CROGL

The State of SecOps

AI in the SOC: What's Working, What Isn't, and What Comes Next

Based on independent research by Ponemon Institute | January 2026 | n=649

Research Report | Sponsored by Crogl | Conducted by Ponemon Institute

Five findings. One clear direction.

Crogl commissioned Ponemon Institute to survey 649 IT and IT security practitioners across North America on the real state of security operations in 2026: how effective SOC teams actually are, where AI is helping, where it isn't, and what the best-performing organizations are doing differently. What follows is what they told us.

What high performers do differently

34% of respondents demonstrate consistently strong security postures. They are more likely to keep SecOps in-house, more likely to have deployed AI in the SOC, and more deliberate about AI governance. 72% of high performers prioritize consistency in AI use, compared to 54% of others. They are not just deploying more AI. They are deploying it with more discipline, more visibility, and more human oversight.

The bottom line

Attackers are already deploying agentic AI: systems that don't just probe and adapt, but autonomously decide and act across entire kill chains without waiting for human instruction. A human analyst reviewing alerts at human speed is not a match for that. Neither is a static AI running yesterday's detection logic. Agentic threats require agentic defense. That means AI that can reason across data in real time, surface what matters before the window closes, and keep a human analyst informed and in control at every step. The organizations building that capability now, on top of data infrastructure that can actually support it, are the ones that will be effective when it matters most.

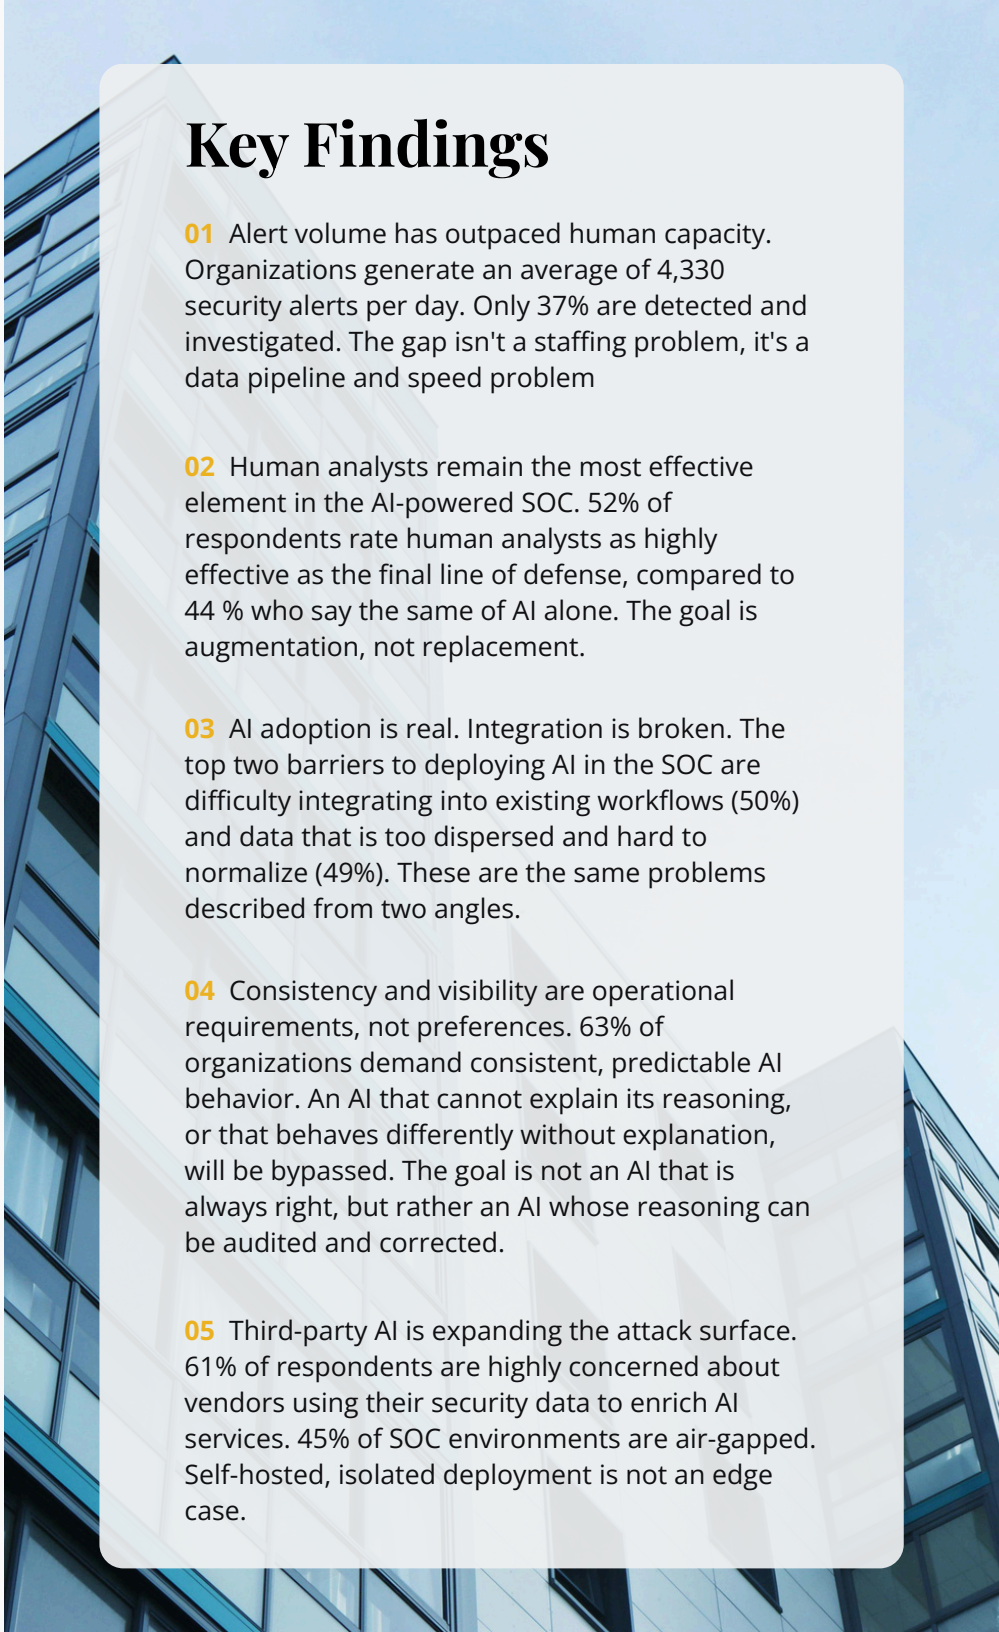

Key Findings

01 Alert volume has outpaced human capacity. Organizations generate an average of 4,330 security alerts per day. Only 37% are detected and investigated. The gap isn't a staffing problem, it's a data pipeline and speed problem

02 Human analysts remain the most effective element in the AI-powered SOC. 52% of respondents rate human analysts as highly effective as the final line of defense, compared to 44 % who say the same of AI alone. The goal is augmentation, not replacement.

03 AI adoption is real. Integration is broken. The top two barriers to deploying AI in the SOC are difficulty integrating into existing workflows (50%) and data that is too dispersed and hard to normalize (49%). These are the same problems described from two angles.

04 Consistency and visibility are operational requirements, not preferences. 63% of organizations demand consistent, predictable AI behavior. An AI that cannot explain its reasoning, or that behaves differently without explanation, will be bypassed. The goal is not an AI that is always right, but rather an AI whose reasoning can be audited and corrected.

05 Third-party AI is expanding the attack surface. 61% of respondents are highly concerned about vendors using their security data to enrich AI services. 45% of SOC environments are air-gapped. Self-hosted, isolated deployment is not an edge case.

The SOC is not failing for lack of effort.

Security operations teams are working harder than ever. The problem is the data. Every day, SOC analysts face thousands of alerts, fragmented across disconnected systems, in formats that don't talk to each other, and they're expected to investigate all of it, in real-time, with the same headcount they had three years ago.

AI was supposed to fix this. And it will, but not by replacing analysts. Not yet, and perhaps not ever in the ways vendors have promised. What the data actually shows is more nuanced, more instructive, and ultimately more actionable.

Crogl sponsored this independent research by Ponemon Institute to get an honest picture of where security operations stand today: the real state of alert management, how AI is actually being used in the SOC, where it's falling short, and what separates high-performing organizations from everyone else.

4,330

Average security alerts generated per organization, per day

Only 37% are detected and investigated.

52%

Rate human analysts as the most effective final line of defense in an AI-powered SOC

44%

Rate AI alone as highly effective in reducing threats

50%

Say difficulty integrating AI into workflows is the biggest deployment barrier

49%

Say data that's too dispersed and hard to normalize is the #2 barrier

61%

Are highly concerned about third-party AI vendors using their security data

Self-hosted and air-gapped deployment options are no longer optional for security-conscious organizations.



FINDING 01

The alert crisis is real and getting worse.

The average organization in this study generates 4,330 security alerts in a single day. Seven IT staff members are responsible for managing them. Those analysts investigate just 37% of what comes in.

That means nearly two-thirds of all security alerts go uninvestigated on an average day. Not because analysts aren't skilled. Not because the tools aren't running. But because the volume is simply unmanageable at human speed.

4,330	37%
Avg. security alerts per day	**Percentage actually detected and investigated**

Only 43% of organizations rate themselves as effective or highly effective at identifying and responding to security alerts. That figure reflects something important: having a SOC and having an effective SOC are two very different things.

The way organizations measure SecOps performance reflects this reality. 79% specifically track the number of security alerts investigated and resolved, which signals that alert throughput has become the de-facto proxy for SOC effectiveness. Organizations know they're leaving threats on the table, and they're measuring the gap.

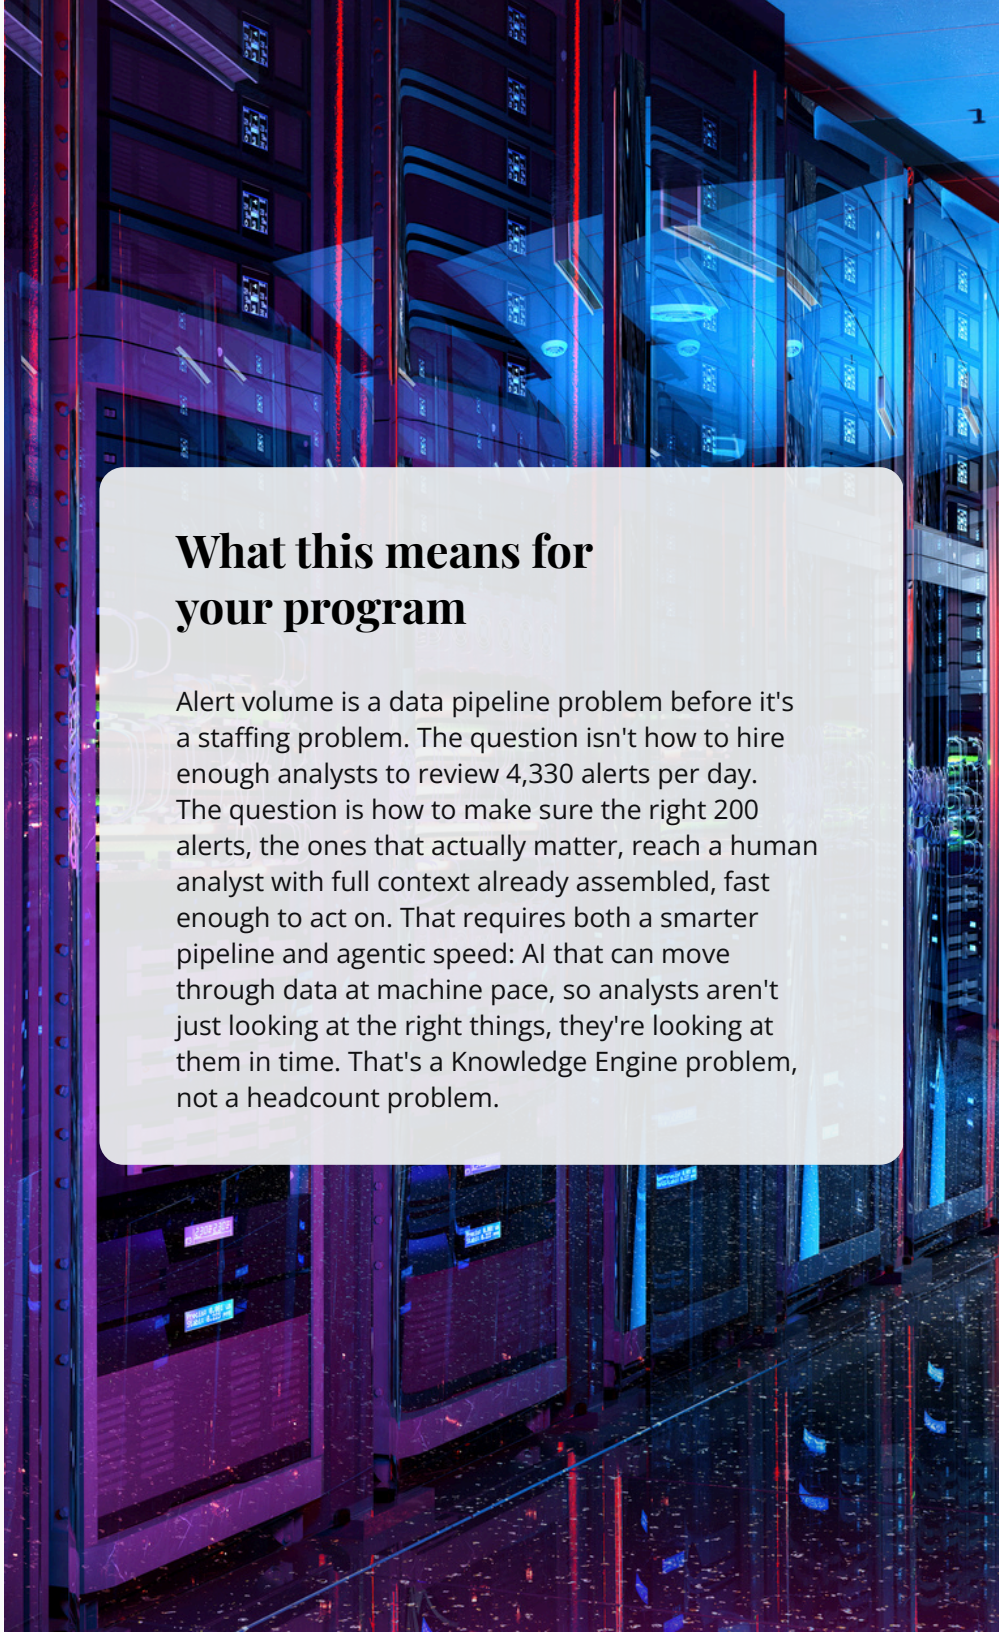

What this means for your program

Alert volume is a data pipeline problem before it's a staffing problem. The question isn't how to hire enough analysts to review 4,330 alerts per day. The question is how to make sure the right 200 alerts, the ones that actually matter, reach a human analyst with full context already assembled, fast enough to act on. That requires both a smarter pipeline and agentic speed: AI that can move through data at machine pace, so analysts aren't just looking at the right things, they're looking at them in time. That's a Knowledge Engine problem, not a headcount problem.

Human analysts are still the most important variable in the AI-powered SOC.

One of the clearest findings in this research: the most effective element in an AI-powered SOC is still the human analyst. 52% of respondents rate human analysts as highly effective as the final line of defense, a higher mark than AI alone, which only 44% rate as highly effective in reducing threats.

52%

Rate human analysts as highly effective as the final line of defense

vs. 44% who rate AI alone as highly effective in reducing threats

This isn't a knock on AI. It's a clarification of what AI is actually good for, and what it isn't. AI excels at pattern recognition, speed, and processing volume. It is not, at this stage, equipped to exercise judgment, interpret ambiguous signals, or make high-stakes decisions with incomplete information. Those are human capabilities.

The most common AI use cases in the SOC reflect this division of labor. Organizations using AI in the SOC report benefits in: reducing complexity of security cases to improve analyst efficiency and collaboration (54%), automatic generation of documentation for complex automated processes (49%), and use of tools that have AI embedded in them (41%).

The primary benefits of an AI-powered SOC are similarly practical: speeding up the time to resolve alerts (67% of respondents), freeing analyst bandwidth for urgent incidents and strategic projects (57%), and improving triage and remediation of Tier-1 and Tier-2 alerts (53%).

Primary benefit of an AI-powered SOC	% Respondents
Speed up time to resolve more alerts faster	**67%**
Free analyst bandwidth for urgent incidents and strategic projects	**57%**
Improve triage, investigation, and remediation of Tier-1 and Tier-2 alerts	**53%**
Increase ability to meet compliance mandates	**50%**
Increase SOC team's ability to detect and respond to threats	**46%**
Decrease cost of cybersecurity operations	**44%**

The data supports a straightforward conclusion: AI should make analysts faster, not make analysts optional. The organizations getting the most out of their AI investments aren't the ones trying to automate their way out of having a SOC team. They're the ones giving their analysts better information, faster.

Augmentation principle

The most effective AI in the SOC operates like an exoskeleton, not an autopilot. It gives human analysts enhanced reach, speed, and information, while keeping them firmly in control of the decisions that matter. That distinction isn't philosophical. It shows up directly in the data.

AI adoption is real. Integration is broken.

62% of organizations have adopted AI in some form. Of those with a SOC, 57% have deployed AI specifically within it. These are not trivial numbers, AI in security is not a future state, it's an operational reality for the majority of organizations in this study.

But there is a large gap between having AI and having AI that works at scale. When asked about the biggest barriers to deploying AI tools in the SOC, the answers were remarkably consistent.

50%

Say difficulty integrating AI into workflows is their biggest deployment barrier

49%

Say data dispersed and hard to normalize is their #2 barrier

These two barriers aren't separate problems. They're the same problem described from two angles. The reason AI is hard to integrate into SOC workflows is precisely because the underlying data is fragmented, inconsistently structured, and scattered across systems that were never designed to talk to each other.

When asked what features matter most when choosing AI tools for the SOC, 63% of respondents cited the ability to integrate into existing workflows and data as the top requirement, above consistency, above cost, above any other feature. And yet this is exactly where most AI deployments fail.

The normalization trap

Traditional AI-in-security approaches require organizations to normalize their data before the AI can use it. But data normalization at SOC scale is a project that takes months, requires dedicated engineering resources, and begins breaking the moment new data sources come online. The organizations stuck in this cycle aren't waiting on AI to improve, they're waiting on their data to be clean enough for AI to start. That's the wrong order of operations.

Regulatory and compliance concerns round out the top barriers at 37%, a figure that will only grow as AI governance frameworks continue to mature. Organizations deploying AI in security need to be able to demonstrate what the AI did, why it did it, and how decisions were made. That requires visibility into the AI lifecycle that most current deployments don't provide.

57% of respondents say visibility into an AI's development, training data, and deployment is highly important, yet only 36% rate their own organization's ability to detect new forms of data leakage from AI tools as high. There is a significant gap between what organizations know they need and what they've been able to build.



Consistency, visibility, and documentation aren't nice-to-haves.

When AI is making decisions that affect security outcomes, the standards for explainability and consistency are different than in other business applications. Security teams need to be able to audit what happened, reproduce findings, and demonstrate compliance to regulators and internal stakeholders. That requires a level of AI discipline that most organizations are still building toward.

63%

Say consistency in AI's use for security operations is very or highly important

58%

Say AI documentation of SecOps tasks is very or highly important



These numbers reflect a genuine operational concern, not an abstract preference for best practices. When an AI system surfaces an alert, marks a case as resolved, or recommends a remediation action, there must be a record of how that determination was made. When the AI behaves differently on Tuesday than it did on Monday with no explanation, analyst trust erodes. And when analyst trust erodes, the AI gets bypassed.

Consistency is the prerequisite for adoption. Organizations that can't guarantee predictable, repeatable AI behavior will find their analysts working around the system, not with it.

But here's what consistency actually requires in a security context. In dynamic environments where threats evolve constantly, an AI system that never makes mistakes isn't robust. It's overfitted. Like security controls tuned to zero false positives, an AI optimized for perfect accuracy will miss real threats. Systems that never experience stress become brittle. When they encounter something genuinely novel, they fail.

This reframes what AI-ready actually means. It's not about having perfect, normalized data before your AI can operate. It's about data infrastructure that lets AI systems learn from errors, adapt to new patterns, and stay effective across the full range of security scenarios, including the ones that don't fit historical patterns. The goal isn't an AI that's always right. It's an AI whose reasoning you can see, audit, and correct when it isn't.

The audit trail imperative

The most important feature in an AI SOC tool isn't its detection accuracy on day one, it's its ability to explain, document, and reproduce its outputs over time. 63% of respondents already understand this. The question is whether their current AI deployments are built to support it.

FINDING 05

Third-party AI is expanding your attack surface. Most organizations already know it.

One of the most underappreciated risks in the AI-in-security conversation is what happens to your data once it leaves your environment. When security teams use cloud-hosted AI tools, especially those trained on or enriched by customer data, they are extending their trust boundary in ways that standard vendor management frameworks weren't built to handle.

The survey respondents are clear-eyed about this risk. 61% are very or highly concerned about third parties using their security data to enrich their AI services. 59% are very or highly concerned about AI vendors using their data for derivative purposes.

61%

Highly concerned about third parties using their security data to enrich AI services

59%

Highly concerned about AI vendors using their data for derivative purposes

These aren't hypothetical concerns. Security data is uniquely sensitive, it reveals infrastructure topology, detection logic, response playbooks, and the specific vulnerabilities an organization is most exposed to. When that data is processed by a third-party AI, it becomes part of a training corpus or enrichment pipeline that the organization no longer controls.

The concern is especially acute for organizations handling regulated data, operating in government and defense environments, or managing air-gapped networks. Of the SOC respondents in this study, 45% indicate their SOC protects an air-gapped network, environments where cloud-first AI tools are simply not an option.

Self-hosted AI is not an edge case

45% of SOC respondents operate in air-gapped environments. For these teams, SaaS-hosted AI tools are off the table by default, not by preference, but by requirement. Any AI solution designed for enterprise security operations needs to be deployable in isolated, self-hosted, and on-premises configurations. The market that demands this is larger than most vendors acknowledge.



High-performing organizations don't just use AI differently. They're structured differently.

34% of respondents report high confidence in their security posture's ability to mitigate risks, vulnerabilities, and attacks across the enterprise. We call these high performers. Their responses diverge meaningfully from the rest of the survey population, and the patterns are instructive.

Metric	High Performers	All Others
Keep SecOps in-house	**47%**	34% in-house (avg.)
Effective at gathering evidence and finding threat sources	**68%**	54%
Effective at detecting and responding to threats	**63%**	53%
Have deployed AI in the SOC	**68%**	46%
Rate human analysts as highly effective final line of defense	**58%**	46%
Rate AI as highly effective in reducing threats	**51%**	37%
Prioritize AI consistency in security operations	**72%**	54%
Prioritize visibility into AI lifecycle	**64%**	50%

The pattern is clear. High performers are more likely to keep their SecOps function in-house, more likely to have deployed AI in the SOC, and more likely to rate both human analysts AND AI as highly effective. These aren't competing conclusions. High performers have figured out how to use AI to make their analysts better, not to replace them.

They're also more deliberate about AI governance. 72% of high performers (vs. 54% of others) say consistency in AI's use is very or highly important. 64% (vs. 50%) prioritize visibility into the AI lifecycle. High performers aren't just deploying more AI — they're deploying it more carefully.

The in-house advantage

Only 20% of high performers outsource their SecOps team, compared to 34% of other respondents. Owning your security operations, including the data it generates, the AI that analyzes it, and the analysts who act on it, is a distinguishing characteristic of organizations that consistently outperform. Control over the data is control over the outcome.

CONCLUSION

The data isn't complicated. The implementation is.

Read across all five findings, a coherent picture emerges. Security operations teams are drowning in alerts they can't investigate. AI is being deployed to help, but it's running into the same data infrastructure problems it was supposed to solve. The organizations getting the most out of AI are the ones that keep humans firmly in the loop, govern their AI deployments carefully, and maintain ownership of their data and infrastructure.

None of this is surprising to practitioners who've lived it. What the data adds is specificity — and a clear set of priorities for organizations looking to make real progress.

What high-performing SOC teams do differently:

- They keep security operations in-house, maintaining direct control over data, tooling, and analyst judgment.
- They deploy AI to augment analyst capacity, not to remove analysts from the decision loop.
- They treat data normalization not as a prerequisite for AI, but as a problem to solve architecturally.
- They demand visibility into how their AI tools work, what data they touch, and how decisions are made.
- They protect their security data from third-party exposure, especially in regulated and air-gapped environments.

The SOC of 2026 is not fully automated. It's not going to be in 2027 either. But the threat landscape is not waiting. Attackers are already deploying agentic AI, systems that probe, adapt, and move at machine speed. A human analyst reviewing alerts at human speed is not a match for that. Neither is a static AI running yesterday's detection logic.

Agentic threats require agentic defense. That means AI that can reason across data in real time, surface what matters before the window closes, and keep a human analyst informed and in control at every step. The organizations building that capability now, on top of data infrastructure that can actually support it, are the ones that will be effective when it matters most.

ABOUT CROGL

Crogl builds an AI-powered Knowledge Engine for enterprise Security Operations Centers. Where other AI tools require security data to be normalized, structured, and pre-processed before they can work, Crogl's patented non-normalization approach connects directly to existing data sources, in whatever format they're in, and surfaces the context analysts need, when they need it.

The result is an agentic defense for SOC analysts: AI that enhances what a skilled analyst can see and do, without removing them from the decisions that matter. Crogl deploys in cloud, on-premises, and fully air-gapped environments, giving security teams full control over their data and their AI.

www.crogl.com



METHODOLOGY

This research was sponsored by Crogl and conducted independently by Ponemon Institute. A sampling frame of 18,602 IT and IT security practitioners in North America were selected as participants. 688 total returns were received; after screening and reliability checks, 649 surveys were included in the final analysis, representing a 3.5% response rate. Respondents were required to be knowledgeable about their organization's security operations and SOC practices. All survey responses were captured in October and November 2025.

Respondents represented a range of organizational levels, with 59% at or above the supervisor level. The largest industry segments represented were financial services (16%), public sector/government (11%), technology and software (11%), and industrial and manufacturing (10%). The largest departmental representation was IT security/InfoSec (21%), followed by IT operations and DevOps/DevSecOps (each at 17%).

Ponemon Institute | research@ponemon.org
© 2026 Crogl, Inc. All rights reserved. Research conducted by Ponemon Institute.

