Gurucul

2025

Pulse of the AI SOC Report

From Alert Fatigue to Actionable Intelligence: How AI Is Reshaping Detection, Response, and Analyst Confidence

Research by

Cybersecurity INSIDERS

Executive Summary

Security operations teams are under pressure from every direction: a relentless data deluge, rising alert volume, the risks from limited identity and cloud visibility, fragmented tooling, and the widening disconnect between threat velocity and analyst capacity.

This report captures a shift in how global SOCs are responding. Identity-based threats are rising, analysts are stretched thin, and tool fragmentation is slowing response. Artificial Intelligence (AI) isn't just the next innovation—it's becoming the practical path to scale, speed, and resilience. Security leaders are no longer debating whether AI belongs in the SOC, they're focused on implementing it in a way that delivers results without introducing new risk.

Based on survey responses from 739 cybersecurity leaders, this research highlights where AI is already delivering results, where it's falling short, and what defines the next phase of SOC modernization. The picture that emerges isn't hype, but momentum grounded in measurable gains and a growing sense of urgency.

Key findings from this report include:

- **Identity and Human Risk are the Top Concerns—and the Least Visible:** Social engineering and phishing is a top concern, with 78% of security leaders still struggling with the human attack vector. Closely aligned are the identity-based threats, with 73% saying it is a top concern, yet 67% still lack visibility into access behavior and lateral movement. The most exploited threat vector is still the least monitored, as many organizations trade visibility for affordability—skipping data sources because ingestion pipelines are too costly or complex to integrate.

- **Alert Volume Keeps Climbing—While SOCs Fall Further Behind:** 88% say alert volume has increased, and 46% report a spike of over 25% in the past year. Alert fatigueis a top challenge for 76% while most SOCs still rely on dozens of platforms that don't correlate signals well.

- **Human Capital Is at a Breaking Point:** 73% of organizations report analyst burnout and persistent staffing shortages. 64% say their detection, triage, and investigation processes are still heavily manual, placing unsustainable pressure on small, overloaded teams already working across fragmented toolsets.

- **AI Adoption Is Accelerating, But Operational Use Is Still Limited:** 87% of organizations are deploying, piloting, or evaluating AI-powered SOC tools. But just 31% use them across core detection and response workflows. Interest is high, but the execution gap remains wide.

- **AI Automation Is Delivering Measurable Gains:** 60% of adopters have cut investigation time by at least 25%. Organizations that have implemented AI-powered automation are already seeing real ROI in the form of reduced investigation times, faster triage, and lower analyst fatigue.

- **Security Leaders Expect AI to Deliver Results:** 72% of CISOs are prioritizing faster investigation, 65% want to reduce alert noise, and 61% are investing in automation. These aren't AI use cases, they're executive priorities.

These insights reflect a SOC environment under pressure, and an industry leaning into AI not as a future promise, but as an operational necessity. What follows is a deeper look at where AI is gaining ground, where trust and integration still lag, and how leading organizations are aligning strategy to execution.

Table of Contents

01

CHAPTER 1:
The Evolving Threat Landscape

Identity: The Expanding Battleground

Attackers have learned that identity is the easiest way in and the hardest for defenders to monitor. Phishing and account takeover tactics are now so effective and cheap to execute that they continue to dominate the threat landscape, even as AI-enabled attack tools grow more sophisticated.

The data shows a clear hierarchy of concern: 78% of security leaders report social engineering and phishing as a top threat, followed closely by identity-based attacks at 73%. Ransomware (69%), cloud risks (56%), and AI-generated threats (48%) round out the top tier, revealing that most organizations are still focused on the risks they encounter most frequently—not just the ones grabbing headlines.

This mirrors what security teams are experiencing on the ground. Identity threats are stealthy, credentialed, and context-dependent, often evading traditional detection logic by mimicking legitimate access behavior. And as attackers increasingly chain techniques across email, identity, cloud, and SaaS surfaces, the blast radius of a single successful phishing lure or MFA bypass continues to grow.

From privilege escalation to lateral movement, attackers are now treating identity as both initial access and persistent control. Without real-time behavioral understanding of how accounts are used—not just which logs are generated—SOC teams are left reacting too late or missing the threat entirely.

Which threats are causing the greatest concern today? (select top 5)

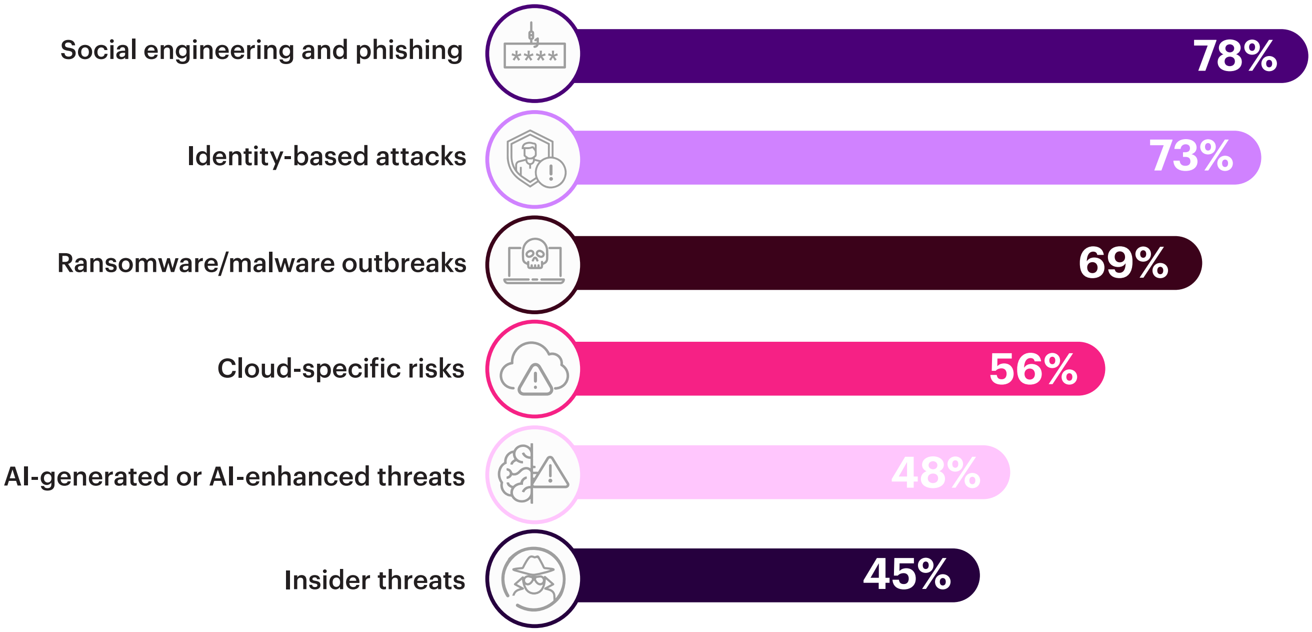

Identity Threats Demand Real-Time Detection

Identity is no longer just a risk factor—it's a frontline battlefield. But while security teams recognize the importance of identity-based threat detection, few are equipped to act fast enough.

Eighty percent of respondents say identity-based detection is highly important, with over half calling it mission — critical to SOC success. Yet most organizations still lack the behavioral visibility, data correlation, and contextual intelligence needed to reliably detect identity misuse in real time.

Unlike signature-based threats, identity-driven attacks unfold quietly. A compromised credential or misused session token can allow an attacker to operate with legitimate access across multiple systems—especially in SaaS and cloud environments where visibility is fragmented. Traditional detection tools often miss these low-noise, high-impact movements.

Preventing these attacks requires platforms that not only collect logs but also understand intent, detect subtle behavioral deviations, and correlate across identity, device, and location, providing SOC teams with actionable context instead of reactive noise.

How important is it for your organization to detect and respond to identity-based threats (e.g., account takeover, privilege abuse, credential misuse)?

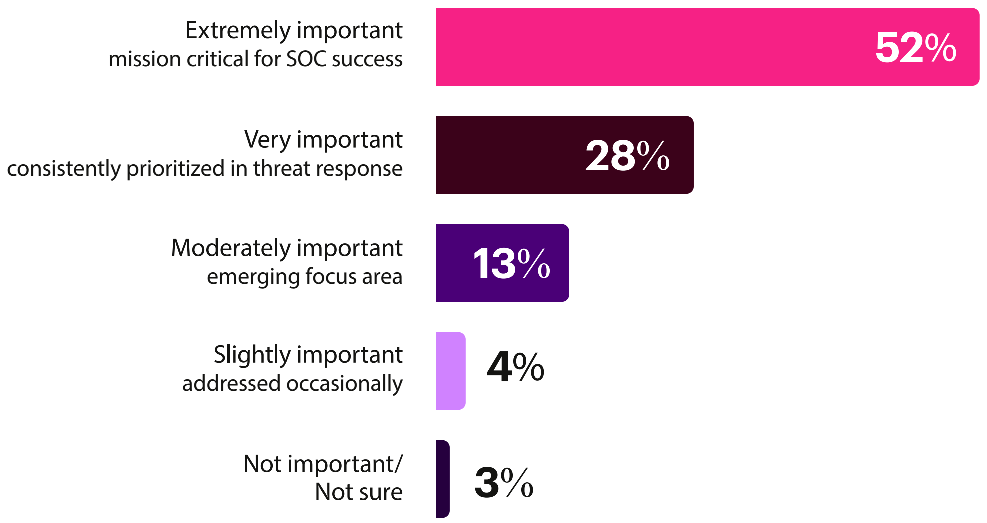

80%
of organizations say detecting identity threats is highly important to SOC success

02

CHAPTER 2:

Why the SOC Is Breaking

Where Security Operations Are Failing Under Pressure

SOCs today are buried—not just in alert volume, but in disconnected tools, fragmented telemetry, expanding cloud workloads, and siloed data that's hard to act on.

The survey confirms what most teams already know: alert fatigue is overwhelming operations. Seventy-six percent cite it as a top challenge, followed closely by analyst burnout (73%). Even well-resourced teams are falling behind, with 64% pointing to manual investigations and 59% citing tool sprawl as a major operational drag.

Notably, evolving threats that include AI-powered techniques are rising fast, yet only 55% of respondents selected them as a top-five challenge. This reflects both their novelty and the fact that existing complexity in detection workflows continues to eclipse newer attack trends for many organizations.

Which of the following are the top 5 most pressing challenges facing your SOC today? (select top 5)

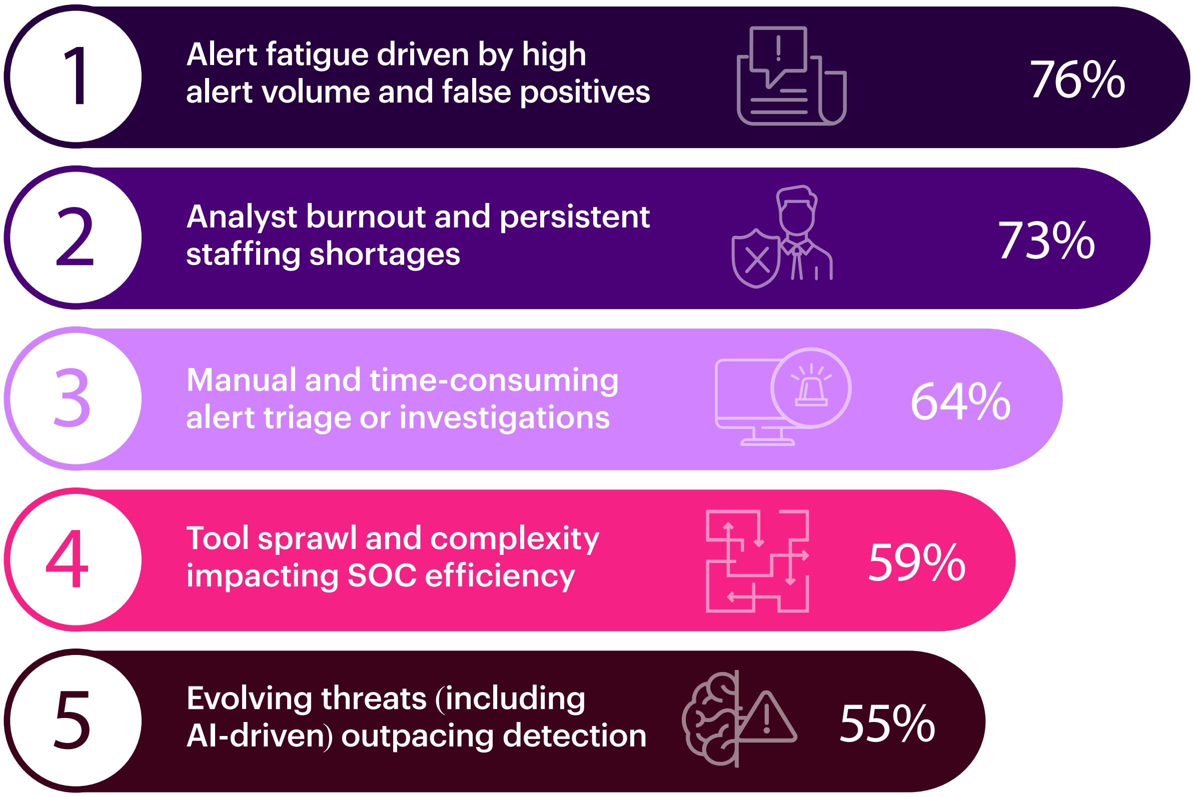

What teams need isn't more dashboards or another layer of tooling. They need unified visibility from a single interface and full context from a platform that can analyze any signal, from any source, in real time. That means moving from fractured tools to integrated workflows, powered by AI that not only prioritizes what matters but helps reduce the cognitive load on analysts.

Improvements in AI-driven correlation, triage, and threat interpretation are already proving essential to managing this complexity. The ability to combine disparate signals into a cohesive picture and act faster with fewer manual steps has become a core requirement for SOC performance at scale.

Alert Volume Keeps Climbing

The survey data shows why SOCs feel so overwhelmed: 88% said alert volume has increased, with nearly half (46%) reporting a spike of more than 25% in just the past 12-24 months. In some cases, new tooling is generating more complexity, more alerts, and more ambiguity. SOCs are suffering at both ends of the spectrum: critical visibility gaps in areas like cloud, SaaS, and identity, and overwhelming noise from the data they do collect. The result is a flood of alerts without the context to prioritize what matters.

Solving this isn't just about reducing data and noise but about elevating signal. Platforms must prioritize what matters, suppress what doesn't, and deliver enriched, contextualized insights that accelerate triage instead of slowing it down. In today's SOC, the biggest threat isn't just an attacker. It's everything competing for your analyst's attention.

How has the volume of security alerts in your SOC changed over the past 12–24 months?

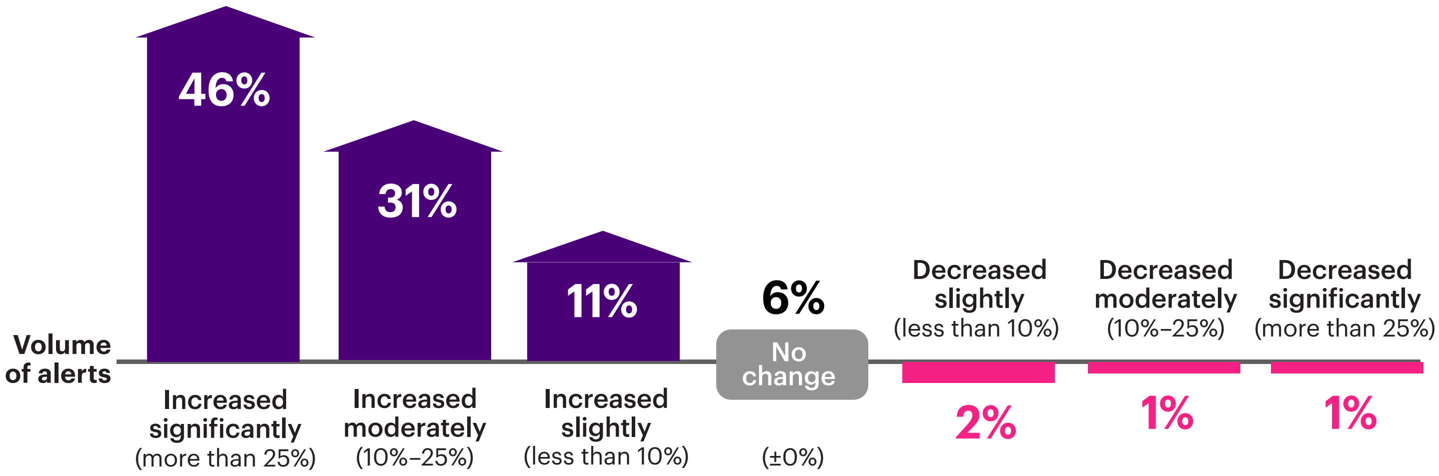

Not sure 2%

SOC Visibility Gaps Are Fueling Risk

Visibility gaps remain one of the most persistent and dangerous weaknesses in the modern SOC. Shockingly, only 4% of respondents say they have full visibility across their security data estate. The remaining 96% report critical blind spots—most commonly in cloud infrastructure (74%) and identity and access behavior (67%).

These gaps are more than operational nuisances; they directly map to three of the top four threat concerns identified in this survey: identity-based attacks like account takeover and MFA bypass, phishing and social engineering (including AI-enhanced lures), and cloud-specific risks stemming from misconfigurations and vulnerabilities. Endpoint telemetry (62%), SaaS application activity (59%), and encrypted east-west traffic (54%) further illustrate how fragmented and incomplete visibility becomes as infrastructure decentralizes.

Where does your SOC face the most significant data visibility gaps? (select all that apply)

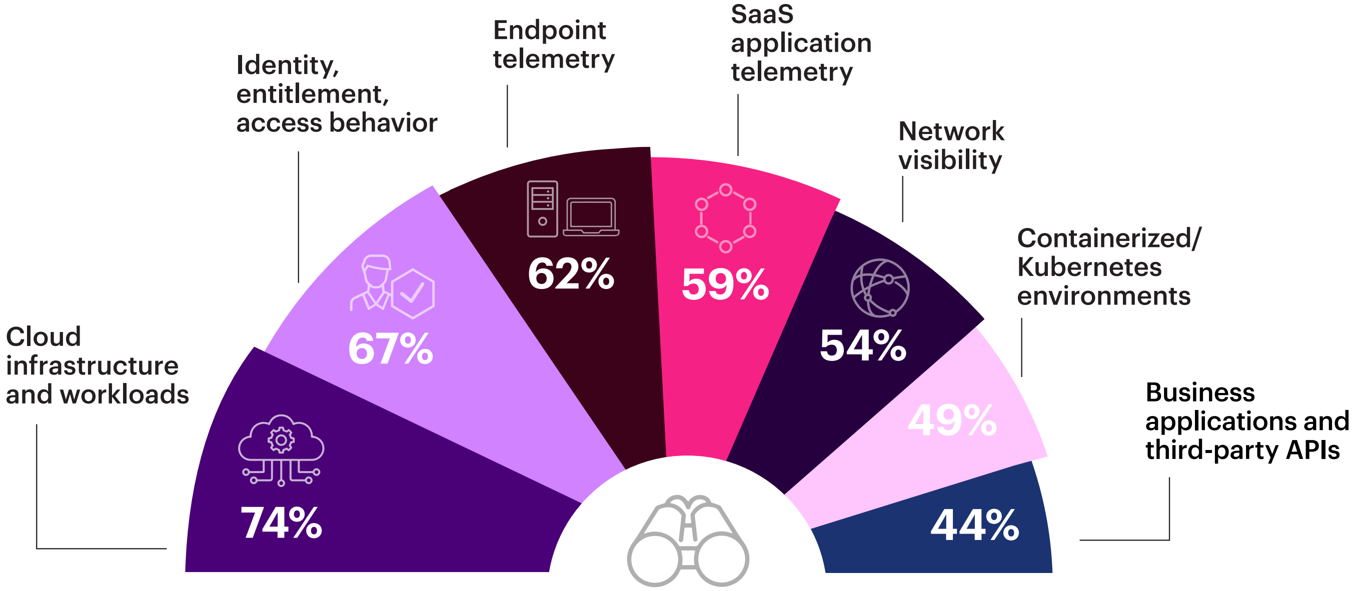

Here's a common failure scenario: a contractor logs into a cloud storage bucket using a valid token. Over the course of an hour, they download thousands of internal documents, including proprietary IP. The activity generates signals in the identity system and cloud telemetry, but without unified correlation, no alert is triggered. The access appears legitimate in isolation and is not flagged in real time. Only later, during a post-incident investigation, is the full scope of the activity pieced together from log evidence.

These telemetry gaps aren't usually caused by technology limitations but by operational tradeoffs. Security leaders often face difficult decisions about which data sources they can realistically afford to ingest. In some cases, teams delay onboarding high-value telemetry because they're waiting on a vendor to build a parser or support a new integration. In others, the decision is budget-driven: the cost of bringing certain data into the SIEM outweighs what current budgets allow. The result is that critical signals—especially from cloud, SaaS, and identity systems—may be missed or underutilized, not because they're unavailable, but because they weren't prioritized at the point of integration. What SOCs need isn't just broader data access, but platforms that can unify, correlate, and contextualize identity, privilege, and behavior in real time across a fractured ecosystem.

Fragmented Tooling Is Undermining Detection

Visibility gaps aren't always the result of missing tools—often they are the product of too many. The previous data showed that cloud and identity remain the least monitored surfaces in the SOC. But even where telemetry is available, fragmentation of tooling continues to undermine effective detection and response.

Only 12% of respondents report using fewer than 10 tools for threat detection, investigation, and response. The majority operate in far more complex environments: 38% use between 10 and 19 tools, 29% report using 20 to 29, and 16% even use 30 or more. This fragmentation isn't just an integration headache, it introduces operational friction at nearly every stage of the incident lifecycle.

How many distinct security tools does your organization currently use for threat detection, investigation and response?

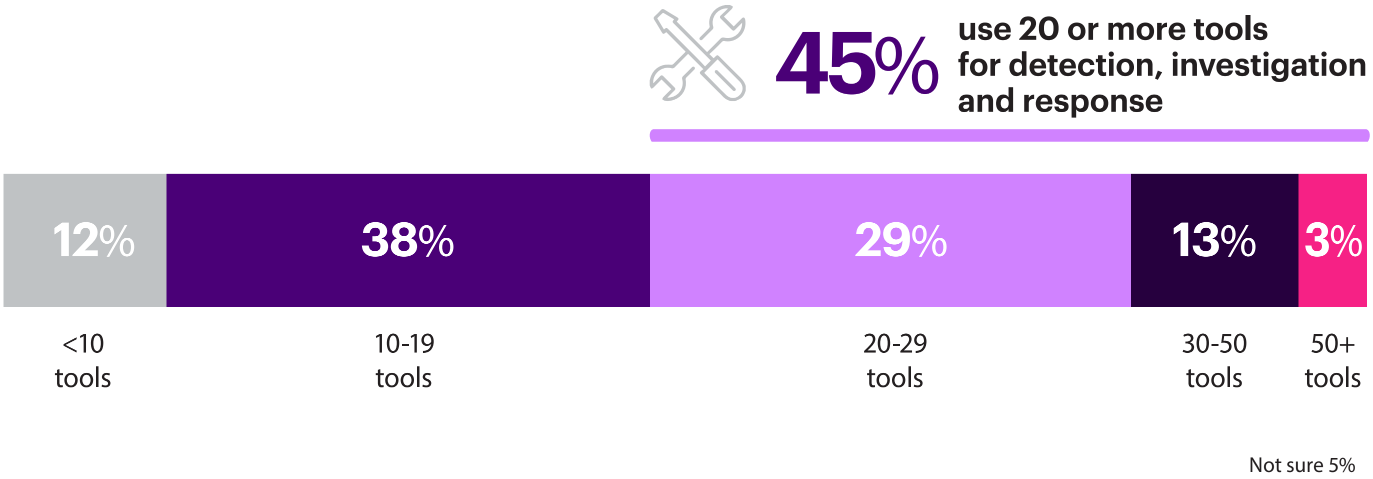

Each tool collects different signals, applies distinct analytics, and focuses on different slices of the threat surface, often looking for different things. An identity outlier in one system, an endpoint anomaly in another, and unusual cloud activity in a third may all describe the same incident, but without correlation, the context is missed across silos. The delay isn't in the availability of data—it's in assembling a coherent picture fast enough to act on it.

Solving this requires more than centralizing logs. It demands platform-native detection with systems built to correlate telemetry, apply behavioral models, and investigate across domains from the start. As infrastructure grows more distributed, response effectiveness will depend less on the number of tools in use and more on how well they integrate and break down silos to function as a unified system.

Slow Data Onboarding Delays Detection

Tool fragmentation isn't the only factor delaying incident response. Even when organizations deploy the right tools, the process of onboarding new data sources into the SIEM remains a major operational bottleneck.

Just 4% of respondents say they can fully onboard a new feed in under a day. The largest group (41%) reports a timeline of one to four weeks, while 32% say it takes between one and six months. That means more than 70% of organizations spend weeks or longer before new telemetry becomes actionable. In practice, that lag often affects high-priority sources like cloud logs, SaaS activity, and identity data.

When a new security data source needs to be ingested into your SIEM, how long does it typically take to fully configure, filter, and optimize the pipeline?

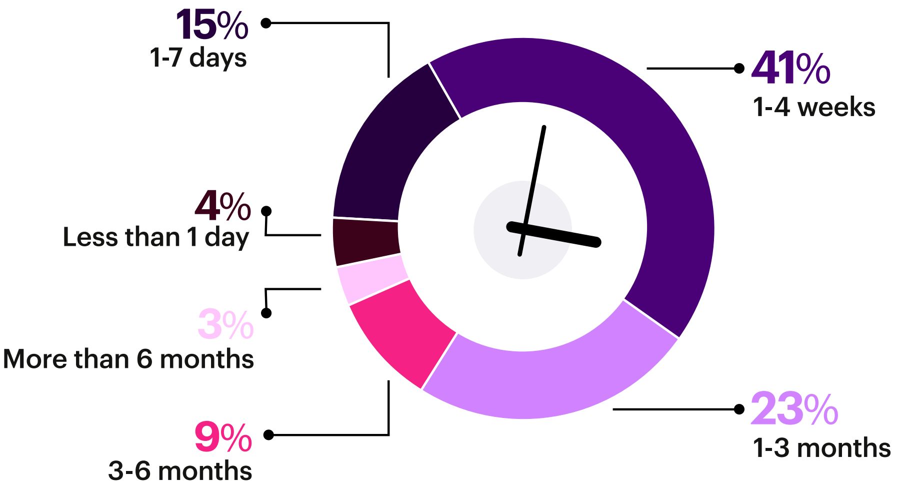

64% take between a week and 3 months to configure, filter and optimize their SIEM pipeline

Don't know 5%

The issue isn't just ingestion. It's the time-consuming process of transforming each feed: normalizing schemas, filtering noise, and mapping data to detection use cases. Until that work is done, analysts are flying blind in the areas where visibility is needed most.

To minimize exposure windows, security teams need data ingestion architectures that are schema-agnostic, dynamically adaptable, and capable of applying filtering and correlation logic before data ever reaches the SIEM. Emerging solutions—such as dedicated data pipeline management platforms and pre-processing layers within modern SIEMs—are starting to address this challenge by normalizing, enriching, and shaping telemetry upstream. These approaches not only reduce onboarding delays but also help manage high-volume data costs while consolidating visibility into a more unified detection layer—ultimately reducing tool sprawl and supporting a single source of truth.

Most Organizations Are Dissatisfied – And Forced to Augment

SIEM remains a foundational part of the SOC, but most current deployments are falling short. Just 22% of organizations say they are fully satisfied with their SIEM and have no plans to change. The other 78% are either dissatisfied, stuck with limitations, or forced to augment to achieve broader detection goals.

Of that 78%, 31% report moderate satisfaction but are actively augmenting their SIEM with additional tools such as UEBA, ITDR, insider risk platforms, data pipeline management, and data protection solutions. Meanwhile, 47% express outright dissatisfaction: 16% are actively evaluating replacements, 15% are blocked by the cost or complexity of switching, and 10% haven't started evaluating yet.

Which best describes your organization's plans for its current SIEM over the next 12–24 months?

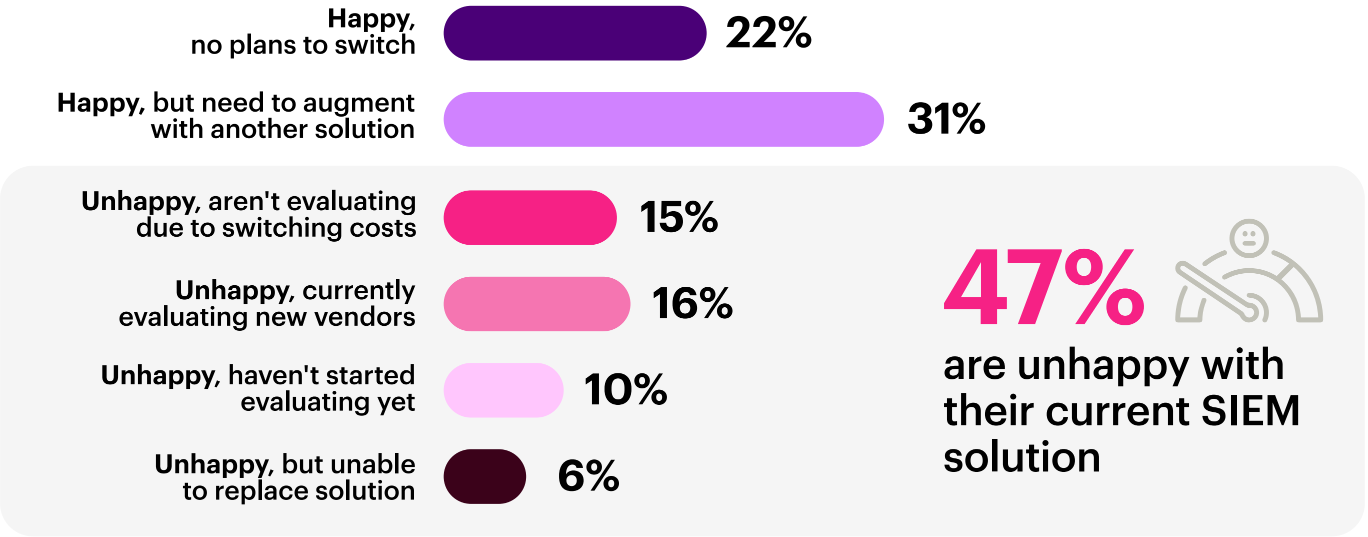

This widespread frustration reflects a deeper structural issue: most SIEMs weren't designed for today's dynamic, identity-driven threat landscape. As visibility requirements grow and threats become more cross-domain, many teams find themselves stitching together detection with auxiliary tools—adding more complexity to already fragmented environments.

While SIEM augmentation can address gaps, it also carries risk. Many SIEMs still act mainly as IT log collectors, meaning a replacement may not fully control the security function. Sophisticated SOCs may add best-of-breed tools to maintain flexibility and avoid vendor lock-in, but each new system can fuel tool sprawl, complexity, and data silos. Without seamless integration, detection suffers and the SOC becomes a fragmented patchwork that can weaken security.

Rather than layering tools indefinitely, organizations are calling for a shift: SIEM must evolve from a log-centric data lake into a unified detection fabric. That means native support for behavioral analytics, cloud identity visibility, context-rich correlation, and real-time investigation workflows, delivered through a platform capable of serving as the central lifeblood of the SOC.

03

CHAPTER 3:
AI Enters the Equation

AI Adoption Is Quickly Gaining Ground

With alerts piling up and talent stretched thin, many SOCs are now utilizing AI—not to replace analysts, but rather to augment and give them a fighting chance to keep pace. Traditional detection methods, such as rules, signatures, and manual triage, are no longer enough to keep up with today's speed and scale of operations, combined with the increasing volume of attacks. AI promises not just gains in productivity, but transformation: triaging noise, correlating behavioral signals, and guiding analysts with contextual precision.

Many teams are already making this shift: 31% of organizations report actively using AI across multiple SOC workflows, from detection and triage to enrichment and response. Another 34% are running targeted pilots, and 22% are actively evaluating use cases. Taken together, 87% of organizations are progressing toward AI adoption and integration within the SOC—not as an experiment, but as a strategic shift in how the SOC operates.

What best describes your organization's current adoption of AI in security operations?

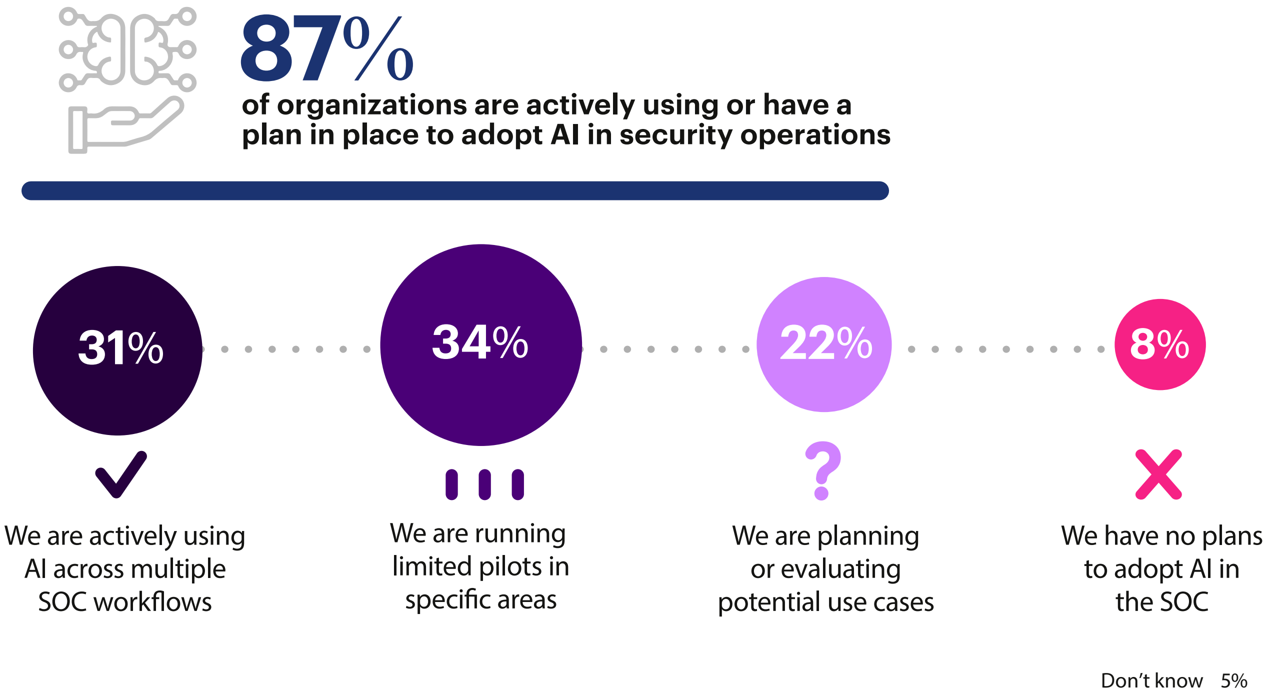

87%
of organizations are actively using or have a plan in place to adopt AI in security operations

The trend is clear: for many security leaders, AI has moved from experiment to necessity, despite ongoing concerns about accuracy, explainability, and risk. The benefits now outweigh the risks, particularly as SOC workloads become unmanageable without automation. As AI adoption increases, the next challenge is operationalizing it. That means embedding AI into daily workflows, aligning it with SOC processes and risk models, and giving analysts the transparency and trust they need to act. Teams that succeed in operationalizing AI won't just resolve incidents faster; they'll detect attacks that would otherwise slip through legacy tools entirely.

AI Automation Is Now a Strategic Imperative

As AI adoption accelerates, automation is emerging as its most urgent and operationally impactful use case in the SOC. Security teams are betting on AI automation to offload repetitive work and respond faster than humans can manage alone.

Seventy-nine percent of respondents say AI-powered automation will be either mission-critical or a key part of their SOC strategy within the next 24 months. Nearly half (45%) say it will be essential. This marks a fundamental shift: AI automation is no longer an enhancement—it's becoming foundational to how modern SOCs function. Organizations are now planning for systems that can triage alerts, suppress false positives, and initiate low-risk responses autonomously.

How important is AI-powered automation for your SOC's effectiveness over the next 24 months?

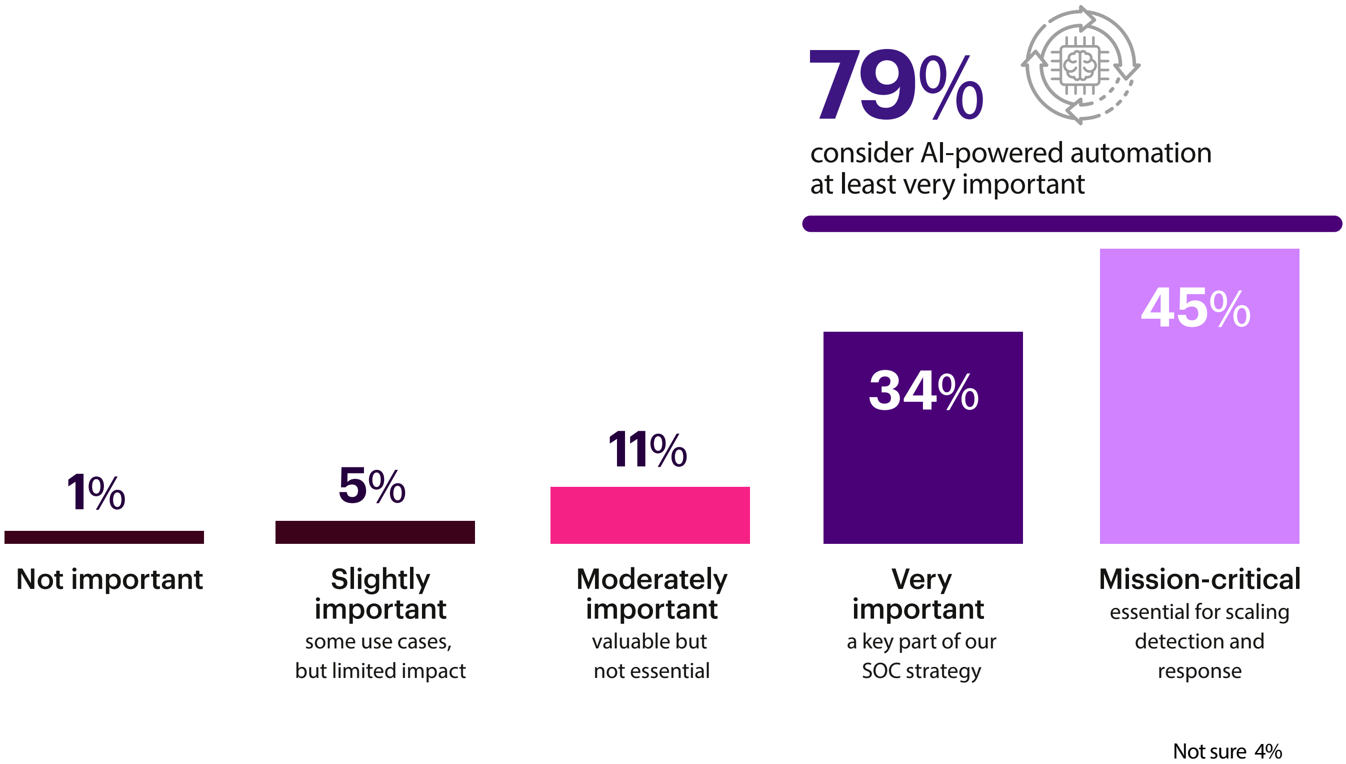

To keep pace with modern threats, SOCs need more than enrichment—they need AI that takes action, often referred to as 'agentic AI' for its ability to operate autonomously within defined guardrails. Automation is not just a cost-saver or convenience layer; it's the only way to scale threat detection, investigation, and response (ITDR) without overwhelming the human layer. The organizations that move first will be the ones that shift from reacting to leading.

AI Automation Is Already Cutting Investigation Time

The shift to AI automation is already delivering measurable gains. Sixty percent of respondents report that automation has reduced investigation time by 25% or more, with 21% seeing reductions greater than 50%. Another 27% report savings between 10–25%.

But with alert volumes rising by more than 25% in many organizations, a new question emerges: are these gains enough to reverse alert fatigue or simply to keep it from getting worse? For many SOCs, automation isn't eliminating the burden—it's keeping operations afloat. Even modest improvements scale quickly when applied across thousands of incidents and a full analyst team. A SOC with 10 analysts saving just 15 minutes per investigation, across hundreds of alerts each week, can reclaim thousands of hours annually. That's the difference between needing to expand headcount and doing more with the same team.

▶ **Looking ahead 2-3 years, which AI-driven advancement do you believe will have the greatest impact on threat detection and response in your SOC?**

RANK			
1	36%	**Analyst augmentation and investigation acceleration**	Empower analysts with auto-triaged alerts, expert guidance, natural language search, and enriched insights with attack blast radius identification.
2	29%	**Adaptive response orchestration**	Dynamically modify and execute response playbooks based on real-time information, evolving threat conditions, and a unique understanding of every incident or attack nuance.
3	20%	**Attack chain learning and model refinement**	Detect evolving attack chains to inform new ML detection models, signatures, and rules while optimizing existing and recommending new data sources.
4	15%	**Autonomous data pipeline optimization**	Minimize the need for manual data onboarding and tuning with autonomous data discovery, classification, normalization, enrichment, and filtering.

Consider phishing response, a daily drag on most SOCs. Before automation, investigating a suspicious email required analysts to inspect headers, review threat intel, check delivery logs, and escalate findings manually. With AI-powered systems in place, low-risk messages are suppressed, high-risk ones are pre-triaged and enriched, and analysts receive a decision-ready incident. What once took an hour now takes 10 minutes, and the time saved is reinvested in more strategic work.

This is how AI proves its value: not by generating more alerts, but by removing workflow friction. Automation restores capacity, accelerates response, and gives analysts the breathing room to focus on what machines can't do.

Security Leaders Are Prioritizing Speed, Clarity, and Scalable Automation

As organizations look to the future of security operations, their goals are pragmatic, focused, and aligned around operational impact.

The number one objective, selected by 72% of respondents, is accelerating the speed of investigations and incident response. That's followed by reducing alert volume and false positives (65%) and increasing automation (61%). These themes reflect a clear intent: improve signal-to-noise ratio, reduce analyst workload, and enable faster, more confident action. Rounding out the top tier are expanding response capacity (57%) and strengthening advanced detection capabilities (51%)—showing a desire to add elasticity to the SOC, but without sacrificing accuracy or context.

Over the next 12 to 24 months, which objectives are most important for improving your organization's security operations? (select up to 5)

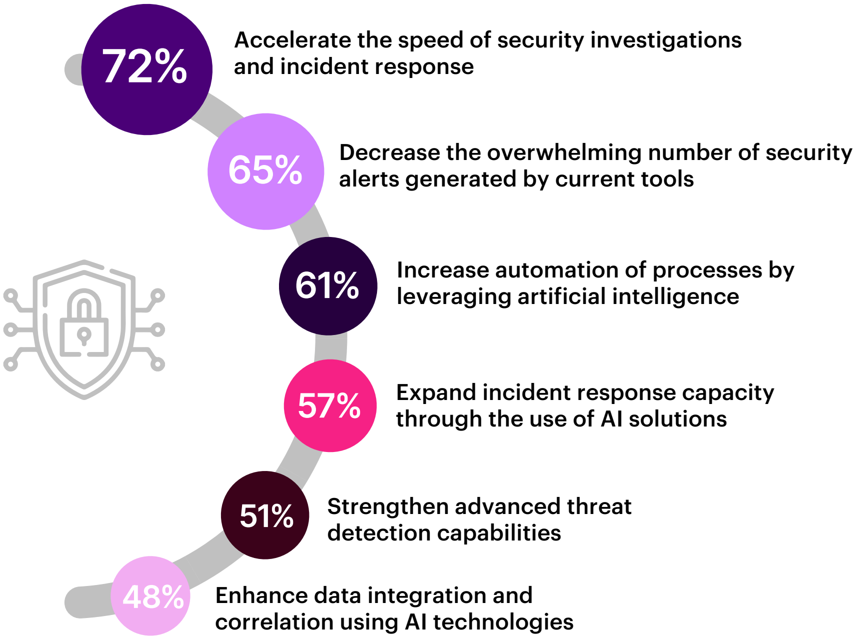

Interestingly, only 48% cite tool consolidation as a top priority. While complexity remains a known issue, most teams appear focused first on measurable operational outcomes, such as speed, clarity, and control, rather than full architectural overhaul.

What's notable is how well these priorities align with the demonstrated strengths of AI. From triage acceleration to enrichment, risk scoring, and low-noise correlation, the tools that are gaining traction are the ones that convert volume into insight and action. The goals haven't shifted—what's changed is the clarity and urgency with which leaders are pursuing them. The platforms that can execute on those priorities will define the next generation of SOC performance.

What Organizations Expect from AI in the SOC

As security teams look ahead to how AI will reshape detection and response, their expectations are focused, measurable, and grounded in operational reality. For all the speed of change, the value organizations expect from AI remains consistent: faster detection, less noise, and stronger analyst productivity.

The top outcome security leaders are targeting is faster detection of real threats, selected by 68% of respondents. Improved analyst productivity and reduced burnout follows closely at 61%, alongside reduced mean time to respond (MTTR) at 56%. Just over half (51%) are seeking a reduction in false positives and alert volume, while 46% are prioritizing streamlined data pipelines and lower ingestion costs. What ties these goals together is a shared focus on scalability without compromise: doing more, with greater clarity and speed, and less human strain.

Which operational outcomes does your organization hope to achieve from adopting AI-powered SOC platforms? (select top 3)

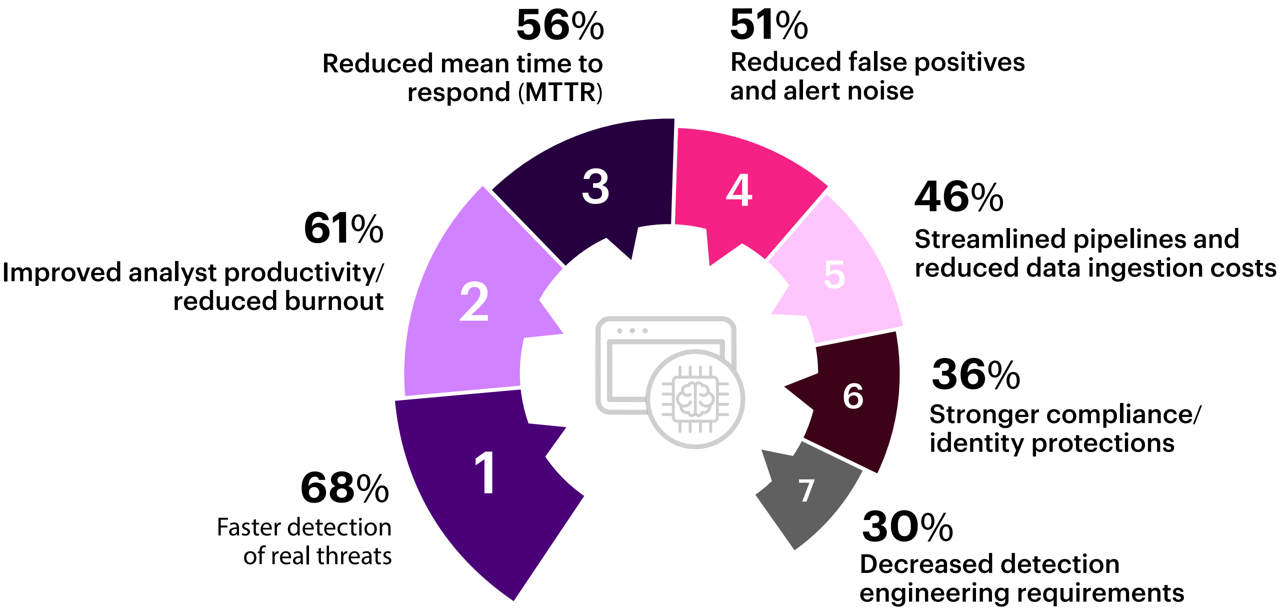

This clarity of intent signals an important shift: AI is no longer being explored for its potential, but for its ability to produce outcomes that legacy SOC tools and processes can't reliably deliver. Organizations aren't chasing innovation for its own sake but investing in impact.

The next wave of investment will flow to platforms that can execute on this mandate: automating triage, enriching context across domains, suppressing low-value alerts, and empowering analysts to focus on decisions, not data. Because as AI becomes more deeply embedded in SOC workflows, its value won't be measured by how much it replaces humans, but by how much more effective it makes them.

AI Budget Growth Reflects Strategic Priority

After years of aspirational hype, AI in the SOC is being funded with intent. The goals are clear—faster detection, reduced analyst strain, and smarter response—and security leaders are putting budget behind platforms that can deliver.

The vast majority of organizations anticipate steady growth: 8% of respondents expect their AI-SOC budgets to increase by more than 20%. Thirty-one percent project moderate increases between 10–20%, and 39% expect increases under 10%.

What are your budget trend expectations for AI-powered SOC solutions over the next 12–18 months?

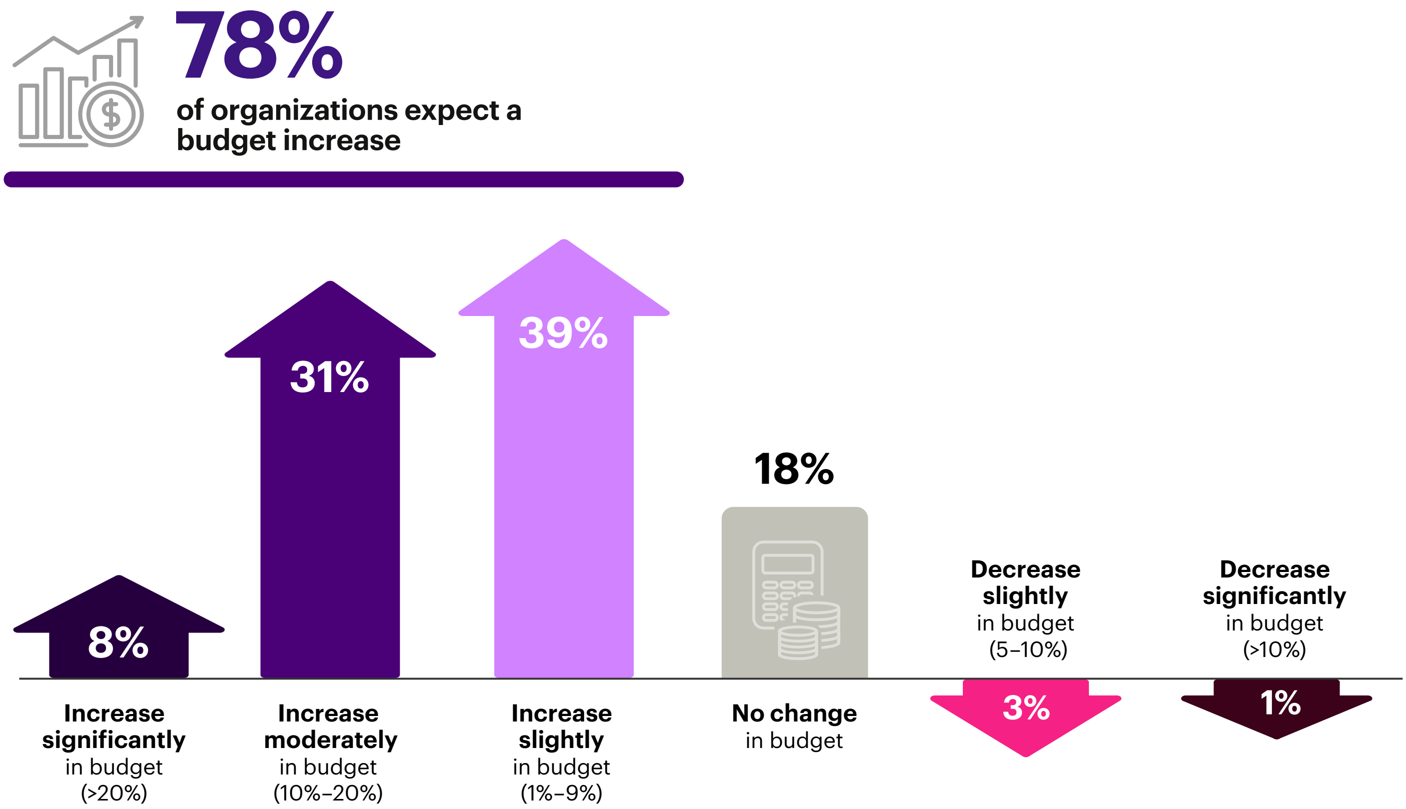

Only 18% expect flat budgets, and almost none anticipate a decline.

This signals a broader shift: AI tools are now being evaluated not by potential but by real performance. Organizations are moving from pilot projects to platform integration, and funding is following results. AI initiatives that reduce false positives, shorten investigation time, or improve incident containment will continue to earn support. Those that don't will be cut.

To maintain this momentum, security teams must align AI investment with measurable outcomes—faster triage, reduced MTTR, improved analyst productivity. In a resource-constrained environment, AI solutions must defend their place in the stack. Not with dashboards, but with impact.

Building Analyst Trust in AI

As organizations increase investment in AI-powered SOC capabilities, the next challenge comes into focus: building trust in the systems analysts are expected to rely on. Confidence isn't yet universal, but it's growing through real-world use, and trust is emerging as the next frontier for AI maturity.

Just 9% of respondents say they are "very confident" in AI-generated alerts and recommendations. Thirty-three percent mostly trust the output with some review, while 41% find AI generally helpful but still require frequent validation. Only 5% express outright distrust. These results point to a healthy learning curve: analysts are working alongside AI, calibrating trust with experience, and gradually increasing confidence in the system's recommendations.

How confidentare your analysts in the accuracy of alerts and recommendations produced by AI-powered detection tools?

83% of organizations are moderatly to very confident in the accuracy of alerts

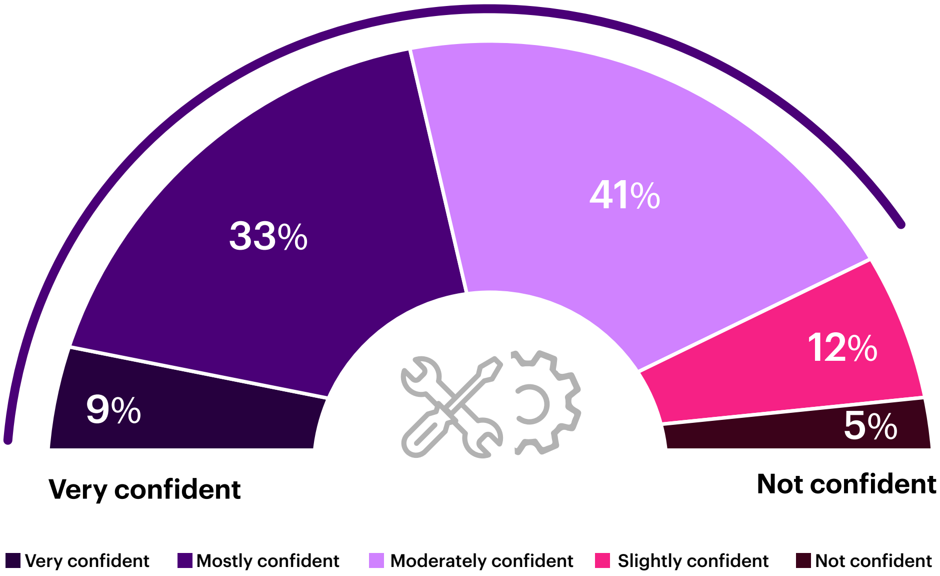

Trust isn't just technical—it's human. Some hesitation stems from uncertainty around accuracy, but much of it is psychological: concerns about being replaced or overridden by systems that largely operate in black boxes. When automation accelerates without clarity, it can breed resistance. But when AI is designed to support—not supplant—human decision-making and is able to explain its reasoning, trust begins to take hold.

That's where the next phase of AI adoption must focus: explainability, context, and collaboration. The most effective AI platforms won't just surface anomalies, they'll show their work. Analysts don't need more alerts. They need systems that reason, communicate, and reinforce trust with every recommendation—keeping humans firmly in the loop.

AI Is Automating Core SOC Workflows

As trust in AI-powered detection grows, security teams are beginning to capitalize on its operational value. The early wins are clear and concentrated where they matter most: repetitive, high-volume workflows that have long drained analyst capacity.

Seventy-three percent of respondents report successful automation of alert triage and prioritization. Sixty-eight percent say enrichment tasks like threat intelligence lookups and asset correlation have been automated effectively. Further down the stack, 59% report AI-driven threat intel correlation, and 50% have automated false positive suppression and detection rule tuning. While deeper-stage tasks like automated response playbooks and incident summary generation are gaining traction, they remain less common, reflecting a thoughtful, phased approach to automation maturity.

Which SOC analyst workflows have been most successfully automated using AI in your organization? (select all that apply)

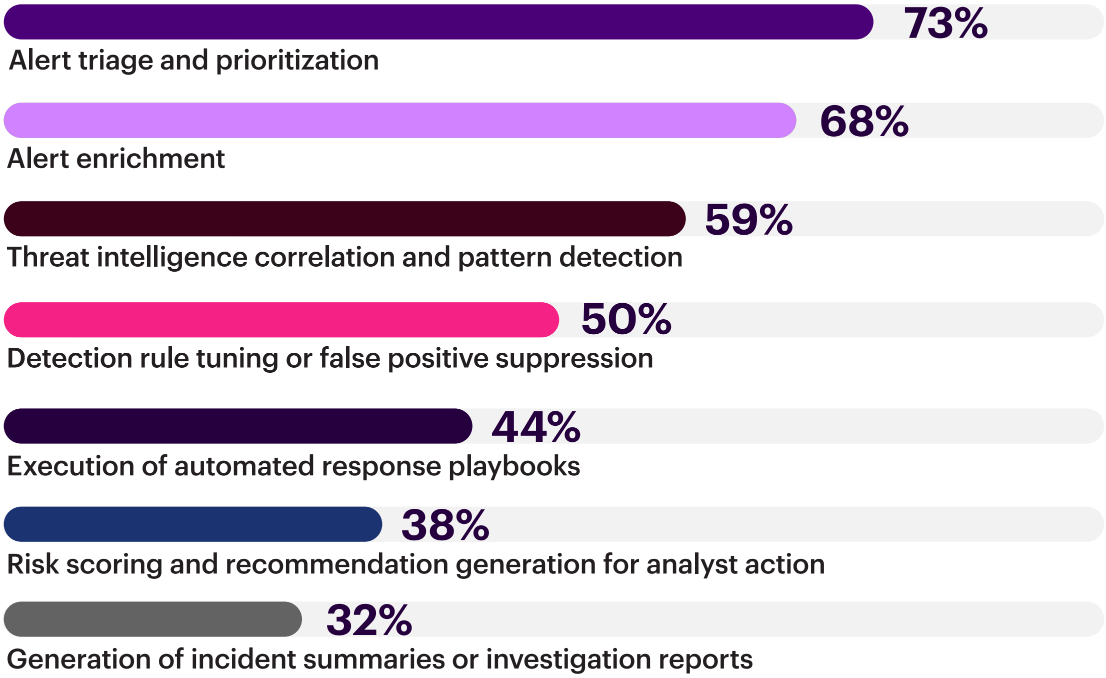

In practice, this means high-volume, low-discretion tasks are increasingly offloaded to AI—giving analysts time to focus on decisions, not data prep. An incoming alert for suspicious outbound traffic, for example, can now be pre-triaged—correlated with known threat intel, enriched with identity and asset context, and scored for risk—within seconds. What once took significant time is now immediate. And instead of a research assignment, the analyst is handed a decision-ready incident.

This is where AI is already delivering value: reclaiming time, reducing noise, and raising confidence at the front lines of security operations. The next stage will require more than workflow logic. It will demand platforms that act transparently, adaptively, and with the trust analysts need to let go of the keyboard and let automation lead.

Best Practices for Operationalizing the AI-Powered SOC

One message from the data is consistent: AI is already being deployed, not just discussed. The real differentiator isn't whether teams are using AI, it's how deeply it's embedded into workflows and whether it's solving real operational pain. Here's where leading organizations are focusing their effort:

1 **START WHERE ANALYSTS ARE OVERWHELMED.**

Seventy-three percent of respondents report successful AI automation of alert triage and prioritization—the clearest early win in the stack. Ambitious playbook automation can come later. The starting point is simpler: eliminate the daily friction that slows triage, drains time, and wears out analysts.

2 **CORRELATE IDENTITY AND BEHAVIOR, NOT JUST EVENTS.**

Cloud infrastructure and identity access are the top two visibility gaps reported in this survey (74% and 67%, respectively). Attackers know that credentials, roles, and behavioral anomalies often pass through undetected by traditional SIEMs. Closing that gap requires platforms that can make sense of who is doing what, across systems—not just what logs say happened.

3 **LET AI AUTOMATION ACT, BUT MAKE SURE IT EXPLAINS ITSELF.**

Sixty percent of organizations using AI report investigation time reductions of 25% or more. But confidence still lags with only 9% of analysts being "very confident" in AI-generated outputs. If analysts can't see how a decision was made, they'll second-guess it or ignore it. The most effective systems show their work, surface reasoning, and support decision-making—not override it.

4 **ALIGN EVERY AI INITIATIVE TO A MEASURABLE OPERATIONAL OUTCOME.**

Executives aren't funding AI for the sake of innovation. They're prioritizing faster investigations (72%), alert reduction (65%), and automation that scales without additional headcount (61%). The AI platforms that earn budget are the ones that directly improve MTTR, analyst productivity, and SOC focus.

5 **DON'T WAIT FOR AI TO BE PERFECT. DEPLOY WHERE IT ALREADY WORKS.**

The SOC doesn't need to automate everything all at once. It needs to automate the right things first. Focus on Tier 1 work, build trust in the early gains, and scale gradually—with visibility, context, and human judgment built in from the start.

AI is already reducing investigation time, offloading repetitive triage, and giving analysts room to breathe—and think more strategically. But turning early wins into lasting impact means moving past pilots, past hype, and into workflows that actually deliver. The next phase is about making AI work everywhere it should—clearly, transparently, and with results that speak for themselves.

Methodology and Demographics

The 2025 Pulse of the AI SOC Report is based on an online survey conducted in early 2025, gathering responses from 739 cybersecurity professionals worldwide. Respondents included CISOs, SOC leaders, analysts, architects, and IT security managers from industries such as financial services, healthcare, manufacturing, government, energy, and technology. The survey specifically targeted organizations that manage their own SIEM environments, either fully in-house or through a hybrid model involving shared responsibility. Organizations that fully outsource their SOC operations were excluded to ensure the findings reflect firsthand operational experience with SIEM tooling, detection workflows, and AI adoption.

A stratified sampling approach ensured balanced representation by role and organization size, yielding a 95% confidence level with a ±3.6% margin of error. Some questions allowed multiple selections, resulting in response totals exceeding 100%. The findings offer a data-driven view into the challenges, priorities, and adoption patterns shaping the next phase of AI-powered SOC operations.

CAREER LEVEL

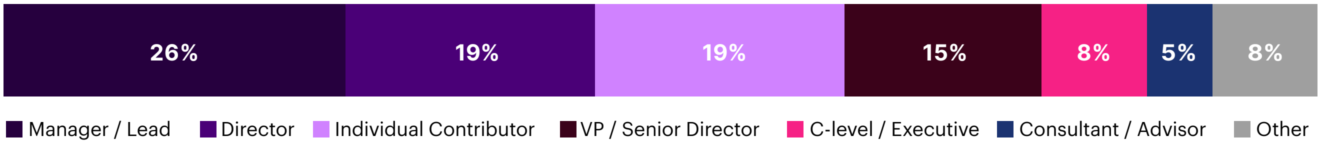

PRIMARY JOB FUNCTION

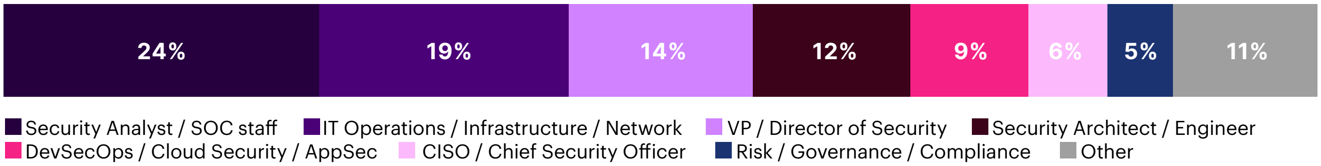

DEPARTMENT

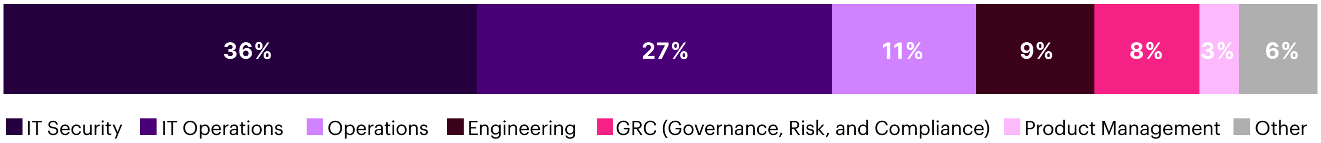

COMPANY SIZE

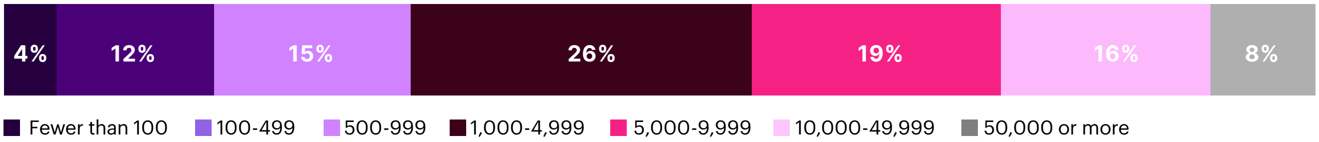

Reuse of content

We encourage the reuse of data, charts, and text published in this report under the terms of this Creative Commons Attribution 4.0 International License. You're free to share and make commercial use of this work as long as you attribute the report as stipulated in the terms of the license. For example: "Source: 2025 Pulse of the AI SOC Report by Gurucul and Cybersecurity Insiders."

A Smart SIEM for the Smarter SOC

The most valuable resources in the SOC are your people, not your tools. Unfortunately, the results from Cybersecurity Insiders 2025 Pulse of the AI-SOC Report tell a concerning story—SOC teams are inundated with false positives, cumbersome investigations, and tool complexity.

Gurucul offers a human focused, AI infused SIEM—capable of reducing costs by at least 40%, false positives by 70%, investigation time by 58%, and MTTR by 83%.

Behavioral Informed Detections, Not Static Rules

4,000+ machine learning detection models put behavioral deviations into context and prioritize true threats with logical risk scoring. You'll discover known and unknown threats, while expanding SOC use cases to cover identity-based threats and insider threats from a unified detection platform.

Optimized to Maximize Analyst Output, Not Hinder Them

An army of AI agents across the entire data and threat lifecycle automates mundane work without relinquishing control. A 24/7 AI SOC analyst who never sleeps offers full transparency into decision making as they auto-triage, escalate, and respond—leaving the critical thinking to your superior human analysts.

Delivers Ultimate Data Democracy, Not Vendor Lock-In

Decoupling analytics from storage with a bring-your-own data lake approach means YOU own your data. Combined with the industry's only native Data Pipeline Manager, costs are within your control. Normalize, enrich, filter, route, and search any data from any source with lean data pipelines acutely aligned to detection models—reducing data and engineering costs for ingestion.

Designed for Customer Choice, Not One-Size-Fits-All

Augment your SIEM or replace it. Deploy it as SaaS, in your cloud or on-prem. Unlock every detection use case, or a subset. The choice of where you start and how you scale is 100% yours. We offer a flexible and modular design with full customizability across data pipelines, detection models, response playbooks and dashboards and reporting. With 10,000+ ready day one content pieces you'll get immediate value and the ability to fine-tune your smarter SOC with ease.

Learn More

You can learn more about Gurucul Next-Gen SIEM here

Request a Demo

Or, request a demo today to discuss your biggest SOC challenges

Gurucul

Gurucul is the only cost-optimized security analytics company founded in data science that delivers radical clarity about cyber risk. Our REVEAL security analytics platform analyzes enterprise data at scale using machine learning and artificial intelligence. Instead of useless alerts, you get real-time, actionable information about true threats and their associated risk. The platform is open, flexible and cloud native. It conforms to your business requirements so you don’t have to compromise. Our technology has earned us recognition from leading industry analysts as the most Visionary platform and an Overall leader in product, market and innovation. Our solutions are used by Global 1000 enterprises and government agencies to minimize their cybersecurity risk.

To learn more, visit Gurucul.com

and follow us on LinkedIn and Twitter.

Cybersecurity
INSIDERS

STRATEGIC INSIGHT FOR CYBERSECURITY LEADERS

Cybersecurity Insiders delivers evidence-backed insights that empower security leaders to make informed, strategic decisions. Backed by over a decade of research and a global network of 600,000+ cybersecurity professionals, we provide actionable intelligence to help leaders navigate emerging threats, evaluate new technologies, and shape forward-looking strategies with confidence.

For cybersecurity vendors, we turn research into results—delivering credibility, visibility, and demand through high-impact formats such as:

- Data-powered market reports that establish thought leadership,
- Webinars that build trust with buyers through credible, expert-led narratives,
- CISO guides that showcase best practices,
- Product reviews that independently validate solutions,
- Thought leadership articles that educate buyers, and
- Award programs that elevate brand reputation.

By combining this content with built-in distribution, we help brands earn trust, amplify awareness, and drive demand in a crowded cybersecurity market.

For more information visit

cybersecurity-insiders.com