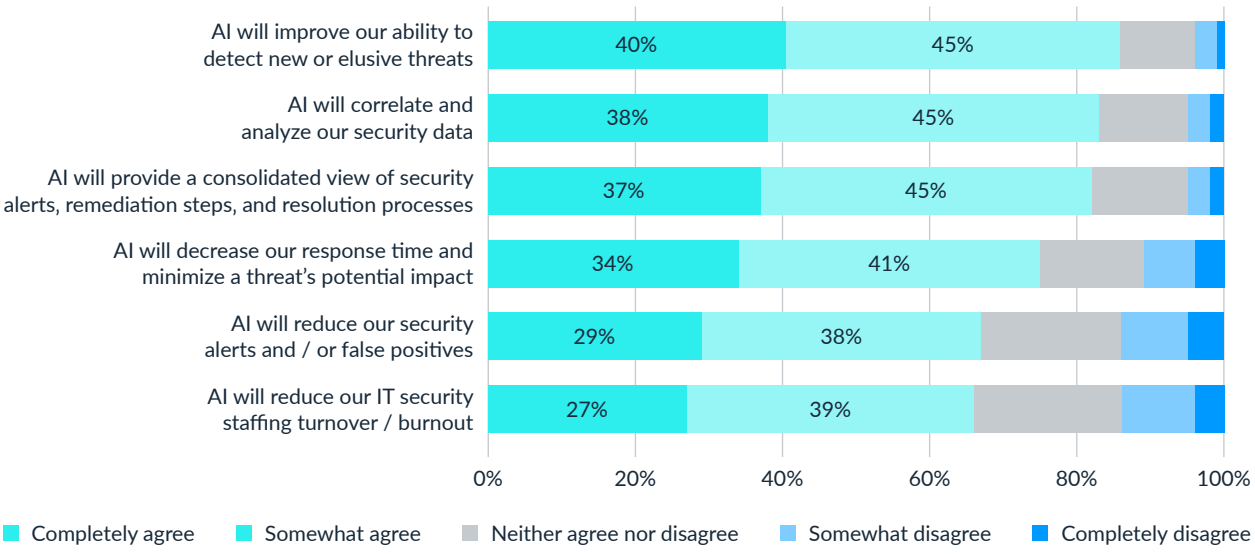

In fact, survey respondents have such faith in AI that they already believe AI is more capable than humans when it comes to improving the organization's cybersecurity.

For four of the security outcomes under consideration, leaders favored AI by more than a two-to-one margin. Notably, these four all relate quite closely to data processing — an area where AI technologies are comparatively mature, and an area of ever-increasing need. For context, across our base of more than 10,000 customers, our data shows that the average customer environment generates nearly 33 billion observations per year.

Things are a bit more even when we move into domains where planning, judgment, and communications are points of emphasis: proactively securing our organization and explaining the context of threats.

AI vs. human capability

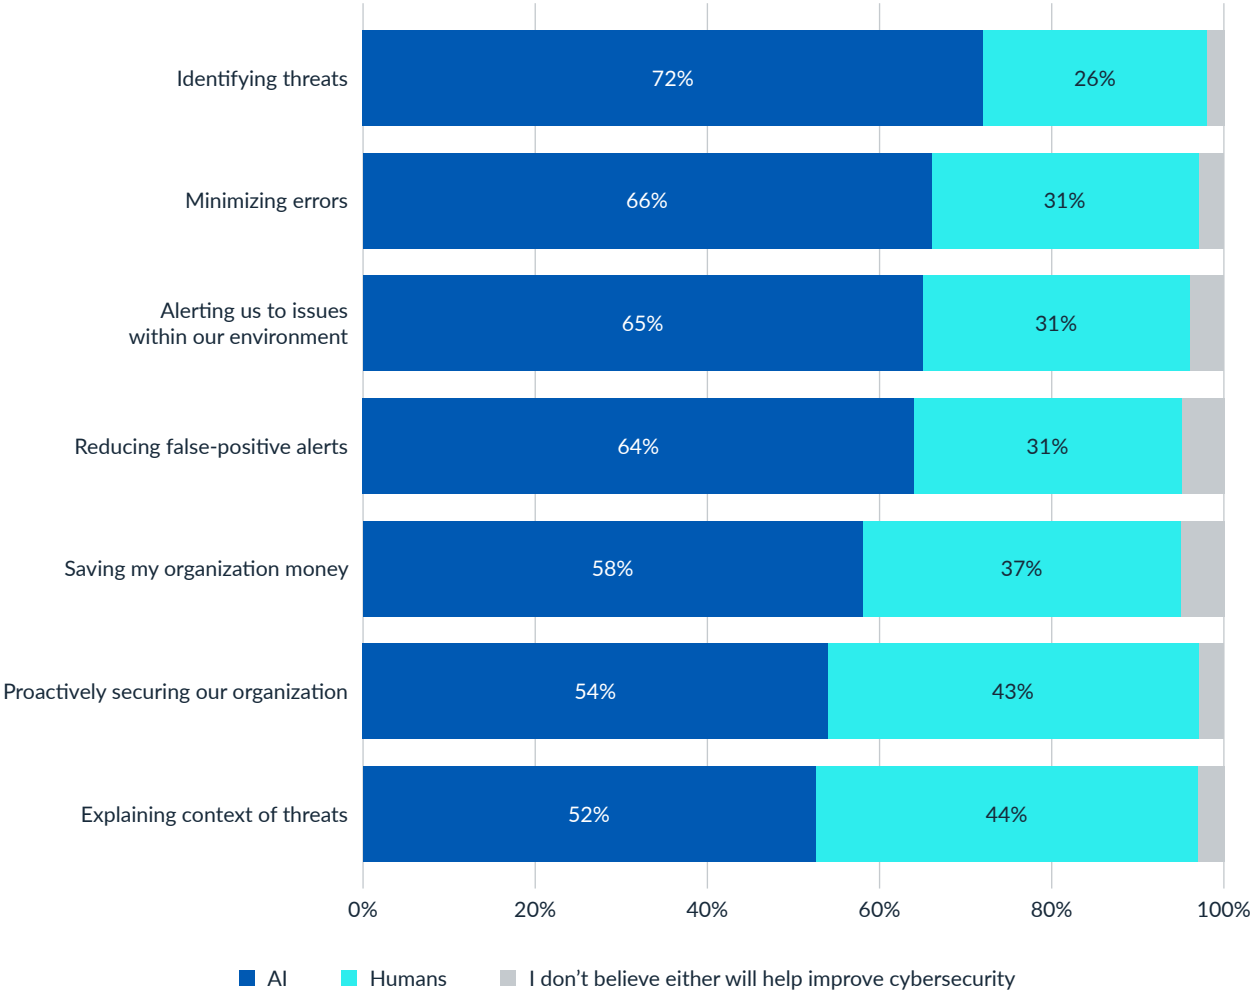

Despite High Expectations for the AI-Powered SOC, Leaders Don't Yet Trust AI to Act Autonomously

Clearly, there's widespread belief among leaders that their organization needs AI.

More specifically, we've seen leaders have high hopes that embedding AI within security operations will have a major positive impact and that there's widespread belief AI can outperform human personnel in contributing to a range of cybersecurity outcomes.

And yet, we also saw that only 14% of organizations have made AI central to their security operations strategy. This suggests that, at least when it comes to handing over the keys to AI and letting agents actually take autonomous action within the organization's environment, there's substantial hesitation.

In fact, blocking malicious IP addresses or domains at the firewall was the only task for which the majority of respondents (although only very slightly, at 53%) trust AI to act without human supervision.

Trusted agentic AI and agent tasks

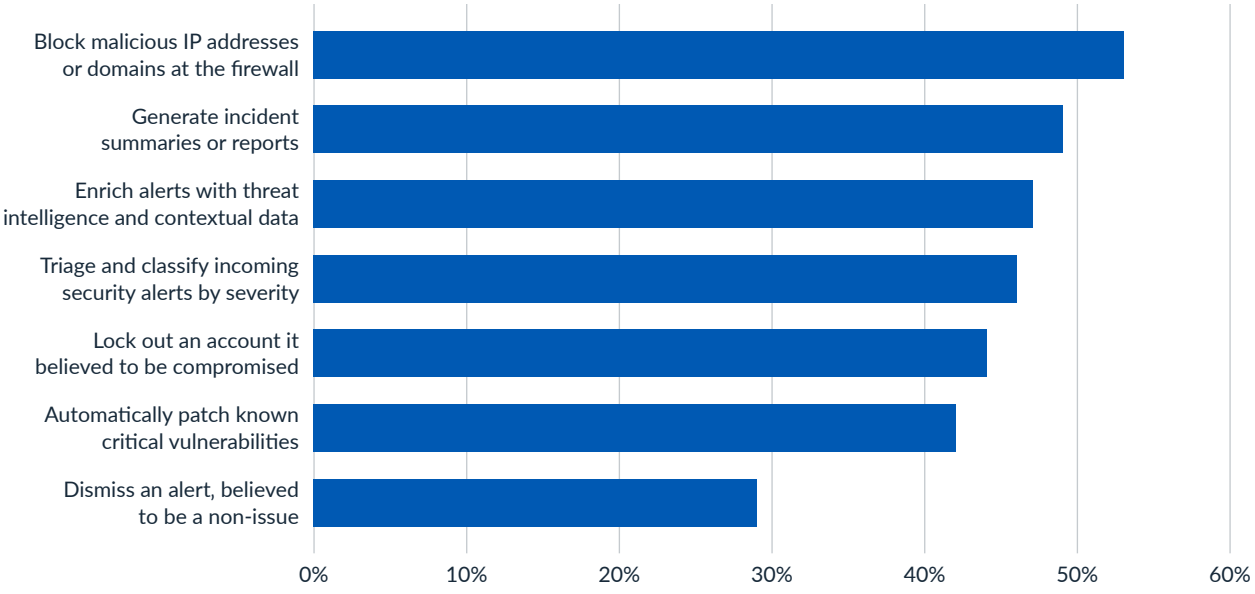

Respondents cited a long list of reasons why they don't yet trust agentic AI to act autonomously within their environment, led by data privacy concerns (51%) and a lack of human intuition (49%).

In today's rapidly evolving environment, security teams need AI for its ability to spot well-hidden threats and to initiate machine-speed response, among other reasons.

And yet, many LLM-powered approaches remain difficult to trust and hard to operationalize. At the same time, concerns about hallucinations, brittle reasoning, a lack of accountability, and signs that costs are poised to soar as AI vendors switch to usage-based billing are absolutely valid until proven otherwise.

In our view, the most effective security operations centers (SOCs) will be those that combine trusted, trustworthy AI with real-world human expertise inside a purpose-built system.

Reasons for distrust in agentic AI

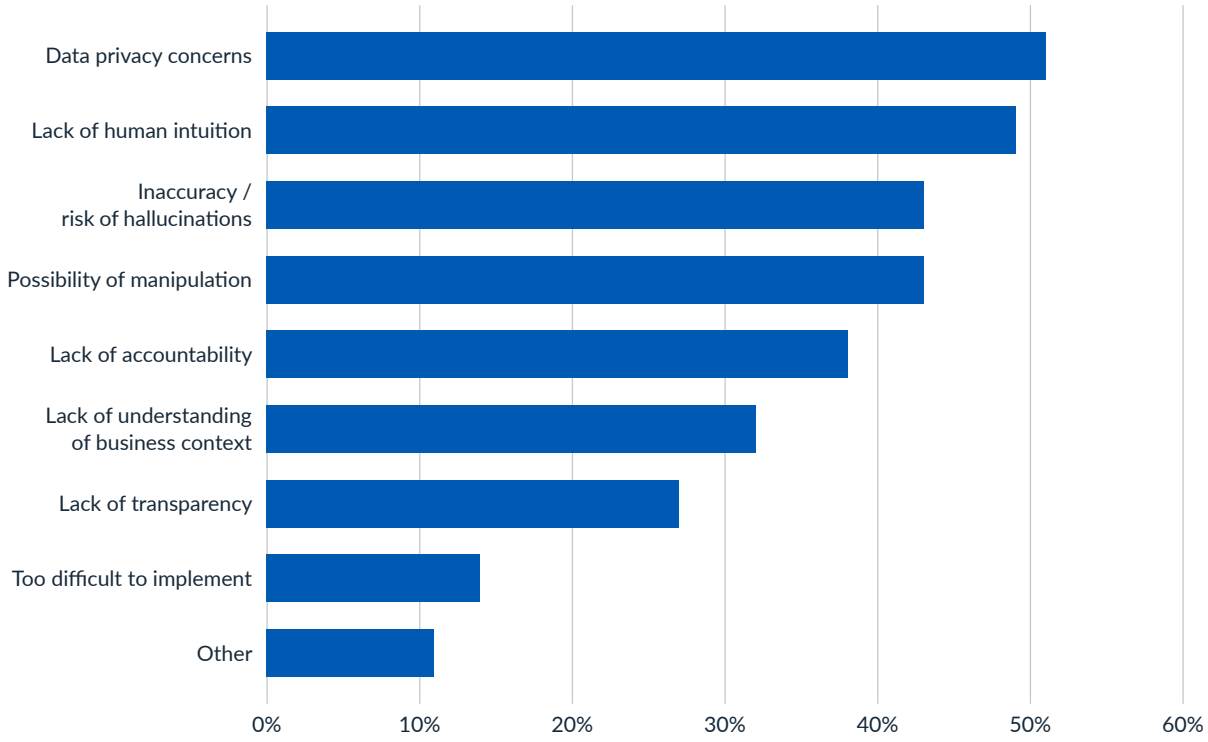

04

SECTION 4:

Regional Trends

EMEA

Digital sovereignty is a critical component of IT and security strategies

Across EMEA, geopolitics has moved from background noise to a board-level input in how organizations build and buy their cybersecurity. 70% of EMEA organizations report that their cybersecurity strategies have been directly shaped by geopolitical developments, and that influence is now visible at the point of purchase. Half are actively evaluating vendor and regional risk, scrutinizing where a provider is headquartered, where their data is stored, and which governments could compel access to it. Digital sovereignty sits at the center of this shift: approximately half of EMEA respondents rate it as "extremely important" to their cybersecurity strategy, and over a third more rate it as "important," making sovereignty a near-universal consideration, rather than a niche concern.

- The U.S. Cloud Act and the Microsoft testimony in France have caused a lot of governments and organizations to carefully consider from whom and where they buy their solutions
- Without digital sovereignty, European companies may be forced to pick which regulation they want to violate: GDPR or ones from the Cloud Act

Importance of digital sovereignty

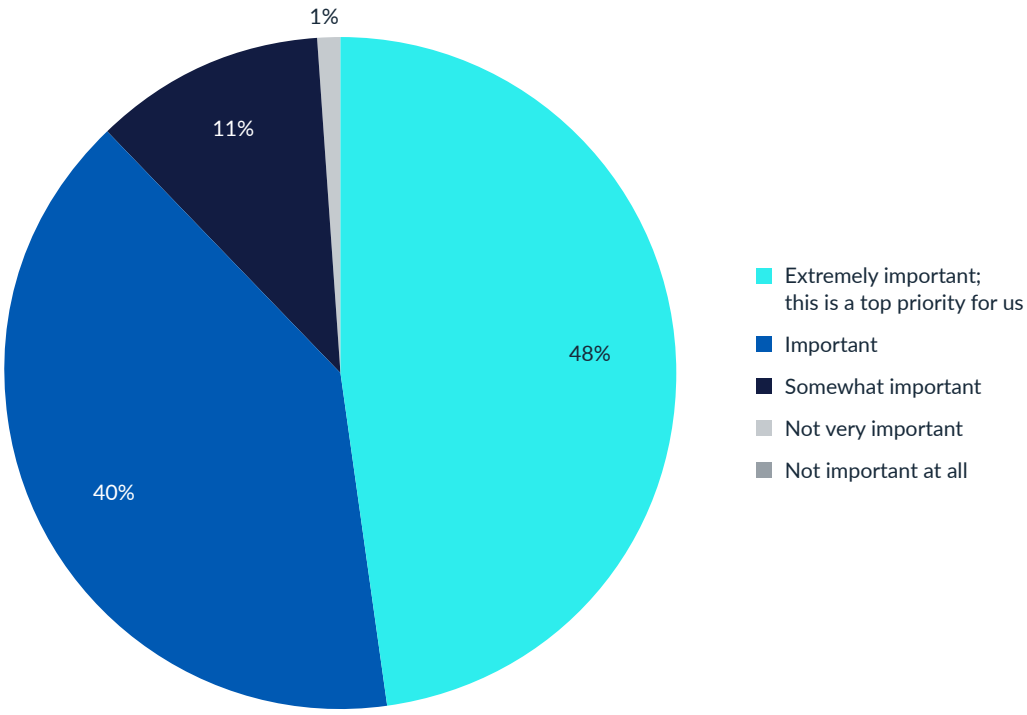

EMEA leaders now rank the U.S. among the biggest nation-state threats to their business

In a notable reflection of the current geopolitical climate, EMEA respondents identified the United States as the second nation posing the greatest cybersecurity threat to their business, behind China and ahead of Russia. Seeing a longstanding ally ranked alongside the adversaries that have traditionally topped these lists underscores how quickly trust has been recalibrated, and helps explain why so many European organizations are reassessing where their security tools, data, and dependencies ultimately reside.

In the DACH region, leaders are nearly twice as concerned about AI as they are about malware

Respondents in Germany, Austria, and Switzerland are close to twice as concerned about the risks of AI as they are about malware itself, which is a striking inversion of traditional threat priorities. That concern is shaping requirements rather than slowing adoption: In DACH, AI must be controllable and regional, large language models must be explicitly permitted, and humans must remain in the loop. The same instinct extends to how organizations buy, with a regional SOC and digital sovereignty weighing heavily on purchase decisions.

For nearly a third of breached DACH organizations, attacks disrupted the business for weeks

When defenses fail in the DACH region, the operational toll is significant. For nearly one-third of companies that suffered a breach, the cyberattack disrupted business continuity for several weeks. This is a reminder that, even as AI dominates the conversation, the real-world cost of a successful attack continues to be measured in lost time and stalled operations.

APJ

The shifting geopolitical landscape is heavily influencing cybersecurity purchase decisions

Across the Asia-Pacific and Japan (APJ) region, the same forces reshaping European strategies are now influencing how organizations evaluate and purchase cybersecurity. Concerns about where solutions are built, where data resides, and which jurisdictions can reach it have pushed sovereignty and supply-chain trust up the agenda. As in EMEA, these concerns are prompting leaders to look more carefully at the provenance of the tools they rely on.

- Similar to EMEA, the U.S. Cloud Act and the Microsoft testimony in France have caused a lot of governments and organizations to carefully consider from whom and where they buy their solutions
- Trusted alliances have been thrown into disarray, and in some cases longstanding allies have been recast as geopolitical adversaries

Impact of geopolitical risk on purchasing decisions

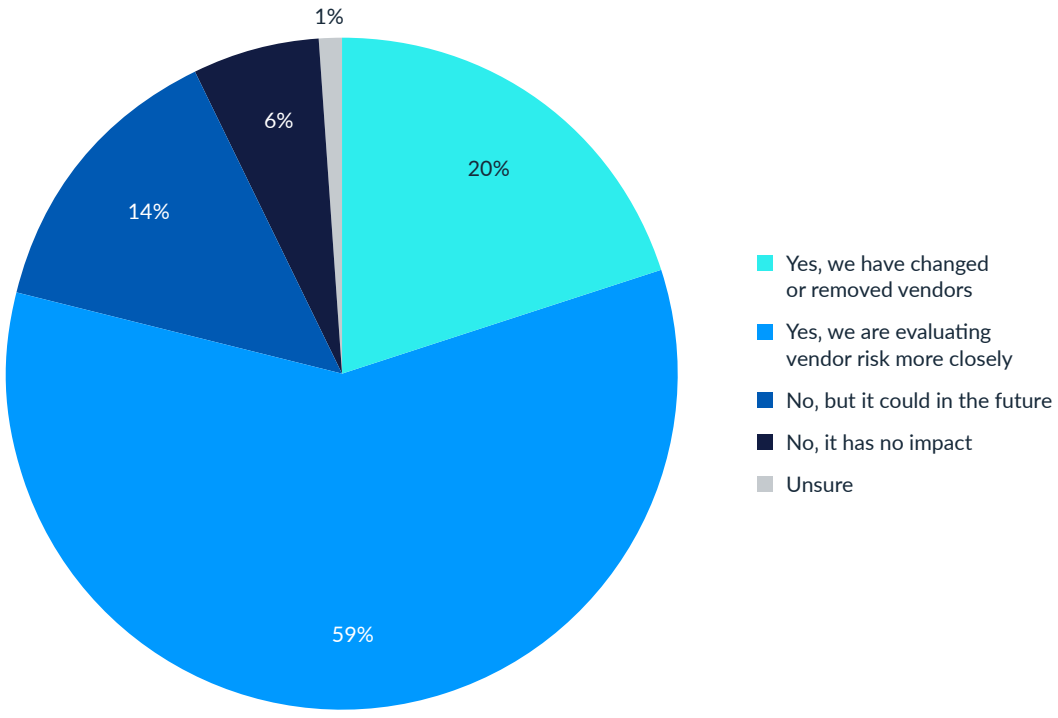

Singapore pairs the world's highest trust in autonomous AI with its sharpest concerns about it

For organizations in Singapore, AI is the headline issue. AI, large language models, agentic AI models, and their privacy implications rank as the single leading cybersecurity concern. AI is also the primary force shaping strategy: 53% cite secure AI adoption and data transformation as their top cybersecurity priority, compared with 41% globally.

That enthusiasm extends to autonomous operation. Sixty percent of Singapore organizations trust agentic AI to independently triage alerts and generate reports, and 56% trust it to patch vulnerabilities without human oversight, which is well above the 42% global average. Yet high adoption has not meant high protection: 61% experienced a major cyber incident, and 13% report severe disruption lasting six to nine months, more than three times the global average. With skills in short supply, 52% are turning to managed services to close the gap, significantly above the 39% global figure.

Japan confronts a disproportionate ransomware and detection challenge, and answers with transparency

Ransomware weighs more heavily in Japan than in most of the world: 29% name it as their dominant threat vector, 12 points above the global average. Visibility is a related worry: Japanese leaders are 12 points more likely than their global peers to believe an incident may have occurred but gone undetected, pointing to gaps in detection. And when incidents do land, they last: 33% report disruption of more than a month, 11 points above the global average.

Japan's response stands out for its openness. Organizations there are 13 points more likely to disclose incidents to the broader security community, reflecting a collaborative, information-sharing culture. On readiness, an incident response (IR) maturity gap is closing: current IR retainer adoption trails the global average (65% versus 74%), but future intent runs 7 points higher — a clear signal of growing recognition and market opportunity. Japanese organizations are also 7 points more likely to consider AI core to security operations, even as they remain more skeptical than their peers that AI will reduce burnout and turnover or improve detection of new threats.

Among ANZ's small and midsize businesses, AI adoption is near-universal, but trust hasn't caught up

AI has swept through Australia and New Zealand's small and medium-sized businesses (SMBs). 95% already use AI or LLMs, 98% have adopted or plan to adopt AI, and 69% say AI will drive their cybersecurity strategies over the next 12 months. Confidence in its defensive value is high; 80% believe AI will improve detection of new or sophisticated threats, though expectations are measured, with 60% not expecting it to reduce alert volume.

Even so, a trust gap is emerging alongside the enthusiasm: 31% of SMEs cite AI and agentic models as a top concern, driven by worries about data privacy, the absence of human judgment, and accuracy. The stakes are real: More than 70% of ANZ SMEs have experienced a significant cybersecurity incident, ranking the region second globally, with many losing up to three weeks of productivity. (Globally, more than half of affected organizations went on to pay a ransom.)

North America

North American organizations are much more likely to pay ransoms

Across North America, ransomware responses are shaped less by principle and more by pragmatism. 74% of organizations that experienced a ransomware attack report that a payment was ultimately made either directly (38%) or through an insurer or third party (35%), giving the region the highest payment rate in the study.

This willingness to pay reflects the operational realities organizations face once an attack is underway. While ransomware represents 13% of the most impactful incidents reported, the downstream effects of disruption, including extended downtime, create strong incentives to resolve incidents quickly.

Ransomware attacks

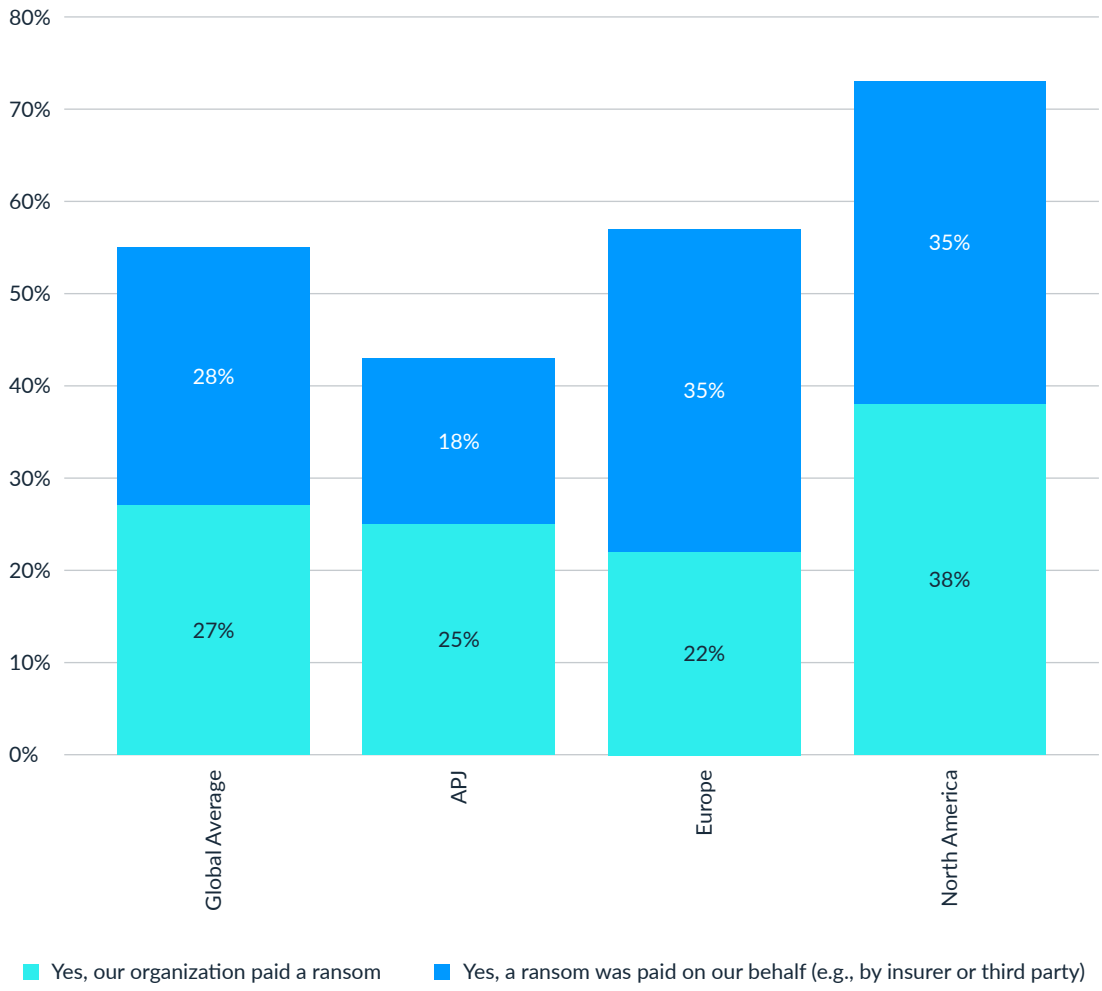

Ransomware response in North America is highly operationalized

Rather than treating ransomware as an exceptional event, North American organizations have embedded it into their broader operating model. 80% maintain an incident response (IR) retainer, and more than half have used those services within the past year, including 68% for active incident support and 58% for non-incident scenarios such as readiness and planning.

This level of maturity extends to disclosure and coordination. 94% of organizations disclose incidents, most commonly due to legal (51%) or insurance-driven (34%) requirements rather than discretionary transparency. Together, these patterns point to a region that is not immune to cyber risk (64% report a significant cybersecurity incident in the past 12 months), but is increasingly structured and repeatable in how it responds.

Most North American incidents are contained in weeks, but disruption is still the norm

While North American organizations are relatively fast to respond, speed does not eliminate impact. Only 13% report no loss of productivity following an incident, while the majority experience measurable disruption, including 32% resolved in less than a week and 28% lasting two to three weeks.

This highlights a key tension in the region: organizations are getting faster at containing threats, but not necessarily better at preventing operational impact. Even incidents resolved within days still introduce downtime, cost, and business disruption. This reinforces why rapid containment, external support, and in many cases ransom payment remain central to the response strategy.

North American leaders show high confidence despite persistent incident rates

Despite the frequency and impact of attacks, North American leaders remain confident in their ability to manage the threat landscape. While 69% of organizations report experiencing a significant incident in the past year, a strong majority say they are confident their teams can keep up with the volume and complexity of threats.

That confidence is underpinned by continued investment in operational capability. Organizations are scaling AI adoption across the SOC, with roughly 69% either piloting or deploying AI within security workflows and 95% saying AI will influence cybersecurity purchasing decisions. But the underlying data suggests these gains have not yet translated into fewer incidents or less disruption.

The result is a region where cybersecurity is both highly active and highly resourced, but still largely reactive. Closing that gap will require moving beyond incremental improvements, combining human expertise with systems that operate at machine speed to not only respond faster, but reduce the likelihood and impact of incidents altogether.

Conclusion

As we've seen:

Disruptive cyber attacks remain all too common

Overconfidence in an organization's ability to keep up with the volume and complexity of modern cyber threats is a very real risk

For most organizations, AI-powered security outcomes remain scarce

The promise of modern security technology is clear. AI, automation, and advanced analytics should help teams move faster, reduce manual work, and stay ahead of adversaries.

With enough tools, integrations, and skilled personnel, the thinking goes, any organization should be able to assemble a modern security operations center. And yet, many organizations are managing an ever-growing security stack, struggling to hire and retain experienced analysts, and fighting a constant backlog of alerts and investigations.

As is often the case, the do-it-yourself approach creates enormous operational burden that never seems to make it into a vendor's keynotes or marketing material.

AI will continue to reshape the cybersecurity landscape, but the challenge here is not reinventing cybersecurity; rather, it is evolving and adapting by reinforcing and extending existing security frameworks and controls, enabling security teams to manage AI-driven risks without overhauling their entire security infrastructure.

We hope that this report has shown you that you're not alone; both in experiencing the stresses of being a security leader and in the ongoing battle against those with malicious intent.

But our larger hope is that the results of this survey and the accompanying analysis and commentary will help you, your team, and your security vendors work together to build a stronger security posture for the rest of 2026 and beyond.

How Arctic Wolf® Can Help

To withstand today's AI-powered threats, organizations need superintelligent outcomes in security operations: faster detection, more accurate investigations, and more decisive response across the entire attack surface.

Achieving this requires trusted AI, real-world expertise, and rigorous validation of all agentic work.

Unfortunately, too many approaches in the market still expect customers to assemble the final mile, whether that's stitching together tools, operationalizing data, or building and governing AI agents themselves.

Our view is simple. If the customer is still doing that work, it's not a finished solution; it's a toolkit.

The Future of AI-Driven Cybersecurity

Arctic Wolf's Aurora® Superintelligence Platform brings together three major AI advancements, grounded in more than 14 years of real-world security operations expertise.

1. **Swarm of Experts™** is an agentic framework that is AI-led with humans in the loop. It is designed to plan and execute actions for most cybersecurity goals, rather than being narrowly focused on individual tasks like triage or investigation.
2. **Security Operations Graph™** is the proprietary data and intelligence foundation for the platform, ingesting more than 10 trillion telemetry events in real time each week and integrating them through an open data pipeline. As a result, it is able to generalize insights across diverse signals without exposing customer-specific data. Critically, the graph also incorporates business context and case memory for each customer environment. This model is part of the Concierge Experience™, ensuring every decision is optimized for each organization's unique context.
3. **AI Trust Engine™** is a process of validation and a series of guardrails that bound the autonomy of the agents to help ensure accurate, reliable, and trustworthy outcomes. Agents in the swarm are built to be deterministic, which guards against guessing or attempting to provide any answer that is not part of their known, validated experiences. Humans remain in the loop to handle tasks beyond an agent's experience, and each validated resolution is fed back into the system, expanding the agent's capabilities and enabling similar tasks to be handled automatically in the future. All decisions, whether agentic or human, are reviewed by an AI judge as an added layer of protection. Before entering the swarm, every new agent must be battle-tested in Arctic Wolf's SOC, outperform human-only workflows, and exceed rigorous performance benchmarks.

Together, these three innovations enable the platform to achieve superintelligence for security operations, helping to deliver outcomes far beyond what human-only or AI-only solutions can provide.

AW_R_2026 TRENDS_0726