prophet

State of AI in Security Operations 2025

Table of Contents

Introduction 3

The Alert Problem 6

The Pain in Organizations 12

The AI SOC Shift 16

Methodology & Demographics 26

Appendix 29

Glossary 32

Introduction

Foreword

In the relentless landscape of modern cybersecurity, the adage "knowledge is power" has never been more critical. Security Operations Centers (SOCs) stand as the frontline defense for organizations worldwide, yet they are increasingly besieged by an escalating "alert problem" and the sheer volume of data generated by a proliferating array of security tools. Traditional approaches which were once sufficient are now strained to their breaking point, leaving organizations vulnerable to sophisticated and rapidly evolving threats. To make matters worse, there is a growing talent shortage and security leaders are being asked to do more with the same.

At Prophet Security, we recognized this growing crisis. Our journey began with a clear vision: to fundamentally transform how security teams operate, moving beyond reactive measures to proactive, intelligent defense. This 'State of AI in SOC 2025' report, based on insights from nearly 300 security leaders, illuminates the urgent need for a paradigm shift. It reveals that an alarming 40% of alerts are routinely ignored, while a staggering 60% of security teams have experienced critical breaches stemming from overlooked alerts. The average time to investigate an alert remains too slow to counter today's agile adversaries. These findings underscore a critical reality: the human-centric SOC model, without powerful augmentation, is simply unsustainable against the modern threat landscape.

This report is more than just a collection of statistics; it is a call to action. It highlights the rapid ascent of AI for Security as a top-three priority for security leaders. Companies are taking note, with nearly 35% of non-AI SOC users implementing an AI SOC solution and ~60% planning to evaluate AI SOC solutions within the next year. This is not merely a trend, but a strategic imperative. AI-native platforms are emerging as the indispensable solution to manage data overload, accelerate triage, and empower analysts to focus on what truly matters. We believe that by embracing AI, organizations can finally move from merely reacting to threats to anticipating and neutralizing them with unprecedented efficiency and efficacy.

We hope through these pages you gain insights necessary to navigate this transformative era in cybersecurity. The future of the SOC is AI-driven, and the time to adapt is now.

Sincerely,

Kamal Shah
Founder & CEO

Key Findings

These findings come from a survey of 282 security leaders (CISOs, Security Directors, Managers, and Analysts) from companies over 1,000 employees primarily in the United States. See detailed demographics on Page 28.

~960

alerts generated daily

Companies of all sizes are experiencing a barrage of alerts. Smaller enterprises and middle-market companies generate ~500 alerts per day causing a substantial management burden while larger companies are facing a tsunami of data from the ~3,000 alerts per day.

A concerning 40% of alerts (on average) are never investigated, leaving organizations vulnerable to significant, avoidable security risks. Roughly 60% of security teams have reported that an ignored alert proved to be critical, leading to a direct impact on the organization's security posture.

~40%

of alerts are never investigated

57%

of companies suppress detection rules

To cope with the alert problem, approximately 57% of companies are actively suppressing detection rules, especially for cloud and identity. Companies are consciously accepting increased risk to manage their current operational limitations, making them more susceptible to sophisticated attacks.

"AI for security" has become a top-three priority for security leaders, following data security and cloud security. The rapid ascent of AI to a top-tier priority indicates a widespread recognition among security leaders that current approaches are unsustainable and that AI is a viable path forward to address the "alert problem" and improve efficiency.

AI for Security

is a top 3 priority

55%

of companies use AI for alert triage & investigation

Most companies use AI in some capacity for triage and investigation, with more starting to use AI-native pure-play solutions. Of the AI non-users, nearly 60% plan to evaluate an AI SOC solution within the next year, and another 30% are evaluating AI SOC solutions. This strong intention to adopt AI signals a clear mandate for the security industry.

Security leaders anticipate AI solutions will handle approximately 60% of SOC workloads within the next 3 years, fundamentally reshaping security operations. It implies a dramatic reduction in manual effort, allowing human analysts to focus on more complex, strategic tasks while AI handles routine, high-volume operations.

~60%

of SOC workloads will be completed by AI in next 3 years

The Alert Problem

A tax on SOC resources

Organizations have 17 tools on average, while larger companies have over 20 alert generating tools

How many security tools do you have deployed today that generate alerts?

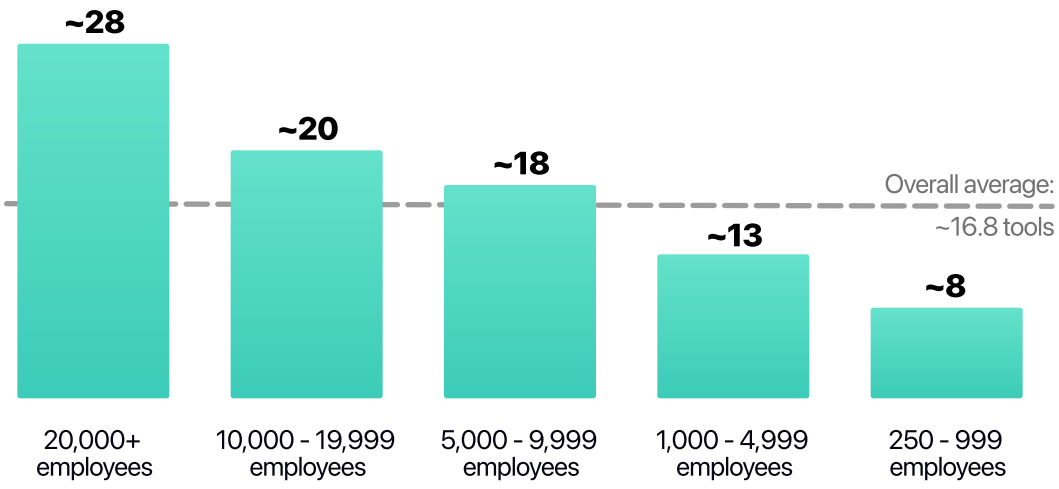

The number of security tools generating alerts deployed by companies directly correlates with their size, with larger enterprises utilizing a significantly higher average number of tools. Companies with over 20,000 employees report an average of nearly 30 security tools that generate alerts, while those with 250-999 employees use fewer than 10. This indicates a growing complexity in the security landscape, particularly for larger organizations, where a multitude of tools generate an overwhelming volume of alerts, making it increasingly challenging for security teams to effectively manage and respond.

Industry commentary:

We have way too many data sources that generate alerts and information. Any kind of incident takes a really long time to analyze and correlate.

Security Lead, Financial Services Co.

Companies generate ~960 alerts daily; Very large enterprises reach over 3,000

How many alerts do you receive per day (of all severities)? Please select the choice that most closely matches your average

Daily average number of alerts received by companies

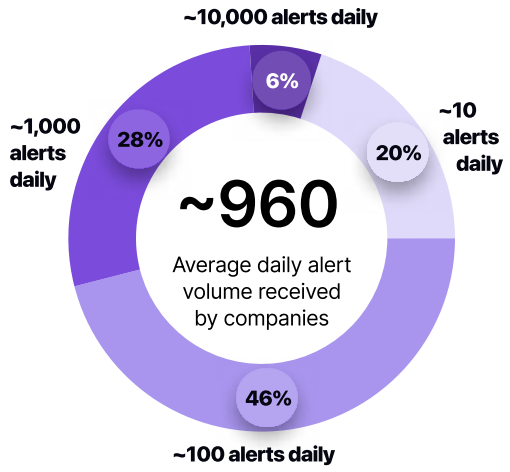

Daily average number of alerts by different sized companies

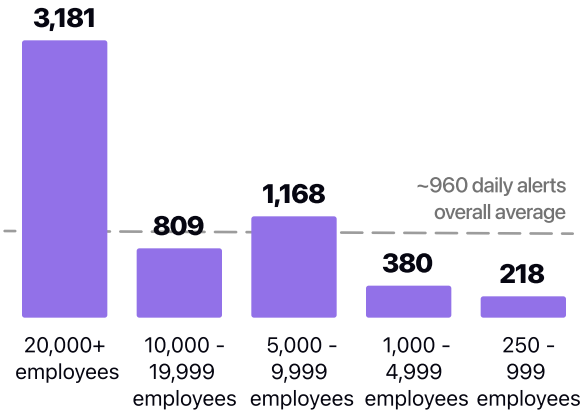

Similarly, the number of security alerts an organization generates directly correlates with their size, with larger organizations using a significantly higher average number of tools and therefore generating more alerts.

Very large enterprises (20,000+ employees) generate an average of ~3,000 alerts per day, while most large enterprises (10,000-19,999 employees and 5,000-9,999 employees) generate around 1,000 alerts daily.

Smaller enterprises and middle-market companies tend to see fewer alerts for three reasons: they have fewer alert generating tools in their security stack, they have fewer users and they tend to be a lesser priority for cybercriminals.

This already uphill battle against alert volumes only keeps getting steeper. As found across annual SANS SOC 2021-2025 surveys, alert volumes continue to grow year over year. While overall volumes inflate, the number of real threats doesn't necessarily follow, leading to a large haystack for SOC teams to search through for threats.

Industry commentary:

Security personnel cite being woken up during the middle of the night for an alert only for it to be a false positive.

This doesn't just happen during the night, the false positives keep pestering analysts 24/7.

SOC Manager, Financial Services Co.

Organizations are taking risks by not investigating ~40% of alerts generated

What percentage of alerts of all severities (low, medium, high critical) are being ignored versus triaged or investigated?

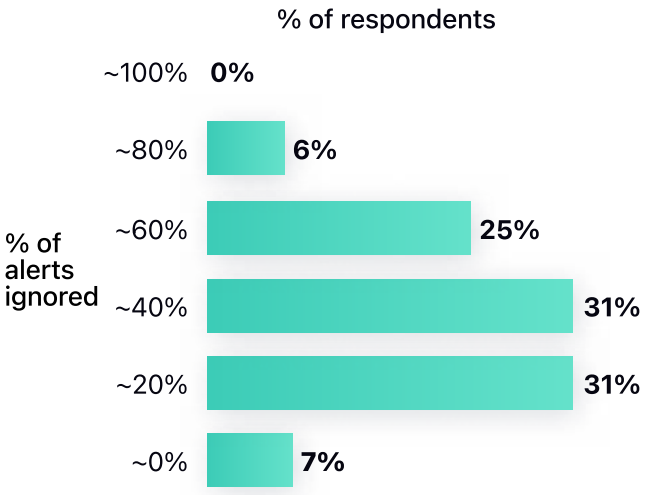

1/3rd of companies ignore more than **half** of their alerts

Industry commentary:

In regulated industries, like healthcare and financial services, you aren't allowed to turn off certain alerts due to compliance rules - even if those alerts tend to be false positives.

As more tools are brought on, it becomes an even bigger haystack to search through for real threats.

CISO, Financial Services Co.

The sheer volume of alerts stemming from an ever-growing array of security tools has created a critical vulnerability: the alarm is ringing, but no one is answering. The feeling of constantly "drowning in noise" fosters an environment where an average of ~40% of all alerts are simply ignored or not investigated. This leads to a scenario where analysts become tired of false alarms and may overlook a legitimate, potentially catastrophic breach.

The percentage of alerts that are ignored tends to increase the smaller the company size is, indicating a potential struggle for smaller SOC teams to keep pace with alert volumes.

	20,000+ employees	10,000 - 19,999 employees	5,000 - 9,999 employees	1,000 - 4,999 employees	250 - 999 employees
Avg. % of alerts ignored	26%	36%	37%	41%	52%

This data highlights that while larger organizations investigate a higher proportion of alerts, companies of all sizes ignore alerts. Even the respondents that claim 0% of alerts are ignored, tend to rely on an MSSP or MDR which likely ignore some alerts. This reinforces the the reality that a significant percentage of alerts are left untriaged or ignored across the board.

This relentless influx of alerts, often exacerbated by a high percentage of false positives and a lack of context, leads to a phenomenon known as "alert fatigue." Security analysts become desensitized and overwhelmed, making it difficult to distinguish genuine threats from background noise. This directly impacts productivity and morale, as analysts spend disproportionate time on low-priority or redundant notifications. The emotional toll is significant, with many reporting stress, burnout, and a reduced ability to relax outside of work. This fatigue translates into critical risks: slower response times, increased mean time to respond (MTTR), and, most dangerously, the potential for truly critical alerts to be missed or ignored.

This isn't just an inefficiency; it represents a profound and unnecessary risk exposure for organizations.

~60% of security teams have ignored an alert that turned out to be critical

The consequences are dire. A staggering 61% of security teams admit to having ignored an alert that subsequently proved to be critical, leading to outcomes as severe as customer data exposure, system downtime, or significant operational interference. This "alert fatigue" is not merely an inconvenience; it is a direct pathway to costly breaches and reputational damage. It highlights a fundamental breakdown in the human-centric SOC model when faced with an unmanageable data deluge.

When segmented by company size, the middle-market companies had the highest rates of ignoring an alert that proved to be critical, at 77%, followed by Enterprises and Large Enterprises at 68% and 64% respectively.

Ignoring alerts directly increases the risk of successful cyberattacks. When critical alerts are missed, it can lead to undetected intrusions, longer breach lifecycles, and higher overall costs. Furthermore, failing to respond to critical issues in a timely manner can trigger significant regulatory penalties and legal liabilities.

Have you/your team ever ignored an alert that turned out to be critical?

Note: critical means it risked exposing customer data, led to a system downtime, or interfered with regular operations

% of respondents

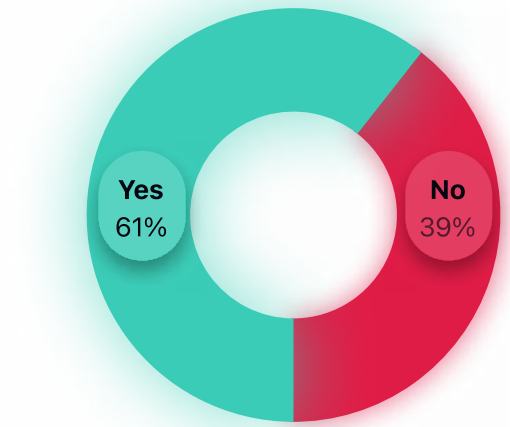

Industry commentary:

Unfortunately, we have to tune out or ignore alerts due to the high volume, and so far we have accepted the risks that comes with that.

Thankfully we haven't been burned yet...but, this is not a game of chance worth playing.

Director of Security, Technology Co.

On average, alert dwell time is ~56 min while MTTI is ~70 min…too slow to stop attackers

On average, how long does it take to investigate an alert (MTTI) and the dwell time?

~56 Minutes

Dwell Time

The period from when an alert is fired from a detection tool to when it is picked up for review

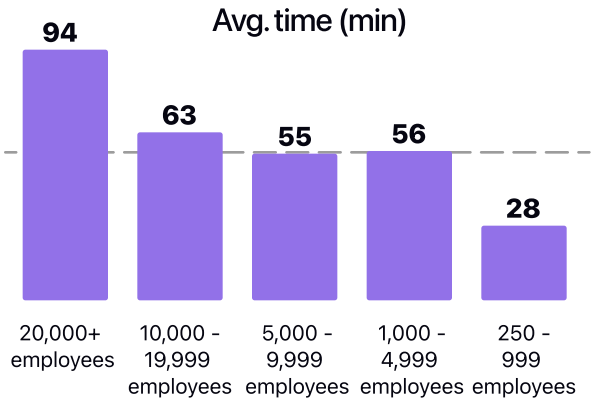

~70 Minutes

Mean Time to Investigate

The Average time it takes for a security team to thoroughly investigate a generated alert

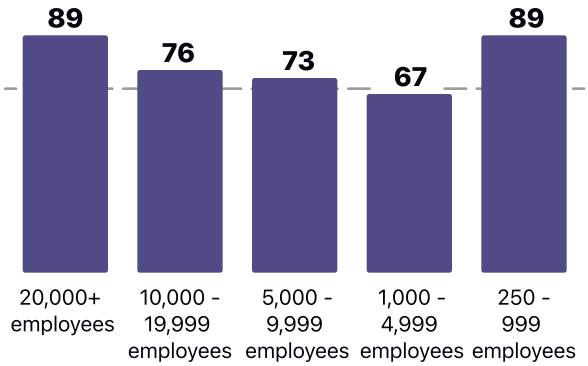

Response times are **"too slow to stop attackers."** Every minute an alert remains undetected or unaddressed increases the likelihood of a successful breach, data exfiltration, or compromise. Attackers can move laterally, escalate privileges, and achieve their objectives faster than these average investigation and dwell times allow for detection and containment. **As found in Crowdstrike's 2025 Global Threat Report, phishing threats take only 48 minutes on average to extract sensitive information (or as little as 51 seconds in fastest recorded time).**

On average, the alert dwell time is approximately 56 minutes, which refers to the time period between an alert being fired and it being acknowledged by a team. Note how this average is longer than the average time it takes for phishing threats to materialize into breaches. The Mean Time To Investigate (MTTI) an alert is about 70 minutes and represents the time it takes for a security team to thoroughly investigate a generated alert. Now imagine spending your entire day in 70 min chunks investigating alerts. How wouldn't a human get fatigued?

These metrics underscore an urgent need for security operations to accelerate their alert triage and investigation processes. Reducing dwell time and MTTI is paramount for effective risk reduction, necessitating solutions that can rapidly identify, contextualize, and prioritize genuine threats amidst the overwhelming volume of alerts.

The Pain in Organizations

Analyst burnout, missed detections, poor automation

Alerts are being ignored because of the constrained SOC resources; often presented as interrelated issues

What are the top 3 challenges that your SOC team currently faces? (Select up to 3 choices)

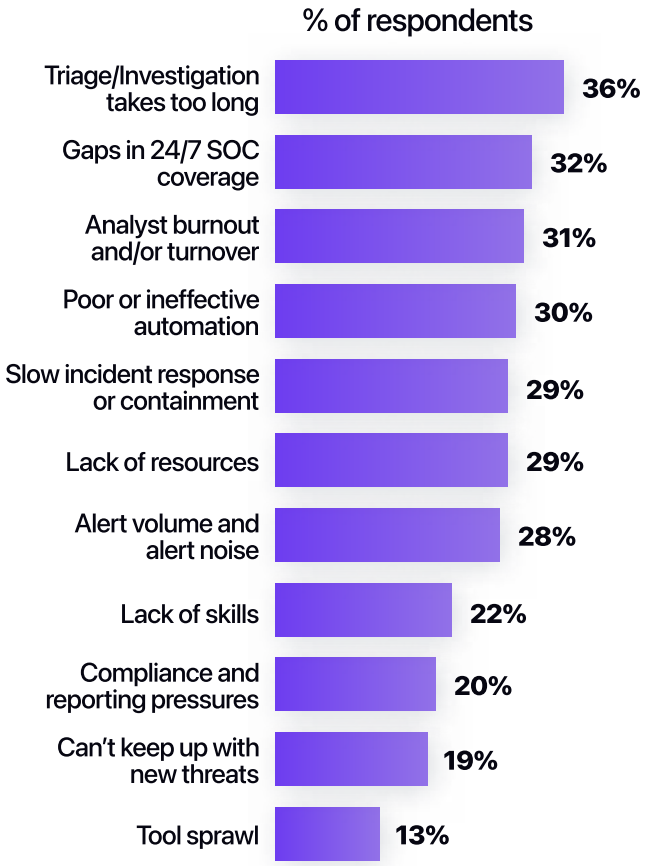

Industry commentary:

Even F1000 companies struggle with the alert volume coming in. It causes challenges with alert fatigue and repetitive investigations.

Head of Security, Retail Co

We want to investigate alerts that were previously ignored due to time constraints.

SOC Manager, Technology Co

SOCs are experiencing a vicious cycle: high alert volumes, coupled with inefficient processes and a lack of automation, lead to prolonged investigation times. In turn, this contributes to analyst burnout and turnover, further exacerbating staffing and coverage gaps. The result is an environment where alerts are by necessity or oversight ignored.

Triage/investigation takes too long (36%): This is the leading challenge, indicating that once an alert is generated, the process of understanding and responding to it is excessively time-consuming. This directly contributes to alerts being ignored, as teams likely can't keep up with the volume.

Gaps in 24/7 SOC coverage (32%): A significant portion of respondents face issues with continuous monitoring, which means that during off-hours, weekends, or holidays, alerts might not be seen or addressed promptly, leading to them being overlooked.

Analyst burnout and/or turnover (31%): This human element is crucial. When analysts are overwhelmed by the volume and complexity of alerts, it leads to burnout, which in turn causes turnover. A constantly rotating or fatigued team will struggle to maintain vigilance, preserve institutional knowledge, and effectively respond to alerts.

These well known challenges all point to the current constraints of SOC teams: there aren't enough SOC analysts and budgets are often strained, especially since cybersecurity is viewed as a non-revenue generating activity. Something has to give...

To cope with the noise, 57% of organizations deliberately suppress detection rules, accepting higher risk to just stay afloat

Q Are you currently limiting the number of detection rules in place due to limited capacity to triage and investigate alerts?

% of respondents

For organizations, "suppressing alerts" means limiting the number of detection rules they have in place. This action is taken specifically to cope with their limited capacity to triage and investigate the overwhelming volume of alerts that their security tools generate. Essentially, rather than having every potential threat or anomaly trigger an alert, organizations are intentionally disabling or not activating certain detection rules because their human security teams lack the resources (time and staff) to investigate all of them effectively.

While it might alleviate immediate alert fatigue, it can give a misleading impression of security posture by not revealing the full scope of potential issues. In essence, suppressing detection rules is a coping mechanism driven by overwhelming alert volumes and strained SOC resources, but it comes at the high cost of reduced visibility and increased susceptibility to cyberattacks.

Industry commentary:

Detection rule suppression tends to be high because of bad detection engineering processes. Lots of detection rules tend to trigger false alerts, and because of the weak internal processes, organizations tend to turn off the rule rather than go through the hassle of tuning it.

Staff Security Engineer, Technology Co

Top two detection gaps companies have due to resource constraints are Cloud and Identity

Survey data shows that if given more triage and investigation resources, security teams would most urgently enable detections for Cloud (65%) and Identity (61%) - well above other areas like Insider Threat (47%) and DLP (29%). This prioritization underscores a critical and growing visibility gap in two of the most dynamic attack surfaces in today's enterprise environments.

The decision by 57% of organizations to mute detection rules is not just tactical triage, it reflects structural overload within modern SOCs. By suppressing alerts from cloud and identity sources, teams are opting for short-term survivability at the cost of long-term visibility debt. Every disabled rule represents a blind spot adversaries can exploit, especially as cloud misconfigurations and credential abuse often evade notice until post-incident. Suppression essentially becomes outsourced risk - a gamble against attackers who iterate faster than most governance can respond.

The economic impact is also substantial: organizations continue investing in top-tier detection tools but run them at reduced capacity, incurring hidden costs in underused licenses and degraded ROI.

Q Given more triage and investigation resources, which detections would you enable? (Select your top 3)

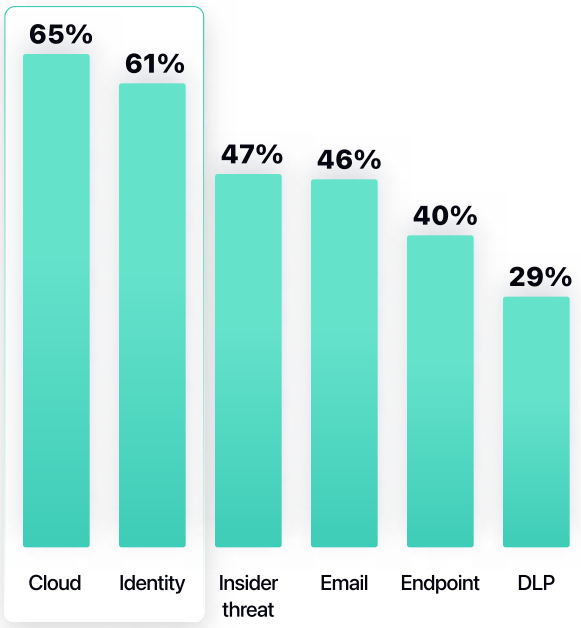

Based on SACR research, the remedy is not simply "more analysts," but smarter signal engineering and automation. AI-powered correlation, enrichment, and scoring can reduce alert volumes dramatically without losing fidelity. Equally vital is a cloud- and identity-focused detection engineering program that treats muted rules as technical debt which is scheduled for review, reactivation, and refinement. By measuring the delta between active and suppressed rules, both in coverage and potential risk, security leaders can turn today's blind spots into tomorrow's strategic advantage.

The AI SOC Shift

How security leaders
envision the future SOC

Though an emerging priority, AI for security already ranks among the top 3 priorities for security leaders

Please select the top security priorities for you? (Select up to 3 choices)

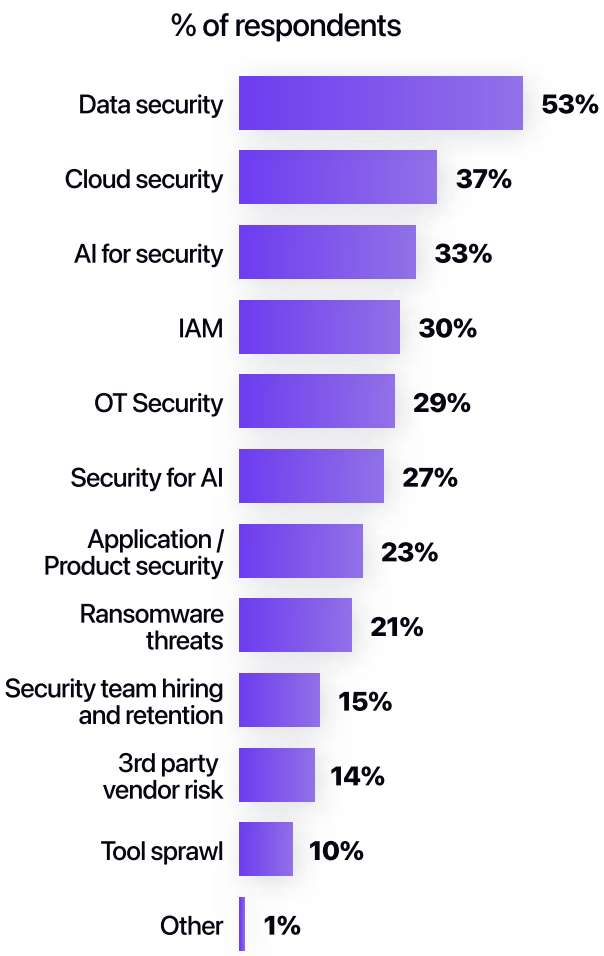

The transition to an AI-driven SOC is no longer a distant vision; it's a rapidly accelerating reality. While Data Security (~53% of respondents) and Cloud Security (~37% of respondents) remain top priorities for security leaders, reflecting long-standing industry concerns, AI for Security has rapidly emerged to claim a spot among the top three (~33% of respondents), signifying its critical importance in modern cybersecurity strategies.

AI for Security has firmly established itself as an emerging priority for security leaders, not just for the SOC but across all of security. This high ranking underscores a growing recognition among security professionals that traditional approaches are insufficient to combat the escalating complexities of the modern threat landscape and the overwhelming alert problem. The immediate prioritization of AI signals a clear market demand for innovative, intelligent solutions to augment and transform security operations.

The top use for AI in the SOC is for Alert triage and investigation; followed by Detection engineering & tuning, and Threat hunting

What use cases do you believe AI is most valuable for in your SOC? (Select the top 3)

1 Alert Triage and Investigation

67% of respondents

Alert Triage and Investigation stands out as the primary application for AI, cited by a significant 67% of respondents. This directly addresses the "alert problem" by automating the sifting through high volumes of alerts, prioritizing critical incidents, and providing context for faster human analysis.

2 Detection Engineering and Tuning

65% of respondents

Detection Engineering and Tuning is the third top use case, favored by 65% of respondents. This indicates a strong desire for AI to refine detection rules, reduce false positives, and ensure the efficacy of security controls, thereby directly mitigating the need to suppress detection rules due to capacity constraints.

3 Threat Hunting

64% of respondents

Threat Hunting follows closely, with 64% of respondents believing AI is most valuable for this proactive security activity. AI's ability to analyze vast datasets for subtle patterns and anomalies can empower analysts to uncover hidden threats before they escalate.

Interestingly, Remediation and Incident Containment was cited as a secondary priority, with only 43% of respondents selecting it. This suggests that while security leaders recognize AI's power in identification and analysis, there's a current tendency to view human intervention as crucial in response and containment phases. So rather than prioritizing use of AI when moving down the investigation lifecycle (towards remediation and incident containment), organizations are looking to apply it earlier (e.g., for detection engineering and tuning).

Security Report Generation was selected by about 45% of respondents for AI's value in the SOC. While AI can add great 'color' beyond typical dashboards, it's seen as having a lower direct impact on overall security posture.

Note: "Other" option not shown above as less than 1% of respondents selected it

~55% of organizations currently use AI in some capacity to triage, investigate, and/or remediate alerts

Q Do you currently use an AI solution to help triage, investigate, and/or remediate alerts?

% of respondents

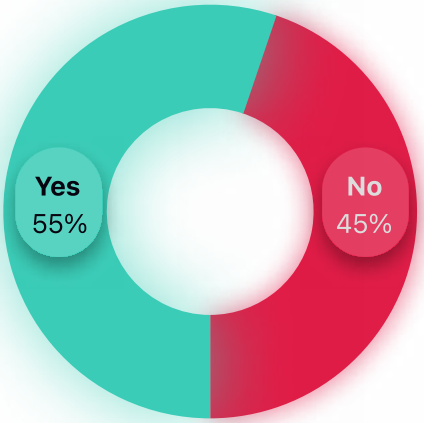

Approximately 55% of organizations now report using AI in some capacity to triage, investigate, or remediate alerts. But digging deeper reveals a crucial divide in how AI is being adopted - and how effective it is.

AI-native, pure-play SOC platforms are emerging as solutions for the value AI can provide in security operations. Purpose-built from the ground up by security professionals, these platforms embed reasoning, context, and transparency into every phase of alert triage, investigation, and response. They're engineered to scale with the realities of the modern SOC, and trained with the kind of operational nuance only SOC leaders, CISOs, and front-line analysts can provide. As these specialized platforms gain traction across more forward-leaning security teams, they're redefining what effective, AI-driven security operations truly look like.

By contrast, many organizations rely on bolt-on AI features from generalist platforms like Microsoft Sentinel or Crowdstrike. While these tools can offer surface-level automation, they often depend on broadly trained models not tailored for security operations. The result is AI that lacks investigative depth, produces more noise than clarity, and ultimately requires analysts to do the heavy lifting. These limitations highlight why generalist solutions often fail to deliver on the promise of faster, more efficient SOC performance.

Despite AI's momentum, 45% of organizations still haven't adopted any AI for alert triage or investigation, a gap that underscores the scale of untapped opportunity. According to the 2025 SANS SOC study, about 30% of security teams are using AI/ML within SOC operations, and another 40% are applying it in adjacent workflows. The data is clear: while the value of AI is widely acknowledged, the opportunity lies in adopting the *right* kind of AI solutions purpose-built to meet the evolving challenges of security operations.

Industry commentary:

Team members see AI SOC solutions as an additional member of the SOC. One that does great work, is easy to interact with, and closes alerts out.

AI SOC Customer

Note: User platform responses were grouped into three categories for anonymization: Security Generalists (broad platforms with wide-ranging cybersecurity capabilities), Security Specialists (vendors with deep focus on specific security areas), and Pure-play SOC Solutions (tools purpose-built to optimize SOC operations).

45% of respondents are non-AI SOC users but have appetite to adopt in the next year

Do you currently use an AI solution to help triage, investigate, and/or remediate alerts?

Do you plan on using/evaluating an AI SOC solution in the next 12 months?

% of respondents

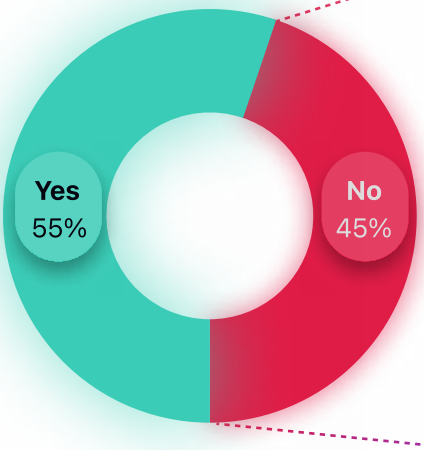

% of current non-AI SOC solution users
("No" respondents from pie chart)

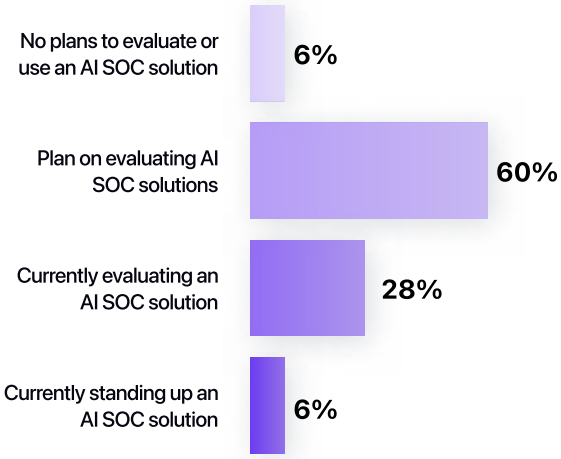

The detailed breakdown of organizations without a current AI solution demonstrates a strong forward momentum and a market clearly poised for rapid AI adoption. A significant majority, approximately 95%, are actively engaged in exploring AI: about 60% are planning on evaluating AI SOC solutions, while another 28% are already evaluating such solutions within the next 12 months. Furthermore, a notable 6% are currently in the process of standing up an AI SOC solution.

This overwhelmingly positive outlook means that organizations are largely past the "if" and are now focused squarely on the "how" and "when" to integrate AI into their security workflows. The critical implication is that only a very small minority, approximately 5%, have no plans whatsoever to evaluate or use an AI SOC solution. This collective energy signals a significant shift in investment and strategy across the industry, firmly positioning AI SOC solutions as an indispensable next step for organizations determined to overcome the alert problem, enhance efficiency, and elevate their overall security posture. The market is not merely interested; it's actively preparing to embrace AI to fundamentally transform its SOC capabilities.

Industry commentary:

To better support our security analysts, we're seeking a solution that can help mitigate the current investigative burden caused by thousands of alerts each month.

Director of Security, Financial Services Co
(Currently evaluating AI SOC Solutions)

The top metrics for tracking ROI of AI SOC tools are: MTTI, MTTR, and 24/7 coverage

The effectiveness of AI-driven SOC tools is most commonly measured by their impact on three critical security metrics, the same security metrics used in every SOC. This indicates a clear focus on efficiency and response capabilities, and in the case of AI SOC Solutions, how seamlessly it integrates with the existing workflow. Security leaders pinpoint these as the top indicators of ROI:

Q How do you measure the ROI or effectiveness of AI-driven SOC tools? (Select the top 3 metrics)

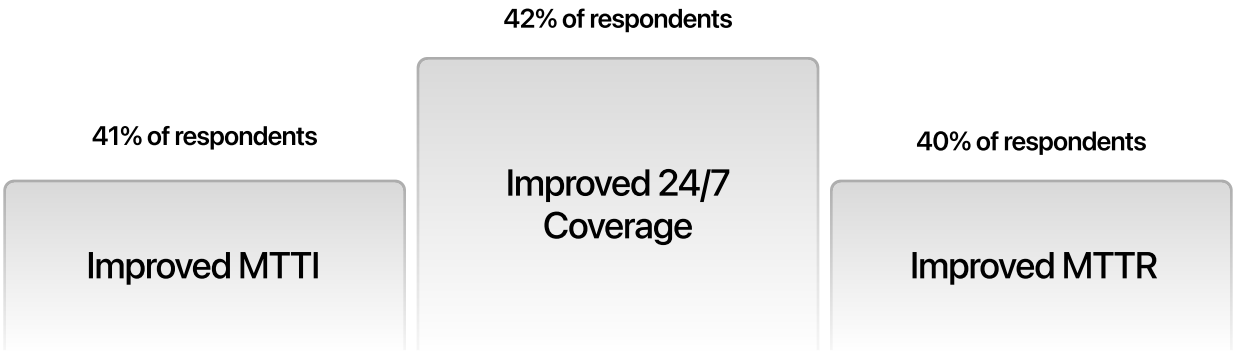

MTTI is the average time it takes for a security team to thoroughly investigate an alert, from its initial trigger to understanding its nature, scope, and potential impact.

This metric refers to the SOC's ability to maintain continuous, round-the-clock monitoring and response capabilities, ensuring no gaps in vigilance regardless of time of day or staffing levels.

MTTR measures the total time from the initial alert firing to its complete containment, eradication, and recovery. It encompasses the investigation phase (MTTI) as well as the response actions.

Note: Answer options not mentioned above include analyst satisfaction (35%), Reduced false positive rate (33%), reduced missed alert rate (30%), Increased time allocation to threat hunting (24%), Reduced spend on MSSPs/MDRs (18%).

Improved MTTI

In a traditional SOC:
This process is often highly manual, involving analysts sifting through disparate logs, correlating events, and performing lengthy analyses to distinguish true threats from false positives.

How AI SOC improves it:
AI accelerates this by automating initial alert triage, enriching alerts with crucial context from various sources (e.g., threat intelligence, asset data), correlating events across multiple security tools, and prioritizing the most critical threats.

Improved 24/7 Coverage

In a traditional SOC:
Achieving 24/7 coverage typically demands large, multi-shift security teams, which are expensive to maintain and prone to challenges like analyst burnout, turnover, and skill shortages, leading to potential coverage gaps.

How AI SOC improves it:
AI solutions can handle repetitive tasks, perform initial automated investigations, and alert human analysts to critical issues even outside of standard working hours. This significantly extends the operational reach of a SOC, allowing for continuous monitoring and faster initial response to threats, even with limited human resources, thereby providing true 24/7 vigilance.

Improved MTTR

In a traditional SOC:
After an alert is investigated, the response often involves manual execution of containment playbooks, isolation of affected systems, and remediation steps. This sequential, human-driven process can be slow, increasing the window of opportunity for attackers to cause damage.

How AI SOC improves it:
By drastically improving MTTI, AI inherently shortens MTTR. Furthermore, AI can suggest and even automate response actions, orchestrate containment measures across various security tools, and initiate remediation steps with minimal human intervention. This accelerates the entire incident lifecycle.

Industry Commentary:

Our MTTI for alerts has now dropped to just 3-4 minutes with Prophet Security. It used to be over 25 minutes.

Director of Security, Technology Co
(AI SOC Customer)

Prophet Security completed an investigation to a greater depth in 9 min, which took the SOC team 2.5-3hrs to complete.

CISO, Retail
(AI SOC Customer)

Data privacy and regulatory concerns are the leading barrier to AI for cybersecurity adoption

What is your biggest concern or barrier around using AI for cybersecurity? (Select all that apply)

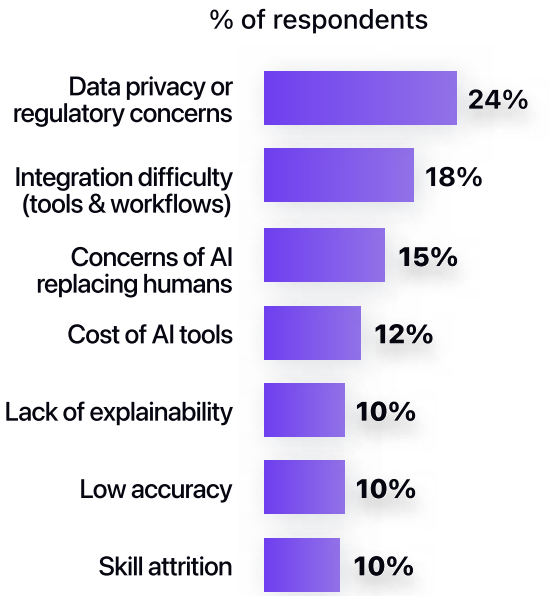

Even with the clear enthusiasm and recognized ROI of AI in the SOC, significant barriers to its adoption persist, primarily rooted in concerns around data and integration. The leading barrier, cited by ~24% of respondents, is data privacy or regulatory concerns. This highlights the complex legal and compliance landscape organizations must navigate when deploying AI solutions that handle sensitive security data.

Following closely, difficulty integrating AI with existing tools and workflows is a major concern for approximately 18% of respondents. This points to the challenge of seamlessly embedding AI into diverse and often fragmented security ecosystems, ensuring interoperability and minimal disruption to current operations.

A surprising finding is how few respondents selected “Low accuracy” (10%). This may imply that security leaders are becoming more comfortable with the capabilities of AI solutions.

Other notable barriers include concerns around AI replacing humans or the risk of team rejection (15%), the cost of AI tools or platforms (12%), and the lack of explainability or transparency - aka AI as a "black box" (10%). These concerns collectively underscore that while the value of AI is increasingly clear, addressing practical implementation challenges and regulatory hurdles is crucial for widespread adoption.

Cybersecurity leaders believe that ~60% of the SOC workload will be performed by AI solutions in the next 3 years

Q What percentage of you SOC workload do you believe will be performed by AI in the next 3 years?

83%

of security leaders believe that

over half

of SOC workload will be completed by AI in the next 3 years

Despite the identified barriers to adopting AI, such as data privacy concerns and integration difficulties, these challenges are likely to be overcome due to the overwhelming conviction among cybersecurity leaders regarding AI's future role in the SOC. Cybersecurity leaders broadly believe that a substantial portion of their SOC workload will be performed by AI solutions in the next three years. Specifically, the average expectation is that 57% of the SOC workload will be handled by AI within this timeframe. Furthermore, a commanding 83% of security leaders anticipate that 50% or more of their SOC workload will be completed by AI. This strong, shared vision underscores AI as a strategic imperative, indicating that organizations are committed to investing in and finding solutions to circumvent current hurdles, as AI is perceived as essential for the future efficacy and scalability of their security operations.

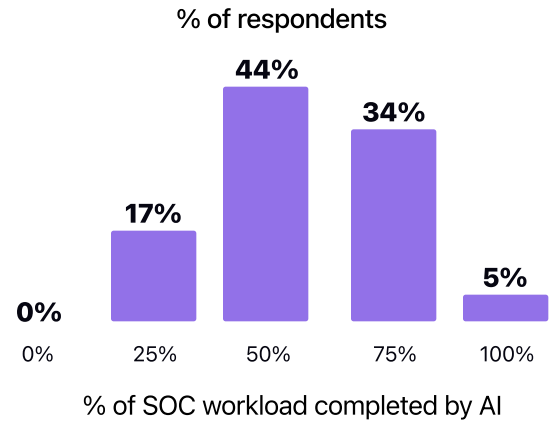

6 Key considerations when evaluating AI SOC solutions

Coverage

Strong coverage means the AI can handle alerts across identity, cloud, endpoint, email, network, and data loss. Broader alert types reduce manual triage, improve ROI, and ensure the AI adapts to different threats and data sources specific to your environment.

Accuracy

Accuracy builds trust. It reduces false positives, catches real threats, and prevents alert fatigue. Testing the AI under real conditions shows how well it handles complex and ambiguous alerts. Reliable accuracy means analysts can focus on response, not second-guessing automation.

Quality

High-quality investigations rely on depth, clear reasoning, and transparency. The AI should gather all relevant evidence, ask the right questions, and explain why it reached its conclusion. Tailoring its logic to your environment improves both accuracy and confidence in the results.

Workflow integration

AI is only useful if it fits into your existing workflows. Look for tools that connect smoothly with your systems and match how your team actually works. Flexible setup and meaningful integrations help avoid operational headaches and speed up real-world impact.

Time to Value

The faster the AI shows results, the better. It should reduce dwell time and investigation time without months of setup or retraining. If value takes too long to appear, that often points to deeper issues in design or deployment.

Data privacy & Security

AI must respect organizational boundaries. Unless otherwise agreed upon, it should never use your data to train external models or share sensitive information. Without clear privacy controls, trust breaks down. Choose tools that prioritize confidentiality and safeguard your internal security data from day one.

Methodology and Demographics

Methodology

This report, "The State of AI in Security Operations 2025," is based on an online survey conducted from May 22, 2025 - June 27, 2025. The primary objective of this survey was to gather comprehensive insights from security leaders and practitioners regarding the evolving challenges within Security Operations Centers (SOCs) and the current and prospective adoption of Artificial Intelligence (AI) solutions to address these issues.

The questionnaire was meticulously designed using Qualtrics and comprised a mix of question types, including Likert scales for sentiment measurement, multiple-choice questions for operational details, and open-ended questions to capture qualitative insights. To ensure clarity and relevance, a pilot study was conducted with a small group of security professionals, and their feedback was incorporated to refine the final survey instrument. Strict measures were implemented to ensure respondent anonymity and data confidentiality; no personally identifiable information was collected, and all findings are presented in aggregated form.

Participants were primarily recruited through targeted email invitations to our proprietary database of cybersecurity professionals, outreach via industry associations, and promotion on relevant professional networking platforms such as LinkedIn. A total of 282 security leaders completed the survey. Data validation involved the exclusion of incomplete (defined as less than 90% completion), unqualified responses, and the identification and removal of any potential duplicate submissions.

For data analysis, quantitative methods, including descriptive statistics (e.g., averages, percentages, frequencies), were primarily utilized to identify key trends and patterns. Cross-tabulations were performed to explore relationships between different variables. Qualitative data derived from open-ended questions was subjected to a thematic analysis to extract recurring sentiments and insights. All statistical analysis was performed using tools such as Microsoft Excel and Python libraries.

It is important to acknowledge that, as with any survey-based research, the findings are based on self-reported data and may be subject to inherent biases. While efforts were made to ensure a diverse respondent base, the generalizability of these findings should be considered within the context of the demographics presented.

Demographics

Survey collected over 3,500 responses but only 282 of those qualified and passed bot detection rules. Qualified respondents were compensated for their time completing the survey. Survey was fielded from May through June 2025. Question flow was designed to minimize response biases. All answer choices appeared in random order to eliminate response bias.

Role

% of respondents

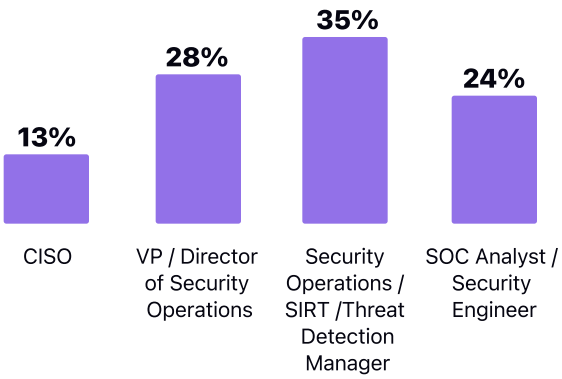

Company size

% of respondents

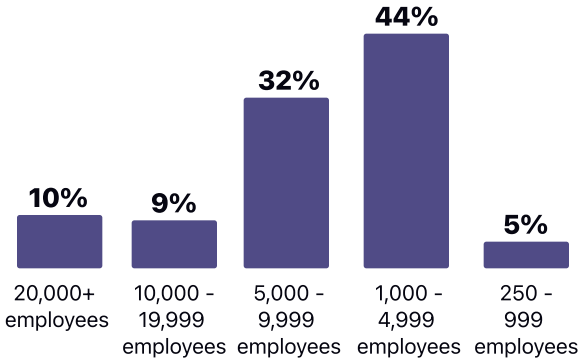

Industry

% of respondents

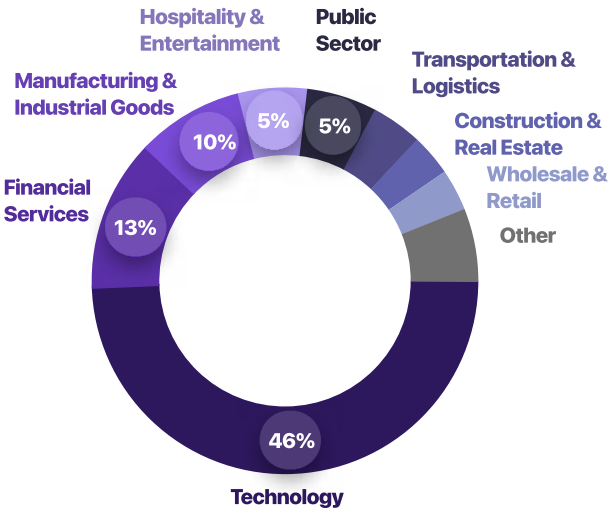

Geography

% of respondents

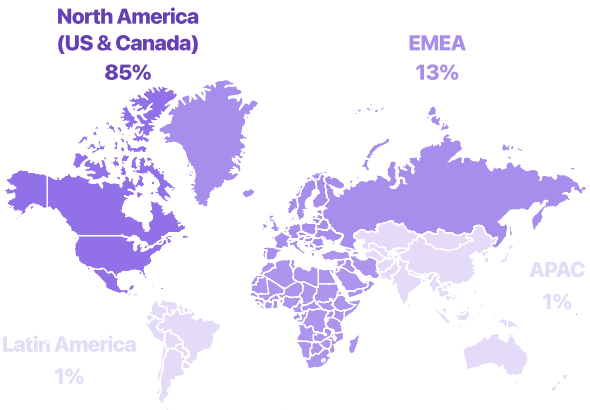

Appendix

The size of the SOC team tends to scale with the size of the organization

How many full-time employees are in your organization's Security Operatison Center (SOC) or equivalent

Number of respondents

Size of SOC team	20,000+ employees	10,000 - 19,999 employees	5,000 - 9,999 employees	1,000 - 4,999 employees	250 - 999 employees
None	0	0	0	3	0
1-4	0	1	0	3	1
5-9	0	0	2	10	1
10-24	5	5	16	31	3
25-49	3	4	26	28	5
50-100	5	6	22	35	2
Over 100	16	9	25	12	1

As expected, the largest enterprises (20,000+ employees) tend to have the largest SOC teams, with notable numbers of organizations reporting 50-100 or even over 100 full-time SOC employees. Conversely, smaller organizations (250-999 employees) predominantly operate with much leaner SOCs, often comprising just 1-4 or 5-9 analysts.

Even with larger teams in bigger organizations, the earlier findings of overwhelming alert volumes (e.g., 3,000+ alerts daily for 20,000+ employees) and a significant percentage of ignored critical alerts (60.6%) highlight that simply increasing human headcount does not fully resolve the "alert problem" or guarantee comprehensive coverage. This indicates that operational challenges persist regardless of team size.

Most organizations tend to rely on in-house SOC analysts, some also have 3rd party support

How is your organization triaging, investigating, and responding to security alerts today?

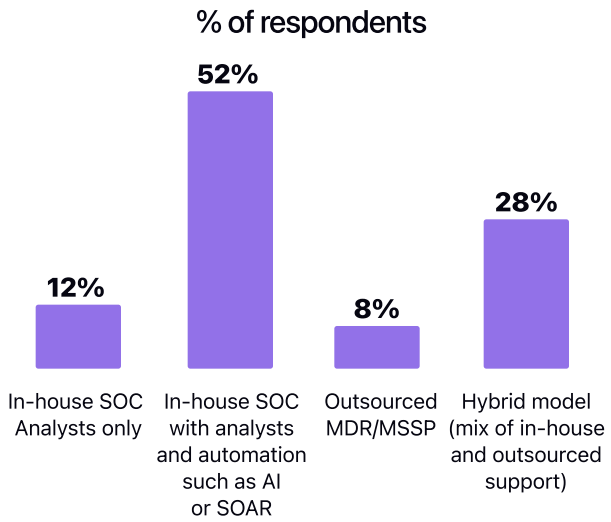

The data underscores a clear and widespread recognition that human capacity alone cannot keep pace with the demands of cybersecurity. The strong adoption of automation within in-house SOCs, combined with the prevalent use of hybrid models, signals a mature understanding that technology and strategic outsourcing are vital for effective alert management. This landscape creates a compelling demand for advanced AI SOC solutions that can seamlessly integrate with and enhance these hybrid and automation-driven in-house environments, promising to further optimize efficiency, reduce analyst burden, and improve overall security posture.

Automation is becoming essential for in-house SOCs: The largest segment, at 52%, indicates that most organizations are triaging, investigating, and responding to alerts with an in-house SOC team supported by automation, such as AI or SOAR. Organizations are actively investing in technology to augment their internal teams, aligning with earlier findings that AI is a top priority for SOC leaders.

Purely manual SOCs are a minority: Only 12% of organizations rely solely on in-house SOC analysts without the aid of automation. This suggests that a purely manual approach to alert management is becoming unsustainable due to alert fatigue, high dwell times, and the risk of missed critical alerts, pushing most organizations towards technological assistance or third-party managed security service providers.

Hybrid Models Bridge Capacity Gaps: A substantial 28% of organizations employ a hybrid model, mixing in-house capabilities with outsourced support (MDR/MSSP). This highlights a common strategy to address challenges like 24/7 coverage, specialized skill shortages, or scaling operations without fully relinquishing internal control. Purely outsourced models (MDR/MSSP) represent a smaller segment at 8%.

Glossary

Definitions

Alert fatigue	Overwhelm and desensitization of security analysts due to excessive alerts, leading to missed threats.
Alert suppression	Intentionally limiting the number of detection rules due to limited capacity to triage and investigate alerts.
Alert triage and investigation	The process of sifting through high volumes of security alerts, prioritizing critical incidents, and providing context for faster human analysis.
Cloud Security	Protecting data, applications, and infrastructure involved in cloud computing.
Data Security	Measures taken to protect digital data from unauthorized access, corruption, or theft.
Detection Engineering and Tuning	The process of refining detection rules to reduce false positives and ensure the efficacy of security controls.
Detection Rules	Specific configurations within security tools that trigger alerts when certain suspicious activities or patterns are identified.
Dwell Time	The period an attacker remains undetected within a system, from initial alert fire to being picked up for review.
False Positive	A security alert that indicates a threat when no actual threat exists.
Mean Time to Investigate (MTTI)	The average time it takes for a security team to thoroughly investigate an alert, from its initial trigger to understanding its nature, scope, and potential impact.
Mean Time to Respond (MTTR)	The total time from the initial alert firing to its complete containment, eradication, and recovery.
SOC (Security Operations Center)	A centralized unit responsible for an organization's security operations, including monitoring, detection, and response to threats.
SOC Workload	The sum of tasks and activities performed by a Security Operations Center team.
Tool sprawl	The proliferation of numerous security tools within an organization's security stack, leading to increased complexity and alert volume.
24/7 Coverage	The SOC's ability to maintain continuous, round-the-clock monitoring and response capabilities.

prophet

State of AI in Security Operations 2025

Key contributors:

George Dimitrov

Ajmal Kohgadai

Eric Swenson

Filip Stojkovski

Francis Odum