



Agentic Security Operations Report

Machine-speed defense: Proven results of building the world's
largest agentic commercial SOC

2026



TABLE OF CONTENTS	Foreword: AI Has Permanently Changed the Face of Security	3
	Key Takeaways: The Agentic SOC Is Here	5
	Data Sources and Methodology	8
	Introduction: The Rise of Agentic Security Operations	9
	Detecting and Responding to Threats at Machine Speed	12
	Building an Agentic SOC Worthy of Trust	16
	Scale Means More Than Just Volume	21
	Delivering Value at a Predictable Cost	23
	Conclusion: Machine-Speed Attacks Require Machine-Speed Security Operations	26



Foreword: AI Has Permanently Changed the Face of Security

Security leaders have already decided that AI is worth buying. They have not decided it's worth trusting. In Arctic Wolf's 2026 AI & Cybersecurity Trends Report, 94% of respondents said AI capabilities now factor into their security buying decisions, and 51% said they won't buy a product without them. Ask those same leaders how much they'd let AI do inside their own environment, though, and the answers get a lot more cautious. Only 14% have made AI central to their security operations strategy. That caution has some basis. In the 2026 SANS AI in Cybersecurity Survey, nearly two-thirds of security professionals said they'd received AI-generated guidance that later turned out to be wrong. Buyers want AI in their tools, but most aren't ready to trust it with their environment.

The Gartner® 2026 Hype Cycle™ for Agentic AI puts AI SOC agents at the Peak of Inflated Expectations, with mainstream adoption still two to five years out.¹ Gartner also predicts that "By 2027, over 40% of agentic AI projects will be canceled due to escalating costs, unclear business value or inadequate risk controls."² Attackers aren't hesitating the same way.

Speed makes defense harder, and the federal government has written that into policy. In June 2026, the Cybersecurity and Infrastructure Security Agency (CISA) issued Binding Operational Directive 26-04, which sets a three-day remediation deadline for the highest-risk exploited vulnerabilities. CISA named the reason directly, writing that threat actors' use of AI may further narrow the time defenders have to react between patch release and possible exploitation. Most organizations are nowhere near that pace, and they're moving away from it. Verizon's 2026 Data Breach Investigations Report (DBIR) found that the median time to fix a known exploited vulnerability rose from 32 days to 43 days, while full remediation fell from 38% to 26%. For the

first time in the report's 19-year history, exploited vulnerabilities passed stolen credentials as the most common way attackers get in.

The downstream effect is already measurable. According to IBM's Cost of a Data Breach Report 2026:

- A quarter of malicious incidents now involve AI, a 56% increase over the prior year
- The average breach now costs a record \$4.99 million (USD), a 12% increase in a single year, with breaches involving AI running about \$1 million higher
- Identifying and containing a breach takes 247 days on average, up from 241, which reverses five straight years of improvement

Most change management policies were written for a different era, and they haven't kept up with the pace of AI. AI is forcing companies to rethink them, along with other processes that have barely changed in 40 years. Everything that gets through becomes a reactive problem that drives urgency and response which must be dealt with outside of a plan. In the same IBM study, organizations that used AI security and automation extensively cut 65 days off the breach lifecycle and roughly \$1.93 million off the cost. That's what the Aurora® Agentic SOC is built for. Speed only works if you can trust the outcomes, so our agents act when they're confident in a decision and not before. When they aren't, they escalate to a human analyst rather than guessing. Over the past year, the time from an alert-generating event to an actionable ticket fell 36%. Because our AI agents run around the clock, the 44% of alerts that arrive outside business hours get the same depth of investigation as the ones that land at 10 a.m. on a Tuesday. Attackers don't keep business hours, and neither do we.



In this emerging world of agentic AI, it is not only attackers with access to this revolutionary technology. When implemented strategically, we have seen many security teams achieving significant defensive gains; from AI-enhanced detections and security investigations to security teams utilizing AI to both harden environments proactively and enhance their reactive capabilities.

In this AI-centric landscape, traditional defenses scaled by adding more tools and headcount can't keep pace with AI-augmented attacker tactics, techniques, and procedures (TTPs). At this speed and scale, you need effective agentic security operations to counter both of these risks at once.

To gain the advantage in the AI security arms race, security teams need to restructure into integrated, AI-driven security operations that continuously detect, investigate, and respond. In that model, AI agents are not bolt-on features or hastily integrated assistants but the central engine, executing workflows, correlating signals, and driving response actions across the environment with human-in-the-loop oversight. Routine work that once consumed costly analyst hours can now run autonomously. Specialized AI agents now handle many of the jobs of the SOC, escalating only the complex investigations that require human experts. When that happens, AI agents provide valuable, pre-enriched context and evidence that previously took traditional analysts hours to assemble by hand.

The Aurora Agentic SOC has also shown us that the role of humans is one that continues to evolve in beneficial ways. Human analysts remain a core part of operations, but they are no longer burdened with manually processing alerts or stitching together data from disparate systems. Freed from these traditional time-consuming processes, they act as the strategic lever in tandem with AI agents, interpreting outcomes, taking decisive action when moments matter, and applying the experience and judgment that novel attacks demand.

Our journey of human-AI interdependency was a design we began initially implementing over two years ago. Today, it powers the world's largest commercial SOC — Arctic Wolf's — and has ushered in a new era of agentic security operations at a scale the industry has never seen. This report presents a framework for customers to quickly adapt to the modern AI-accelerated threat landscape. We detail the key capabilities, features, and outcomes needed to build business resiliency and trust. AI is changing the economics of cyber attacks, and security teams are under pressure to detect, investigate, and respond faster than ever. But transforming security operations takes more than adding AI. The real question is whether AI can deliver better outcomes at the speed and scale of modern attacks without sacrificing trust or control.

This report answers that question with evidence from Arctic Wolf's global agentic SOC, examining the four factors that determine whether agentic AI delivers real security outcomes: speed, trust, scale, and value. For security leaders deciding where and how to deploy it, this report offers a practical look at what the technology delivers today and what to expect from the next generation of security operations.



LAURA ELLIS
Senior Vice President,
Artificial Intelligence
Arctic Wolf



LISA TETRAULT
Senior Vice President,
Security Services
Arctic Wolf



Key Takeaways: The Agentic SOC Is Here

To meet the demands of today's AI-powered threat landscape, the security operations center (SOC) must adapt, completely transforming with AI at its very core. Three numbers define the problem and the fix in the same breath.

⬆️ 21%

Monthly investigation volume grew 21% in the eight months after the Aurora® Agentic SOC launched, as AI-powered threats and AI-generated noise from security tools keep flooding the pipeline with more to examine.

⬇️ 19%


Yet, alert volume fell roughly 19% over the same period, despite record data volume. This figure shows that the agentic SOC is cutting through that noise, and getting sharper at recognizing which signals are genuine threats. The work that remains is real investigations requiring human expertise, not time consuming noise.

⬇️ 36%

The average time to ticket fell 36%, as AI delivers not just faster detection, but stronger guidance leading to better-informed decisions, helping customers move from alert to resolution faster than ever.

Rather than more visibility translating to more alerts, our AI agents are designed to ensure the noise gets filtered before it reaches a person. Arctic Wolf built and operates the Aurora Agentic SOC, the world's largest agentic commercial SOC, with AI agents and human experts working together across every stage of detection and response. The figure to the right shows the results of this AI-native architecture and how we achieved them.

Arctic Wolf Agentic SOC Investigation Resolutions



FINDING

The Aurora Agentic SOC Escalated 13 Investigations for Every 1 Billion Observations Ingested

Stripping benign activity and false positives at machine speed, so analysts see only what matters.

IMPACT

Only 0.0000013% of raw data observations ever become an investigation. That's the equivalent of watching nearly 32 years of security camera footage to catch one 13-second theft. At that ratio, precision isn't optional, it's existential. Miss the signal and an attack goes unnoticed. Chase every anomaly and your reaction time slows before the real threat surfaces. That's the problem the Aurora Agentic SOC is built to handle, filtering at that scale so what reaches an analyst is worth their time, and nothing that matters gets lost in the noise.



FINDING

IMPACT

60% of Security Investigations Resolved by Agentic AI Alone

The majority of our investigations are researched and resolved by agentic AI, requiring no human intervention.

Customers see faster resolution while their analysts' get focused time back for what actually matters. With over 200,000 investigations resolved every week, cases move up to 15x faster to resolution, and internal security teams can stop triaging routine alerts and start focusing on the threats that need human judgment.

The Human-AI Partnership Takes Center Stage

More than one-third of complex security investigations benefit from final human verification.

38% of cases now get AI-supplemented investigation with a final human check before resolution. On these complicated cases, human experts verify the answer before anything escalates to the organization. Customers get AI speed with analyst expert judgment when it's needed the most.

Drastic Reduction of Human-Only Investigations

Due to the enhanced speed and benefits of AI-supplemented investigations, we saw a 98% reduction in human-only investigations.

During this research period, the agentic SOC allowed us to transform from one where human analysts investigated nearly 81% of cases without AI. Today, through expert reinforcement and the continual enhancements we've made, that amount is now down to 2%. That's a ~98% drop. For customers, that means fewer cases need to queue for a human at all. Machine-speed response doesn't depend on case volume or time of day, and an attacker can get shut down just as fast at 2 a.m. on a Sunday as on a Tuesday morning.

3x Increase in Total Agentic Investigation Resolution

We observed a more than 200% increase in security investigations that are now fully resolved by agentic AI agents.

Between September 2025 and September 2026, agentic AI's share of resolved investigations more than tripled. Real-world data, continuous human oversight, and real-time tuning have compounded over the course of a year. For customers, their environment got three times better at stopping attacks autonomously without the need to buy anything additional, add headcount, or touch a settings page.

The Average Time to Ticket Fell 36%

Highlighting the speed at which AI-supported investigations turn raw alerts into actionable tickets.

The time from an alert-generating event to an actionable ticket dropped by 36%. This is closing the gap where damage compounds: an attacker moving laterally, compromised accounts being used to reach more systems, or data leaking before anyone's caught it. Closing more of that gap doesn't just mean a faster dashboard; it means less time for an attacker to turn initial access into damaging incidents.



FINDING

IMPACT

Fewer Alerts, Yet More Investigations: Alert Volume Fell ~19%

Fewer alerts in a threat landscape where attacks are continually increasing thanks to AI-enhanced true-positive detection capabilities.

Our detections get sharper as the threat landscape grows more complex. Enhanced detection logic and AI-driven refinement now filter out false positives before they ever become alerts. At the same time, monthly investigations grew 21%, as human-AI efforts shifted away from dismissing noise and toward examining genuine threat activity. For customers, fewer alerts mean less noise and alert fatigue. More investigations mean more real activity getting a second look rather than chasing false leads. Fewer alerts, real threats, caught faster.

Alert Volume Trends Downward Despite Increased Mountains of Data “Noise”

One alert surfaced for every 175.6 million observations in FY26.

Despite the continued growth of telemetry being ingested, the implementation of the agentic SOC has inversely resulted in fewer alerts generated year over year. Organizations can now expand visibility without expanding alert fatigue.

44% of Alerts Arrive Outside of Traditional Working Hours

AI agents work 24x7 and scale up and down on demand.

Attackers' use of the traditional off-hours advantage is still common, but less dominant. It dropped from 51% last year to 44% this year, as more of them shift toward techniques that blend in with legitimate user behavior during the workday. Whether it's an attacker's hope that defenses run thinner at 3 a.m., or that unusual activity gets lost in the noise of a busy Tuesday, the Aurora Agentic SOC runs at the same depth of investigation 24x7. There's no schedule or noise to exploit.

Properly Configured Endpoint Defense Pays Off

AI-enhanced endpoint defense detections prevented the execution of 250 threats per Aurora Endpoint Defense customer in FY26.

The Aurora Agentic SOC's AI-enhanced detection and response is increasing the success of finding and eliminating threats across customer environments at machine speed. In FY26, customers with Aurora Endpoint Defense saw an average of 250 threats per customer prevented from executing.

Agentic SOC Services Are Effective and Cost Efficient

Offering a roughly 12x cost advantage over building your own agentic SOC.

Building an agentic SOC can be a substantial undertaking full of hidden costs, from initial purchase to the cost of tuning, managing, and modifying models. The Aurora Agentic SOC delivers customers enterprise-grade agentic capability at a fraction of the cost of building and maintaining their own in-house program.



Data Sources and Methodology

While this report cites other Arctic Wolf publications and third-party sources for context, the majority of facts, figures, and statistics presented here come from the Aurora Agentic SOC, drawn from our customer base of 11,000+ organizations in 100+ countries.

More than 10 trillion telemetry events flow through the Aurora Superintelligence Platform each week. We utilized this telemetry flow's metadata for this report, which focuses on comparisons from the launch of the Aurora Agentic SOC in August 2025 through September 2026.

Alongside internal telemetry, we draw on anonymized customer investigation outcomes, performance metrics from agentic operations, and assessments from our security analysts on the effectiveness of the Aurora Agentic SOC.

Key Components and Terms

Aurora Agentic SOC	The world's largest commercial agentic SOC, delivering trusted outcomes with turnkey deployment.
Aurora Superintelligence Platform	Combines real-world security data and defense expertise compiled over more than 14 years with specialized AI agents and 24x7 humans in the loop to deliver the speed, context, and trust needed to outperform both AI-only and human-only approaches.
Security Operations Graph™	The Platform's proprietary data and intelligence foundation. It connects more than 10 trillion telemetry events each week with a large body of validated security outcomes, customer-specific business context, and case memory developed by operating the largest commercial agentic SOC in the world.
Swarm of Experts™	Arctic Wolf's patent-pending architecture that coordinates specialized AI agents working in tandem to investigate, analyze, validate, and resolve security activity, with humans in and on the loop 24x7.
AI Trust Engine™	A validation framework that governs how our agents behave, checks their decisions, bounds their autonomy, and brings in a human at the right moment rather than after the fact.
Concierge Experience™	Gives customers triaged alerts that reduce fatigue, continuous tailored guidance that strengthens their security program, and on-demand access to named security experts who know their environment.
Open Data Pipeline™	Collects telemetry across an organization's entire attack surface, enriches it into high-fidelity signals, and builds customer-specific context.
Observation	A single security-relevant event drawn from customer telemetry. The raw material of everything downstream.
Investigation	A case the Aurora Agentic SOC opens to determine whether observed activity is malicious. Used interchangeably with "case" throughout this report.
Alert	A validated finding delivered to a customer with analysis, business context, and recommended action, including a record of any response already taken.



Introduction: The Rise of Agentic Security Operations

KEY TAKEAWAYS

Security operations need to be rebuilt for a threat environment increasingly shaped by automated and AI-powered attacks moving at machine speed

Visibility across the IT environment remains the precondition for everything else, since any gap in telemetry is both an opening for attackers and a blind spot in an investigation

An agentic SOC has to deliver speed, trust, scale, and value to be a practical alternative to human-led security operations

For the last few decades, security operations ran on a simple model: Tools collected telemetry and generated alerts, and human analysts reviewed those alerts and decided what to do. As alert volumes grew, teams added platforms like SIEM to pull data together and cut down the noise. Those tools helped, but they were layered onto a process that was still built around people doing most of the work by hand.

Even at its best, that pipeline is slow, error-prone, and a reliable source of analyst burnout. Worse, any meaningful increase in alert volume, whether from a shifting threat landscape, an expanded environment, or a new telemetry source, demands a proportional increase in staffing.

That model no longer works.

A Rebuilt SecOps Pipeline

The agentic SOC is not an incremental iteration of the traditional model. It is a rebuild from the ground up, designed so defensive capacity can grow beyond the size of the team once needed to deliver it. The path is more demanding, but it is the one modern threats now require.

The speed, scale, and sophistication of automated and AI-powered threats require autonomous agents to operationalize much of the security operations pipeline, with human-led manual review reserved for quality control of the agent output and novel or especially complex attacks.

Arctic Wolf's Swarm of Experts is an AI-led agentic framework with humans in the loop. It is designed to plan and execute actions for many SOC functions, rather than being narrowly focused on individual tasks like triage or investigation. The Swarm consists of hundreds of agents that handle tasks end-to-end and is constantly growing and adapting with humans in the loop. This model allows seasoned analysts to serve as expert validation points for incident escalations and to continually advance model performance by providing ongoing expertise derived from serving over 11,000 global customers.

The diagram below shows how a single case moves through the Swarm of Experts, from initial signal to validated response. It is a process view rather than an architectural one, and it maps the agentic steps described throughout this report.



Reading it from left to right:

- Event ingestion and enrichment, along with behavioral analysis and correlation, run autonomously.
- Autonomous agents then gather context, validate threats, and prepare recommendations, and in many cases initiate some response actions themselves, such as closing the investigation or creating a ticket. In a time when 44% of alerts arrive outside of traditional working hours, agents work 24x7 and scale up and down on demand.
- Human review and investigation are reserved for the highly complex, genuinely novel, and/or the high-risk cases.

Swarm of Experts™

STEP 01

Signal Ingestion

A case begins as the open data pipeline ingests security data, surfaces a potential threat, and creates the foundation for investigation.

STEP 02

Swarm Orchestrator

The Swarm Orchestrator breaks the case into tasks, assigns the right agents, and directs the swarm to move in parallel.

STEP 03

Agents Investigate

Specialized agents work the case, assembling context and evidence to recommend the path towards resolution.

STEP 04

Validate Response

The Swarm Judge validates the suggested outcome. Human experts step in when the case calls for added judgment, context, or oversight.

STEP 05

Continuous Reinforcement

The validated response feeds back into the Security Operations Graph, strengthening performance of the Swarm of Experts over time.

Beyond the performance gains, this structure changes where human expertise goes. Analysts who are not clearing queues can hunt, harden, and advise, which is how prevention finally gets the attention it has always deserved.



Visibility Remains the Foundation

To improve efficiency and speed, Arctic Wolf rebuilt the ingestion pipeline around autonomous agents. No agent can act on what it cannot see. Broad, deep visibility across the attack surface remains essential.

Modern attacks rarely stay on a single surface. An intrusion may begin with a compromised identity, move through cloud applications, and reach endpoints before its objective becomes clear. Detecting and responding accurately requires telemetry from endpoints, mobile devices, networks, clouds, applications, identity and access management (IAM) services, infrastructure components, human signals, plus any additional log sources and third-party alerts available.

Four Principles Define a Trusted Security Outcome

Much of the cybersecurity industry still focuses its core efforts on detection speed alone. We built the modern agentic SOC around four principles that together define a trusted security outcome, and the chapters that follow take them in order.





Detecting and Responding to Threats at Machine Speed

KEY TAKEAWAYS	Attacks execute faster than ever, because threat actors have added AI and automation to a proven playbook
	Defenders need more than a notification that something malicious happened; they need context, a recommendation, and a response tailored to their organization
	60% of investigations are automatically researched and resolved by agentic AI with no human intervention
	Human-only investigations fell from roughly 81% of the caseload in the first full month after launch to 2% one year later, while AI-human partnership investigations grew to 38%
	The average time to open a ticket fell 36% while the median fell 24%, which means the difficult cases improved the most

Fast detection and response is a security control in its own right, since stopping an attack in its earliest stages limits lateral movement and contains the blast radius.

Threat actors know this, which is why they have now combined AI and automation to compress their attack timelines. Human-led operations cannot match that pace, so agents working at machine speed have to take the lead.

Speed Without Context Is Ineffective

Detecting something is only half the job. A fast alert with no context just hands the problem to a person, who still must figure out what it means and what to do about it. Mean time to detect (MTTD) looks better, but mean time to respond (MTTR) gets worse in practice. That trade-off favors the attacker.

Whether the defender is a human analyst or an AI agent, the requirement is the same: Understand what is happening, understand what matters about it, and know how to act on it within the specific business and operating context of the organization.

That is the standard we hold our own operations to. An Arctic Wolf alert explains what happened, what it means, and exactly what to do next. Where we have preapproval, it arrives with the response already taken.



The Handoff: AI Speed Meets Human Expertise

Since the first agentic elements of the Aurora Agentic SOC went live in August 2025, the distribution of work has evolved significantly.

At that time, humans handled roughly 81% of our caseload without AI assistance. A year later, that figure has now fallen to 2%. Over the same period, investigations resolved entirely by agentic AI have tripled to 60% of total volume, while investigations run as an AI-human partnership grew from less than 1% to 38%.

The key element here is that AI has not replaced the analyst. Nearly 40% of all investigations still require a human, either alone or in partnership with AI, to handle unseen or ambiguous threats. AI acts autonomously where it can be trusted to, and defers to expert analysts where it lacks the knowledge to decide with confidence. Agentic AI absorbs the volume and toil, so humans can focus their expertise on the novel threats that demand judgment.

For customers, this means routine threats are resolved at machine speed, while the ambiguous, high-stakes cases remain the focus of our experienced analysts.

When Agents Act

Many cases can be closed by AI agents without the need for human escalation, including cases finding that the detected activity was never malicious. Recurring examples include:

- ✓ Expected administrative actions flagged as anomalies
- ✓ Failed authentication attempts from recognized locations
- ✓ Routine network and endpoint behavior matching known safe patterns
- ✓ Scheduled maintenance activity classified incorrectly

In other cases, agents are able to confirm malicious intent, and response actions follow quickly. Such response actions may include:

- ✓ Containing or isolating a device
- ✓ Pulling a malicious email back out of every mailbox it reached
- ✓ Disabling or suspending a user account
- ✓ Blocking reconnaissance activity and payload transfers at the network layer
- ✓ Forcing a credential reset



Where the Partnership Continues

The 38% of investigations that involve both agents and humans are, with few exceptions, the cases where complexity is a defining feature. The situation may be novel, new, highly complex, or involve conflicting signals. In any case, an attack technique that is recent enough to be absent from years of historical operational data has no precedent for an agent to match against, and precedent is what agents are good at.

In these investigations, agents gather evidence, correlate signals, and generate hypotheses before surfacing the case to a human analyst who validates the findings, assesses business impact, and approves the response. The categories that consistently land here include:

- Multistage attacks whose connections are not obvious from any single signal
- Patterns consistent with nation-state or geopolitically motivated threat actors
- Attack techniques novel enough to fall outside our validated experience
- Complex privilege escalation chains
- Activity whose severity depends entirely on sensitive business or customer context

This is what bounded autonomy looks like in practice. When an agent reaches the edge of its validated experience, it does not guess; it escalates.

Increased Speed Across All Case Levels

The average time to open a ticket fell 36%, with our research finding the largest improvements in speed coming from the most complex cases. Traditional automation in a SOC usually does the opposite, clearing simple work quickly, while difficult cases take just as long as they always do. The partnership between humans and AI corrects this, with some workflows now showing resolution up to 15x faster than under the traditional model.

The Off-Hours Threat Actor Advantage Is Shrinking

Roughly 44% of all alerts arrive outside standard business hours in the customer's local time zone, down seven percentage points year over year. Broken out by region as of mid-September 2026, the pattern holds at roughly 46% in North America; 38% in Europe, the Middle East, and Africa (EMEA); and 42% in the Asia-Pacific region (APAC).

Historically, threat actors favored nights and weekends because off-hours activity offered dwell time before anyone noticed. That calculation assumed traditional human defenders kept regular office hours before going home. As AI accelerates attacks and improves defense, the timing advantage matters less than it did, because agents do not keep office hours, and the response is already underway before the on-call phone rings.



MEGA EVENT SPOTLIGHT:

Payroll Pirates – One Signal, Thousands Protected

In July 2026, a phishing operation calling itself Payroll Pirates moved through hundreds of organizations across healthcare, education, manufacturing, government, and professional services in the United States, Canada, and Europe. The campaign's entry point was a voicemail-themed phishing email that routed victims through a multistage redirect chain to an adversary-in-the-middle proxy, one that relayed Microsoft's real authentication flow in real time and intercepted the session token behind it, defeating multi-factor authentication (MFA) without ever needing to break it. Once inside a mailbox, the operation went quiet. No password resets. No new devices. No forwarding rules, in most cases. Just a script maintaining the stolen session while it used Microsoft Graph to enumerate anyone in payroll, HR, or finance. The goal was narrow and specific, redirecting employee salaries to accounts that the attackers controlled.

The pattern started with a single flagged account. An Arctic Wolf analyst caught an account pulling up a long list of HR and payroll employees: an action that looked unremarkable on its own but did not fit the account's normal behavior. Pulled together, the surrounding signals told a consistent story:

- Sign-ins on Firefox where the account had only ever used Edge
- A login error tied to a missing OAuth parameter that was exceedingly rare across Arctic Wolf's global Aurora Managed Detection and Response (MDR) telemetry outside this campaign
- The same session ID persisting across residential proxy addresses that rotated every eight hours
- Mailbox-access events clustering across separate tenants within seconds of each other; in one case 18 events across three tenants inside a 26-second window.

None of those signals alone would trigger a response, but together they were the campaign's signature.

Arctic Wolf conducted a global threat hunt across thousands of customer environments, confirmed the indicators, coordinated containment and response actions for hundreds of compromised accounts, and turned the pattern into six new detections. Every impacted customer received them, along with hardening guidance, before the campaign breached further in their environment. For the rest of our customers, one flagged mailbox became coverage and protection the attackers hadn't gotten to yet.

**CARL ADASA**

Vice President of Security Operations
Arctic Wolf



Building an Agentic SOC Worthy of Trust

KEY TAKEAWAYS

Security leaders are right to be cautious about agentic systems, especially when it isn't clear how those systems reach their decisions

Trust is the organizing constraint of the Aurora Agentic SOC, which pairs agentic systems with human analysts to use the strengths of each

The Security Operations Graph grounds agent decisions in 14+ years of real-world security operations, a large body of validated outcomes, and context specific to each customer

The AI Trust Engine bounds agent autonomy, keeps agents deterministic, checks every decision, and routes to a human expert when judgment is required

The 2% handled by a human alone are the cases we escalate immediately because the attack is novel, new, or complex enough to warrant immediate attention from expert analysts

In Arctic Wolf's 2026 AI & Cybersecurity Trends Report, built on a global survey of 1,350 security and IT decision-makers, 94% of respondents said AI capabilities influenced their cybersecurity purchasing decisions and 51% called AI functionality a requirement when evaluating vendors. The same survey found additional results, however, that seem contradictory:

- Only 14% of these leaders have made AI central to their security operations strategy
- Only 53% said they trusted AI to autonomously perform even a narrowly defined security action such as blocking a malicious IP address at a firewall

Compare that with the 2026 SANS AI in Cybersecurity Survey, in which nearly two-thirds of security professionals reported receiving AI-generated guidance they later determined was wrong, and the hesitation looks less like conservatism and more like an accurate read of the market.

Trust, not capability, is often the barrier to agentic security operations. The rest of this chapter is an account of how ours is built.

"Security leaders want the speed, scale, and efficiency AI can deliver, but they also need confidence that the outcomes are trustworthy, explainable, and aligned to their business. The future of security operations will belong to organizations that combine trusted AI with human expertise to move faster, make better decisions, and build greater resilience."



ADAM MARRE
Chief Information
Security Officer,
Arctic Wolf



Human and AI Collaboration: Complementary Strengths

The most effective security operations are neither agent-only nor human-only, but a deliberately designed division of labor.



Where agentic systems excel

- Processing millions of events instantly
- Applying rules consistently
- 24x7, tireless, deep investigations
- Routine triage and escalation
- Pattern matching across years of data



Where human analysts excel

- Recognizing genuinely new or novel attack patterns
- Contextual judgment calls
- Strategic threat analysis
- Prioritizing by business impact
- Adapting to emerging threats

In the rebuilt alert-to-investigation pipeline, agents turn raw signal into an investigation-ready package that supports a fast decision, whether that decision belongs to an authorized agent or to a person.

- Autonomous systems flag the potential event
- Autonomous agents gather historical, behavioral, and environmental context
- Scoring assesses the likelihood of malicious activity
- Routing directs the case to the right specialist agent for endpoint, network, cloud, or identity

What reaches a human analyst is a fully contextualized package rather than a raw alert.

Unlike traditional systems that generate high volumes of low-fidelity alerts, this architecture applies business context during autonomous triage to suppress benign activity, so only high-confidence findings reach an analyst. In last year's Security Operations Report we noted that 71% of all ingested alerts were suppressed by applying customer context and threat intelligence to identify expected or benign behavior.

Every validated response then feeds back into the Security Operations Graph, which is how the Swarm of Experts gets better at the thing it just did.

When We Pick Up the Phone

The 2% of investigations a human handles directly are cases where immediate action and human reasoning is required.

Escalation of that kind is a design requirement rather than a failure state. An agentic SOC that cannot recognize the cases it should not handle alone is not trustworthy, regardless of how well it performs on the other 98%.



The Foundation: The Security Operations Graph

The Security Operations Graph is the proprietary data and intelligence foundation inside the Aurora Superintelligence Platform, where the agents run. To make a trustworthy decision possible, it equips every agent and workflow with three things:

01

Real-world experience

More than 14 years of operating the world's largest commercial SOC across 11,000+ customer environments gives the platform a foundation of telemetry and threat signals that no synthetic dataset reproduces.

02

Known outcomes at volume

The Graph holds an enormous body of security events whose real resolution is already established, which is the ground truth required to build agents that clear rigorous performance benchmarks, and to know when they have not.

03

Customer context and memory

Through the Concierge Experience, knowledge of each customer's environment, business, and history are designed to ensure agents produce outcomes that are relevant to that organization rather than generically correct.

None of these are static. The Security Operations Graph updates continuously from observations and validated responses across our global customer base.

Specialized Agents: The Swarm of Experts

At the core of the Aurora Agentic SOC, the Swarm of Experts coordinates hundreds of specialized AI agents that execute and validate security operations work in parallel. Coverage spans core SOC functions: triage, investigation, response, threat hunting, threat intelligence, detection engineering, and the application of customer-specific context.

Oversight Agents coordinate that work and review its output, keeping it controlled and aligned to defined standards while keeping human experts central to the critical decisions. They take two forms:

01

Swarm Orchestrator

Directs all work inside the Swarm of Experts and is designed to ensure SOC functions, including calling a human if needed, operates in a coordinated way.

02

Swarm Judge

Reviews outcomes and decisions across the Swarm, validating results against defined standards to reinforce quality, consistency, and trust.

Beneath them, hundreds of Process Agents support the domain specialists with focused tasks, operating in parallel at machine speed to gather context, enrich signals, execute discrete investigative steps, and feed intelligence upward. Each specialist agent has an intentionally bounded scope, which is what makes precise action and reliable escalation possible.



Building Trust into Every Outcome: The AI Trust Engine

The AI Trust Engine™ is the quality framework behind the Aurora Agentic SOC, built on the Aurora Superintelligence Platform. It is a closed loop of five checkpoints that every agent and every change must pass through.

- ✓ **Golden datasets:** Curated, expert-validated investigations, drawn from millions of resolved cases, set the benchmark for every model, prompt, workflow, and agent.
- ✓ **Offline evaluation:** Every change is tested against the golden datasets before it reaches production and is scored on verdict accuracy, speed, and cost. Changes that fall short of the quality threshold do not move forward.
- ✓ **Response grading:** A separate judge model scores production responses against a fixed rubric, checking for hallucinations, reasoning consistency, and the quality of summaries and recommendations. Humans validate the rubric and audit the judge.
- ✓ **Earned autonomy in production:** Agents enter the Swarm of Experts in stages. They start in silent mode, then recommend actions to analysts, and only once validated can they then act on their own under ongoing human audit. Low-confidence cases always escalate to a human.
- ✓ **Continual feedback:** Analysts flag or correct agent decisions directly in their tools, and that feedback is added to the golden datasets along with sampled spot checks.

When a case falls outside an agent's validated experience, the agent then escalates to an analyst instead of guessing. Each human decision becomes a new tuning signal that expands what agents can handle on their own. This loop drives the growth in autonomous resolution described earlier.

Trust compounds with context. Through the Concierge Experience, named Concierge Security Team experts carry deep knowledge of each customer's environment, history, and risk tolerance, and the Aurora Agentic SOC bakes that same context into every investigation, not only the checkpoints above.



AGENTIC SOC SPOTLIGHT:

Why Humans Will Always Be Essential

Not every case begins with an attacker. One recent example began with a third party's AI-generated assessment that proved to be inaccurate.

In this example, an organization chose to outsource its IT environment to an external IT consultancy, which used AI to assess the environment. The AI tool concluded that the customer had been compromised. This AI tool reached that verdict from a thin set of indicators and publicly available reporting, while presenting the conclusion without qualification. The consultancy then accepted this verdict without independent validation.

The Aurora Agentic SOC additionally investigated and found no evidence of compromise. The consultancy disagreed with the assessment and requested further investigation. Upon that request, the case was escalated to the Incident Response team for additional verification. Arctic Wolf Incident Response (IR) forensic analysis, configuration reviews, and artifact validation independently confirmed the Aurora Agentic SOC findings, ultimately reaching the same conclusion; the environment had not been compromised. Throughout the engagement, the consultancy relied on AI-generated assessments that were not substantiated by the available evidence. Following a comprehensive review, including validation by an independent third party, the consultancy aligned with our conclusions.

The financial costs to the organization were real, and none were due to an actual breach. The result was a full investigation, an incident response engagement, and a third-party review; all spent disproving something that never happened. The customer took the findings to its board as a governance matter.

This is the case for AI abilities partnering with human judgment. An AI system produced a specific, serious, and incorrect conclusion, then presented it with complete confidence. What corrected it was human domain expertise, evidence, and the willingness to contradict a machine. That same discipline governs how we build and operate our agentic systems, where AI output is evidence for a human to validate rather than a verdict to accept.



MARK THOMAS

Regional Vice President of APAC Security Services
Arctic Wolf



Scale Means More Than Just Volume

KEY TAKEAWAYS

Handling telemetry volume is only one dimension of scale, and doing it well means balancing increased telemetry without dramatic increases in alert volume as a result of increased false positives

In FY26, 13 investigations arose from every 1 billion observations, and one alert from every 175.6 million

Monthly investigation volume grew 21% while alert volume fell 19%, and the share of high- and critical-severity tickets held steady at 4.63% in FY26 against 4.67% in FY27 to date, meaning we cut noise rather than signal

As of September 2026, our Open Data Pipeline uses 250+ out-of-the-box integrations, a vendor-neutral approach, and predictable pricing to remove the incentive to see less

The Security Operations Graph creates a rising tide effect, where a novel threat against one of 11,000+ customers becomes protection for all of them

Organizations face an unprecedented volume of threats. However, the more useful fact sits underneath it: Benign, everyday business activity drowns out threat signal by many orders of magnitude, and processing that volume is only the first requirement of scale.

From Noise to Signal to Alert

Finding threats without exhausting a security team on false alarms means achieving two things that pull against each other:



Increased visibility through a high volume of telemetry ingestion



Low rate of false positive alerting that does not spike with increased ingestion

The Aurora Agentic SOC parses, enriches, and analyzes the telemetry stream to raise the signal-to-noise ratio before the Swarm of Experts begins investigating.

In FY26, 0.0000013% of observations became an investigation. Put another way, the Aurora Agentic SOC opens 13 investigations for every 1 billion observations it processes, and surfaces one alert for every 175.6 million. For customers, that means the noise stops with us. Their teams are not buried in raw telemetry or false positives, and when an alert does arrive, it carries enough weight to act on.

High- and critical-severity tickets held steady, moving only from 4.63% in FY26 to 4.67% in FY27 to date, while alert volume fell 19%. That means fewer alerts with the same share of serious findings. Comparing month to month since launch, monthly investigation volume grew 21% and alert volume fell 19%, which dropped the share of investigations escalating into an alert by 13 percentage points, from 40% to 27%. More noise is now resolved before it ever reaches an Arctic Wolf analyst.



Each investigation, whether escalated or not, gets added to our golden datasets alongside the tens of millions of other cases we have resolved over more than 14 years conducting security operations. As these golden datasets grow broader and more representative, every future improvement gets measured against a tougher and more realistic bar. The system improves instead of drifting.

Scaling Telemetry: The Open Data Pipeline

Visibility gaps create openings for attackers and strip the investigation pipeline of the context it needs, which is why telemetry should come from every corner of the environment rather than the corners that are convenient to instrument.

Three problems reliably hamper homegrown security operations programs:

01

Visibility that stops partway, leaving gaps that are opportunities for an attacker and blind spots in an investigation

02

Security tools that do not share data with one another, or do not share it in a format that can be correlated with anything else

03

Pricing based on data volume, which creates a direct financial incentive to collect less telemetry than the environment warrants

Every one of those is a decision to see less, and only one of them is a technical problem.

We have long championed a vendor-neutral approach and volume-agnostic pricing, because the alternative asks a customer to choose between cost control and coverage. Our Open Data Pipeline uses 250+ out-of-the-box integrations, which is what makes a turnkey deployment possible in as little as 10 days and lets customers get more out of the security tooling they already own.

Scaling Global Visibility: The Rising Tide

The Security Operations Graph holds more than 14 years of operational data gained firsthand by running the world's largest commercial agentic SOC, which gives the Swarm of Experts historical knowledge of attacks and threat actor TTPs that cannot be bought.

Because it updates continuously from observations and validated responses, that knowledge compounds in a specific way. When any one of our 11,000+ customers encounters a novel threat, every other customer benefits from what we learned resolving it. A large enterprise with the resources to build and staff its own SOC still sees only its own environment, and closes that gap by subscribing to third-party intelligence feeds that describe threats after someone else has already met them.

The same logic runs in the other direction, toward the attack surface itself. The Security Operations Graph can only reason about assets and exposures that were reported, so the organizations that get the most from it are the ones whose visibility extends to the unmanaged, the unscanned, and the forgotten. Knowing what you have is not a separate discipline from defending it.



Delivering Value at a Predictable Cost

KEY TAKEAWAYS

The average customer receives less than one security alert per day, and it arrives with analysis, a recommendation, and a record of any response already taken

Between January and August 2026, roughly 1.3 million manual reviews were eliminated through automation, including agentic workflows, a 51% increase over the prior year

As of mid-September 2026, Aurora Endpoint Security has stopped an average of 250 threats per customer this year, pre-execution

Predictable pricing and unlimited investigations remove the trade-off between saving tokens and seeing threats

We've done the math, and we invite you to do it yourself with our Agentic SOC Calculator. The Aurora Agentic SOC carries a roughly 12x cost advantage over building and operating an agentic SOC internally.

The following is an example of one of our customers' alert management funnels, where 38 billion observations from their environment in a 12-month window were orchestrated down to a manageable workload for their security team. Comparing against the equivalent costs required for an in-house DIY SIEM or MSSP approach, the customer is receiving \$229,000 of value in the alert management process alone:

SIEM configuration and management costs	\$15,000
Agentic tooling configuration and management costs	\$200,000
Remaining triage, alerting, and deep-dive investigation costs	\$14,000





Now considering the 114 TB annual ingestion volume for this same customer, the equivalent value of our unlimited ingestion model under the typical usage model is a substantial \$502,000. These costs are made up of the following equivalents:

Log Source Ingestion	\$128,500
Agentic Tooling	\$92,000 (before additional token charges)
Analytics Logs	\$81,000
Log Retention	\$3,500
Endpoint Tooling	\$44,000
Threat Intelligence Feeds	\$60,000
Attack Surface Management Consolidation and Reporting	\$93,000

Finally, adding the annual equivalent costs of providing 24x7 coverage with proactive security hardening and ongoing expertise from our concierge and IR teams of \$572,000, it becomes apparent that the operational value from the Aurora Superintelligence Platform and Aurora Agentic SOC adds up to a sum well-beyond what most customers are in a position to pay to build equivalent in-house agentic security operations on their own.

Detection and Response: The Direct Outcomes

The Aurora Agentic SOC has resolved more than three million security cases to date, and AI-driven triage processes over 200,000 investigations every week.

For the average customer, that arithmetic produces:

<1

Less than one security alert per day, complete with analysis and recommended action, rather than an inundation of noise that burns out an IT team

~250

Roughly 250 threats per customer stopped pre-execution by Aurora Endpoint Security, because endpoint enforcement remains one of the most effective places to stop reconnaissance, lateral movement, and malware execution

98%

A 98% catch rate against penetration testing through Aurora Managed Detection and Response

7x

A 7x lower likelihood of a major incident than with endpoint detection and response (EDR) alone




In the rare cases where incident response (IR) is required, the process runs faster and cleaner because the agents did the pre-investigation work. The investigation documentation generates automatically, which supports the evidence requests that arrive from regulators, cyber insurers, and other third parties at the worst possible moment.

Through the Concierge Experience, all of it comes with named human experts who know the environment, so customers make continuous progress against cyber risk without absorbing operational burden.

While operational savings are of utmost importance to customers' budgets, it is the reduction of risk that ultimately determines whether the investment is worth it. Using the same example customer as above, the following is a conservatively projected annualized risk reduction by combining Arctic Wolf's agentic solutions as well as the amount of risk transfer that comes with the investment.

How Arctic Wolf Solutions Reduce Sample Customer's Likelihood & Impact of a Cyber Attack



AI has made threats faster but it has not changed the underlying value of mitigating those threats, as long as the solutions protecting against them stay just as effective. The Aurora Superintelligence Platform handles detection and the Aurora Agentic SOC handles response, with Arctic Wolf managing the orchestration burden between them. That combination is designed to give organizations the same level of risk reduction and transfer they have always received, even in the AI era.

Prevention Takes Center Stage

The most consequential number in this report is the one that describes work that stopped happening. Between January and August 2026, roughly 1.3 million manual reviews were eliminated through automation, including agentic workflows, a 51% increase over the prior year.

Every manual review that no longer needs a person is analyst time returned to the work that prevents incidents rather than processes them: threat hunting, detection engineering, posture assessment, and the unglamorous configuration review that quietly keeps an intrusion from starting.



Conclusion: Machine-Speed Attacks Require Machine-Speed Security Operations

KEY TAKEAWAYS

Speed: Machine-speed attacks demand machine-speed defense. Agentic AI now fully resolves over 60% of investigations, while human-only investigations fell from 81% to 2% in a single year.

Trust: Humans remain essential. Nearly 40% of investigations still involve an analyst's expert judgment, either working alone or in partnership with AI, and agents escalate rather than guess when a case falls outside their validated experience.

Scale: More visibility no longer means more noise. Monthly investigations grew 21% while alert volume fell 19%, with the share of serious findings holding steady.

Value: Building an agentic SOC in-house is costly, and consumption-based pricing pushes organizations to see less. The Aurora Agentic SOC delivers a roughly 12x cost advantage compared to building one internally.

Urgency: Adversaries have already adopted AI. The decision to modernize security operations belongs to this year, not next.

On top of a proven arsenal of TTPs, today's threat actors use AI and automation to find vulnerabilities, assemble novel attack chains, execute at machine speed, and target the AI assistants now embedded throughout the software stack.

You cannot fight machine-speed threats with human-led operations alone. You also cannot take the humans out, because expertise, judgment, and accountability are what make an outcome trustworthy rather than merely fast.

The question in front of security leaders has moved accordingly, from whether to adopt agentic security operations to how quickly it can be done and how to know the outcomes can be trusted.

Building and Operating an Agentic SOC Is Out of Reach for Many Teams

For most organizations, building and running an agentic SOC internally is a costly and resource-intensive undertaking that requires significant, sustained investment.

It takes real expertise and real expense to design, engineer, and tune the systems, then an ongoing commitment to keep them operating and evolving as attacker tradecraft moves.

Most organizations willing to take on the engineering still run into the economics, because most AI security tools are priced the way the underlying models are priced, by token consumption. Even with today's heavily subsidized AI pricing, that meter runs.

Worse, consumption pricing pushes a do-it-yourself program toward exactly the wrong decisions, restricting the telemetry fed to AI workflows, capping the number of investigations opened, and leaving human analysts under the same barrage of alerts the project was meant to relieve.

No security team should have to choose between saving tokens and seeing threats.



World-Class Security Operations, Within Reach for Every Team

We built the Aurora Agentic SOC so that AI-native, human-backed security operations are available to organizations of any size and any level of maturity, at a fraction of what an internal program costs.

Through the Aurora Agentic SOC and the Aurora Superintelligence Platform, customers get predictable pricing, unlimited data ingestion, and unlimited investigations, while we run the underlying AI, infrastructure, and data on their behalf.

Unlike homegrown programs that stay a few months from ready, ours deploys in as little as 10 days and plugs into the tooling an organization already has.

The figures in this report describe 13 months of operating this model at scale:

✓ **Human-only investigations fell from four in five to one in 50**

✓ **Autonomous resolution more than tripled**

✓ **Significant improvements were made in investigation speed across all case levels**

Security operations built on broad visibility and machine-speed response will manage this shift. Point-in-time defenses running through a human queue will be outpaced by adversaries who have already made the change. The window for deciding which one your organization runs is now.

“Ask a CISO what’s actually kept agentic security out of reach, and it comes down to this: they can’t afford to build it, and even if they could, they’d struggle to trust it.”



DAN SCHIAPPA
President of Technology and Services
Arctic Wolf

Building and running an agentic SOC takes real expertise and ongoing investment. We invite you to see how that compares for your organization. Our [Agentic SOC Calculator](#) lets you work through the numbers yourself.

[Spend 15 to 20 minutes with one of our security experts.](#) We'll talk through how you're approaching AI in security operations and what we're seeing work.

This report may include forward-looking statements. These reflect our current views and are subject to change. They are not guarantees, and actual results may vary. Statements regarding potential cost savings or performance improvements are based on internal data and modeling and may vary by organization. Results are not guaranteed and are dependent on customer environment and implementation. Customer examples, calculator outputs, and dollar figures are illustrative, rely on stated assumptions, and do not predict results for any other organization. Unless noted, figures in this report reflect our internal data for August 2025 through September 2026 and may change over time. Percentages describing how investigations are resolved reflect our operations during that period. Pricing and usage terms described herein reflect our current offerings and may change, including through usage-based charges or limits. Your agreement and Order Form govern what you pay. Outcomes from AI-driven and human-supported security operations depend on customer environment, configuration, available telemetry, and customer-authorized actions. This report is provided for informational purposes only and does not constitute a guarantee of detection or prevention. Defensive effectiveness varies by environment, configuration, and available telemetry.

Arctic Wolf is not a licensed insurance producer and does not sell, solicit or negotiate insurance products.

Gartner Citations

1. Gartner Hype Cycle for Security Operations, 2026, Published 5 June 2026 - ID G00846365
2. Gartner 2027 Planning Guide for AI, Agents and Analytics 7 October 2026 - ID G00858393

Gartner Disclaimer

GARTNER and HYPE CYCLE are trademarks of Gartner, Inc. and/or its affiliates.

©2026 Arctic Wolf Networks, Inc., all rights reserved. Arctic Wolf, Arctic Wolf Platform, Arctic Wolf Security Operations Cloud, Arctic Wolf Managed Detection and Response, Arctic Wolf Managed Risk, Arctic Wolf Managed Security Awareness, Arctic Wolf Incident Response, Arctic Wolf Concierge Security Team, and Aurora are either trademarks or registered trademarks of Arctic Wolf Networks, Inc. or Arctic Wolf Networks Canada, Inc. and any subsidiaries in Canada, the United States, and/or other countries.