



AI SOC FOR MDR: THE STRUCTURAL EVOLUTION OF MANAGED DETECTION AND RESPONSE

AQSA TAYLOR
FRANCIS ODUM
KYLE KURDZIOLEK
JOSH TRUP

AiStrike[®]

Research

We explore the newest frontiers of cybersecurity.

Whether you're looking at emerging vendors, evolving threats, or shifting architectures, our timely, opinionated insights help modern security leaders make smarter, faster decisions.

About Software Analyst Cybersecurity Research

SACR is a modern research and advisory firm built for today's cybersecurity leaders. We deliver in-depth, timely analysis across SOC operations, Identity, Network, Cloud, Application Security, Data, and AI Security; equipping CISOs, security teams, founders, investors, and practitioners with the insight they need to navigate high-stakes decisions.

With an engaged community of over **80,000** readers and followers, SACR connects with a global network of cybersecurity decision-makers and innovators. Our access to leaders across categories and industries gives us a direct line to the conversations shaping the market. By pairing these insights with rigorous technical analysis and continuous market tracking, we produce research that is both data-driven and grounded in the realities of modern security operations.

Whether you're seeking clarity on emerging technologies, evaluating vendors, or tracking market shifts, SACR delivers trusted, independent research designed to help you see clearly and decide with confidence.




Authors:

- **Aqsa Taylor** is the Chief Research Officer at Software Analyst Cyber Research, where she leads research initiatives and the CISO Arm (security leaders community). She is a published author with over a decade of experience building cloud security platforms and experience consulting more than 44% of Fortune 100 organizations on their security posture.
- **Francis Odum** is the Founder/CEO of the Software Analyst Cyber Research, where he leads the firm's research and engagement with cybersecurity leaders.
- **Kyle Kurdziolek** is a VP of Security at BigID shaping the future of AI-driven cybersecurity. He translates complex threats into practical, data-driven strategies that streamline operations, reduce alert fatigue, and modernize risk management. Known for building high-impact teams and developing the next generation of security leaders, he blends technical expertise with a forward-thinking vision for the industry.
- **Josh Trup** is Principal at Greenfield Partners, a global early-growth technology investment fund. Greenfield was spun out of TPG Growth in 2020, backing category-defining tech companies worldwide.

Table of Contents

Executive Summary	6
What CISOs Should Take Away.....	8
What Is AISOC for MDR?.....	8
How AISOC Changes MDR Economics and Delivery	9
Evolution of Security Operations.....	10
Voice of Security Leaders.....	12
The State of the SOC and MDR Market.....	15
Buyer’s Guide to finding the right platform for you	20
In-Depth Vendor Insights.....	26
AiStrike	28
Conclusion.....	35




The managed detection and response market reached \$9.6 billion in 2025 and is projected to grow to \$46.9 billion by 2035, reflecting a 17.2% compound annual growth rate driven by rising cyber threats and increasing highly skilled analyst shortage. Despite this growth, operational challenges still remain. According to the 2025 SANS SOC Survey, 69% of SOC's still rely on manual or mostly manual processes to report metrics, and only 42% of organizations use AI and machine learning tools with any customization. The consequences are measurable. Research shows that 61% of organizations admit to ignoring critical alerts that later caused breaches, exposing the gap between detection capability and actual response capacity. The market is now splitting between organizations that can staff internal SOC's and those that cannot. Large enterprises are exploring agentic AI platforms to augment existing security teams, while small and mid-market organizations increasingly turn to AI powered MDR providers for 24/7 coverage they cannot build internally.



AI is fundamentally transforming managed detection and response. What makes this moment particularly significant is the convergence happening between AI SOC platforms and MDR services. For years, security teams relied on a mix of SIEM, EDR, and MDR vendors, but these stacks created their own problems including endless alert noise, long investigation times, and overworked analyst teams stuck in repetitive triage. Now, AI SOC platform vendors are recognizing that many organizations need more than tools, they need the accountability and outcomes that come with managed services. This is driving AI SOC companies to expand into MDR service delivery, blurring the line between platforms and services. The result is a new category: AI-native MDR that combines the automation capabilities of agentic AI platforms with the human oversight and accountability transfer that organizations require.

This report examines how AI is changing MDR services and why it matters for security leaders. We look at two paths organizations are taking: building internal SOC capabilities with AI platforms versus outsourcing to AI-powered MDR providers. The key question is whether mid-market companies will bring these tools in-house or continue paying for managed services that deliver outcomes and accountability. We analyze how vendors are moving between selling platforms and offering services, why traditional MDR economics are breaking down, and what use cases are driving adoption across different company sizes. Security leaders will find practical guidance on evaluating AI-powered security operations, important questions to ask, understanding when to build versus buy, and bridging the gap between threats and the solutions vendors are actually delivering today.



Disclaimer: This report is not a market map and should not be interpreted as a comprehensive survey of the AI SOC or MDR vendor landscape. The AI SOC ecosystem is broad and rapidly evolving, and its exhaustive coverage was not the objective. Instead, this analysis focuses on understanding how AI SOC-powered platforms are reshaping Managed Detection and Response by examining a representative sample of vendors selected to highlight meaningful differences in technical architecture, operational models, and customer depth.

Executive Summary

Here are some key insights from the report:

Why MDR Is Under Pressure

Managed Detection and Response has become harder to run and harder to justify. Alert volumes keep growing, environments are more complex, and skilled analysts are still hard to hire and retain. Traditional MDR services rely on people doing triage and investigations around the clock. That model does not scale cleanly. Costs rise with headcount, investigation quality depends on which analyst you get, and consistency drops as providers grow.

Most organizations did not adopt MDR because it was perfect. They adopted it because running a 24/7 SOC in-house was not realistic. That problem has not gone away. What has changed is the technology available to solve it.

What AI SOC Means for MDR

AI-powered SOC platforms are changing how MDR is delivered. Instead of using automation to assist humans, these platforms let AI handle most of the investigation work directly. The system gathers context, correlates activity across tools, tests hypotheses, and reaches a conclusion. Humans step in when confidence is low, when business impact is high, or when response decisions need approval.

This shift matters because it removes the biggest bottleneck in MDR. Investigations no longer depend on how many analysts are on shift or how experienced they are. AI can investigate every alert, every time, using the same logic and the same depth. That leads to faster response, fewer missed signals, and far less noise reaching security teams.

Better Decisions, Not Just Faster Ones

The value of AI SOC in MDR is not just speed. It is decision quality. The platforms reviewed in this report show that AI can connect weak signals across identity, cloud, endpoint, SaaS, and network

data in ways humans struggle to do at scale. Investigations are more complete because the system does not get tired, does not cut corners, and does not ignore low-priority signals that later turn out to matter.

Several vendors now claim that only a small percentage of cases need human review. While this varies by environment and maturity, the direction is clear. AI is taking on the work that causes fatigue and inconsistency, while humans focus on judgment and accountability.

Detection Quality Becomes Part of the Service

Another important shift is that detection quality is no longer treated as a separate problem. AI SOC platforms observe which alerts are useful, which are noisy, and which gaps exist. Over time, they tune detections, suppress false positives, and in some cases generate new detection logic.

For CISOs, this changes the conversation. Instead of asking why analysts are overwhelmed, the question becomes why so many low-value alerts exist in the first place. AI-driven MDR services begin to fix the problem at the source rather than managing around it.

Context Is the Difference Between Noise and Risk

Attacks today often look like normal behavior. Without context, MDR services either escalate too much or miss real issues. The strongest AI SOC platforms build an understanding of how each organization works. They learn which tools are approved, how different teams behave, what normal looks like for new hires versus senior staff, and where risk tolerance differs.

This context allows the system to make better calls. It also allows MDR services to deliver fewer, higher-quality escalations that security teams can trust.

Trust Requires Transparency


As AI takes on more responsibility, trust becomes critical. Security leaders need to know how decisions are made. The leading platforms in this report make explainability a core feature. Every investigation shows what data was used, what steps were taken, and why a conclusion was reached.

This level of transparency is essential for audits, compliance, and board discussions. It is also how teams build confidence in AI over time.

Different Models for Different Organizations

The report shows a clear split in how organizations adopt AI SOC capabilities. Large enterprises tend to use AI to support internal SOC teams. They keep decision authority in-house and gate automation carefully. Mid-market and smaller organizations are more likely to adopt AI-native MDR services where responsibility is shared or transferred to the provider.

Neither approach is right or wrong. The difference comes down to risk ownership, governance, and operational capacity.



What CISOs Should Take Away

AI SOC-powered MDR is not about removing people. It is about using people where they add the most value. AI handles the volume and the repetition. Humans handle judgment, business context, and accountability. For security leaders, the key questions are not about features. They are about ownership. Who owns the decision. Who owns the failure. How does the system explain itself? How does it improve over time?

Organizations that align AI SOC adoption with their maturity and risk tolerance will see real gains in speed, consistency, and cost. Those that treat AI as a simple add-on will not. The future of MDR is not fully automated and it is not fully human. It is a deliberate balance between the two. As organizations confront an accelerating volume of alerts, expanding attack surfaces, and a persistent shortage of skilled security talent, the traditional Managed Detection and Response (MDR) model is approaching its structural limits. Human-led MDR services built around 24/7 analyst staffing, manual triage, and labor-intensive investigations that do not scale cleanly. Costs rise linearly with headcount, service quality varies by analyst assignment, and consistency degrades as providers grow.

AI-powered Security Operations Centers (AISOCs) are emerging as the next evolutionary step, redefining how MDR services are delivered. Rather than layering automation onto human-centric workflows, AISOC-driven MDRs rebuild the service model around machine-led investigation and response, with humans operating primarily in supervisory, exception-handling, and governance roles.

What Is AISOC for MDR?

AISOC for MDR refers to the application of advanced AI ranging from large language models to agentic automation across the full SOC lifecycle: detection, enrichment, triage, investigation, and response. Unlike legacy MDR offerings, which depend on continuous human attention, AISOC-driven MDRs rely on AI systems to conduct the majority of investigative work autonomously, escalating only high-uncertainty or high-impact cases to human analysts.

This shift materially changes both operational outcomes and service economics. Machine-led investigation eliminates analyst fatigue and variability, enabling true 24/7 coverage with consistent decision quality. Mean time to detect and respond (MTTD/MTTR) improves not simply because actions are faster, but because investigations are executed deterministically, with complete context stitched across telemetry, identity, endpoint, and cloud signals.

Several AI-native MDR providers claim that their AI analysts investigate nearly all incoming alerts, with only a small fraction often cited at ~3% requiring human intervention. This allows dramatically higher customer-to-analyst ratios and more predictable service quality, making the model particularly attractive to mid-market organizations that cannot afford to staff or manage an internal SOC.

How AISOC Changes MDR Economics and Delivery

AISOC-driven MDR fundamentally alters the unit economics of managed security services. In traditional MDR, quality is constrained by human availability, skill distribution, and burnout. As providers scale, maintaining consistent investigation depth becomes increasingly difficult. In an AI-native model, the opposite dynamic becomes possible: investigative logic improves as systems process more incidents and environments, allowing learning to compound at the platform level rather than being fragmented across individual analysts.

Operationally, AISOC platforms increasingly support automated containment, isolation, and remediation actions. Some platforms enable autonomous response for clearly defined, high-confidence scenarios while preserving human approval for

sensitive actions. Others focus on explainable, natural language-driven playbooks that reduce the need for brittle SOAR engineering and make response logic easier to audit, tune, and govern. Deterministic automation frameworks further ensure that every action is traceable and defensible; an essential requirement as MDR providers assume greater responsibility for outcomes.

The net effect is a service model that delivers faster response times, more consistent investigations, and lower marginal cost per customer. Importantly, these benefits are not driven solely by cost reduction; they emerge because automation replaces linear human labor with scalable systems capable of operating under sustained load.



Evolution of Security Operations

Traditional security operations centers were built for a very different era of cybersecurity. They were designed around signatures, rules, and discrete alerts generated by a growing ecosystem of vendor technologies. Analysts relied on correlation engines and complex queries to stitch together activity across log sources, identity systems, endpoints, and networks. As telemetry volumes exploded, even the most resourced SOCs found themselves overwhelmed. To cope, teams narrowed detection scopes and tuned aggressively, sacrificing visibility in exchange for what felt like a manageable workload.

This operational burden fell squarely on analysts. Large portions of their time were spent on repetitive enrichment tasks, manually pulling context from logs, threat intelligence feeds, asset inventories, and identity systems just to determine whether an alert was meaningful. SOAR platforms promised relief, but in practice they automated only small, brittle segments of the workflow. Fatigue continued to mount, while proactive threat hunting, environmental learning, and post-incident analysis were deprioritized simply to keep up with the backlog of alerts.

Several forces are now driving a rethinking of this model. The volume and velocity of data continue to increase. Attack techniques are more adaptive and blend into normal behavior.

Talent remains scarce, expensive, and hard to fill. At the same time, organizations have become more comfortable running AI in production across revenue generating and critical systems. That comfort has not appeared overnight. It has been shaped by years of operating machine learning driven fraud detection, recommendation engines, and capacity planning systems. Security is no longer the first place AI is trusted, but it is no longer the last either.

The growing operational comfort is what makes AI adoption in the SOC more realistic today. Preparedness comes from starting with augmentation rather than autonomy, demanding explainability, and grounding models in environment specific data. When AI is trained on how an organization actually operates, rather than


generic threat patterns, trust develops naturally through repeated validation. Analysts see better prioritization, fewer dead ends, and clearer reasoning. Leadership sees reduced dwell time and more consistent outcomes without surrendering accountability.

The comparison between traditional in-house SOC platforms and MDR offering further accelerates this shift. Many organizations already rely on MDR providers because they cannot sustainably staff, operate 24/7, and also don't want liability for the organization. MDR has already proved that outsourcing analysis and triage could work when paired with human expertise and clear escalation paths. Yet a problem still existed, as it caused fatigue when MDRs do not translate that human expertise to their technologies. Now, AI-MDRs can make this shift for the better and offer similar capabilities as an in-house AI-SOC. However, in-house AI-SOC capabilities represent an evolution, but one that keeps institutional knowledge inside the organization. Instead of sending telemetry out, intelligence is brought inward, embedded directly into daily workflows and decision making. It is all about organizational preference, how the data is handled, amount of risk willing to accept, and if they want to accept liability for bringing it in house. Transparency and clear escalation pathways are always non-negotiable no matter which approach is chosen.

The financial story reinforces this transition. Organizations are already spending heavily on MDR services, overlapping tools, and labor intensive operations. AI-driven SOC capabilities are increasingly viewed as an investment in efficiency rather than experimental expense. Savings come not only from reduced reliance on external services, but from better utilization of existing teams. Analysts spend less time validating noise and more time solving meaningful problems. Over time, improved decision quality reduces incident impact, repeat findings, and operational drag, creating a compounding return that leaders understand and support.

The introduction of AI-SOC in MDR allows security operations to be reimaged without narrowing detection breadth. Instead of suppressing data, teams can allow AI to process streams of telemetry and surface the alerts that actually matter. The system can learn what normal

looks like in a specific environment and identify deviations that warrant attention. What once required hours of manual effort can now be presented in moments, allowing humans to engage earlier with far better information.



Voice of Security Leaders

In my last report on [Security Data Pipeline Platforms](#), I talked about the fundamental practitioner concerns related to SOC operations, based on interviews, surveys, and insights from security leaders. Security Data Pipeline Platforms (SDPPs) address the first half of the challenge, which is the data clarity problem. The second half concerns detection and response platforms. Francis and Rafal's report on the AI SOC market landscape does a good job of explaining the fundamentals of what constitutes an agentic AI SOC platform and the architecture evolution. It highlights. This report focuses on how AI SOC platforms are transforming the MDR (Managed Detection and Response) industry.


We spoke with security leaders about their biggest challenges within the SOC. According to them, these are the top challenges that create opportunities for AI SOC and AI MDR platforms to address:

Alert overload and staffing constraints

Security teams spend four times more budget on people than tools, yet they still drown in alerts. Many organizations receive thousands of alerts per week and cannot investigate them thoroughly. Internal SOCs can cost more than a million dollars a year. SMBs simply cannot afford this, and traditional MDR models still rely heavily on humans. This creates slow, inconsistent triage and burnout.

Lack of skilled analysts

For organizations, finding L3 analysts and a team with deep skills has been brought up as an additional gap. The real challenge is that SOC work has fundamentally changed, but training and hiring haven't kept pace. Analysts face alerts from 28+ tools on average, must understand cloud-native architectures, need to correlate identity and runtime signals, and should think strategically about coverage gaps, all while the industry still hires primarily for SIEM skills and endpoint knowledge.



Source Ref: <https://cloud.google.com/blog/products/identity-security/the-dawn-of-agentic-ai-in-security-operations-at-rsac-2025>

Data complexity and inefficient investigations

Raw logs from cloud, SaaS, identity, and endpoints require analysts to write complex queries, pivot through multiple systems, and manually reconstruct what happened. Some SIEMs still require SQL-style queries for basic searches. Investigations often take more than an hour because tools do not translate events into real-world meaning.

Detection coverage gaps

New services such as Snowflake, GitHub, and Google Workspace generate critical activity with little native detection. Many teams also deploy rules without knowing how they map to their environment. Threat intelligence remains disconnected from actual coverage. The result is blind spots and reactive security.

Institutional knowledge loss

SOC workflows depend heavily on human memory. Playbooks live in individual brains, not systems. When people leave, context leaves with them. SOAR proves too rigid because it requires constant upkeep. AI tools that don't learn from past investigations repeat mistakes and lose context.

A reset moment for SOC teams

Across interviews, one theme is clear: SOC teams are entering a reset phase. Data management is changing, cloud detection is shifting, and AI introduces an entirely new layer of observability and risk that most organizations have not yet prepared for.


Understanding SOC and Agentic Impact

The Central Question that security leaders are asking: Is the Agentic SOC Real or Marketing?

Across the conversation, there is strong consensus that:

- Tier 1 automation (triage, enrichment, false-positive reduction) is real and many cases production-ready
- Tier 2+ autonomous investigation and response remains fragile, narrow, and highly context-dependent
- Full “lights-out SOC” are not trusted by practitioners, especially in regulated or lean security teams


The question is no longer “Can agents do SOC work?” It is “Who owns the failure when agents are wrong?” This becomes decisive in adoption patterns.



Decision Quality vs Speed

AI is changing how security operations are built, but it is not pushing organizations toward a single operating model. Instead, it is expanding the range of viable choices based on maturity, risk tolerance, and governance capacity. The most important factor is not whether AI is used, but where responsibility for outcome resides. Some organizations need

immediate coverage and risk transfer. The ideal state is being able to internalize decision making and compound learning over time. Where the future is making humans faster, more consistent, and far better prepared when incidents occur. There is an underlying distinction today, and is critical to evaluate whether AI SOC platforms or AI MDR services are the right fit at a given stage. This is where ideal models would align as organizations evolve.



Reducing Burden to Invest in Efficiency

Across all maturity stages, one of the clearest benefits of AI driven operations is the space it creates. When repetitive enrichment, correlation, and initial assessment are handled consistently and accurately, security teams can redirect effort toward areas that compound the initial AI SOC value.

Detection engineering improves when the teams have time to analyze missed signals and refine logic. Automation and remediation workflows become more robust when they are designed thoughtfully rather than reactively. Incident learnings are more likely to feed back into architecture and controls. AI does not eliminate the work, but it elevates it.

The State of the SOC and MDR Market

Traditional MDR has reached meaningful scale, with leading vendors generating hundreds of millions in annual recurring revenue. However, the model remains structurally constrained. Margins average approximately 10%, reflecting a service delivery approach that is heavily dependent on human labor. As a result, revenue growth remains closely tied to incremental headcount. AI-native MDR providers are challenging this structure by rebuilding operations around machine-led investigation.

Traditional MDR's Scale and Weakness:

Over the past two decades, the MDR market has produced several large, scaled providers, such as Arctic Wolf, Expel, Secureworks, and others, with hundreds of millions of dollars in recurring revenue. This scale validates sustained demand for outsourced security operations. MDR exists because staffing, training, and retaining a 24/7 SOC is operationally complex, talent-constrained, and economically inefficient for most organizations.

Yet the same model that enabled MDR to scale now imposes a hard ceiling on its economics and performance. Traditional MDR delivery remains fundamentally human-led. Analysts are responsible for alert triage, investigation,

escalation, and response across heterogeneous customer environments. While incumbent providers have invested heavily in tooling, playbooks, and automation, these investments have improved efficiency only at the margins. They have not altered the core cost structure. Incremental growth still requires incremental headcount, particularly for Tier 2 investigation and continuous coverage.

This creates a persistent mismatch between value delivered and value captured. Customers expect improving detection and response outcomes at stable or declining prices, while providers face rising labor costs, analyst burnout, and chronic retention challenges. In practice, this tension manifests in a way buyers know well: high contract prices paired with uneven service quality. The effectiveness of a traditional MDR engagement often hinges less on the underlying platform and more on the specific analysts assigned to the account. Strong teams deliver acceptable outcomes; weaker teams generate alert fatigue, excessive false positives, and slow or shallow investigations.

At scale, this variability is unavoidable. Human-led services struggle to deliver consistent depth, speed, and judgment across thousands of customers simultaneously. Even best-in-class incumbents are constrained by the realities of analyst availability, experience distribution, and operational load. The result is a market where many organizations pay a premium for protection that is sufficient most of the time, but fragile under pressure.

The AI-First Disruption

AI-native MDR is emerging as a direct response to these structural limitations. Rather than layering automation onto an analyst-centric operating model, a new class of providers are rebuilding MDR around a different unit of work: machine-led investigation and response, with humans operating primarily in supervisory, exception-handling, and escalation roles.

The defining shift is not simply faster execution or lower cost. It is a change in how service quality behaves at scale. As mentioned, in traditional MDR, quality is constrained by human variability:


analyst skill, fatigue, turnover, and staffing ratios. As providers grow, maintaining consistent investigation depth and response quality becomes increasingly difficult. In an AI-native model, the opposite dynamic becomes possible. As systems process more incidents, environments, and adversary techniques, investigative logic can improve systematically. Learning compounds at the system level rather than being fragmented across individual teams.

Where this approach works, service quality becomes less dependent on which analysts are assigned to an account and more dependent on the maturity of the underlying system. False positives can be reduced through improved correlation, context stitching, and confidence scoring. Investigations can become deeper and more consistent as evidence gathering and hypothesis formation are standardized. Over time, the service improves not through staffing optimization, but through model iteration and system-level learning.

This shift carries secondary economic consequences, but they are downstream of the

quality argument. If investigation and initial response can be handled reliably by automation, the marginal cost of service delivery drops materially. Providers can support more customers per human, prices can move down relative to traditional MDR, and gross margins can expand. These economics are not the primary value proposition; they are a natural byproduct of replacing linear labor with scalable systems.

That said, this outcome is not guaranteed. Where automation fails to handle real-world complexity, heterogeneous environments, ambiguous signals, novel attack paths, AI-native MDR reverts to human-heavy workflows and inherits the same constraints as incumbents. In those cases, AI becomes a productivity layer rather than a structural advantage. But where automation proves durable, the implications are clear: more consistent detection and response, improving service quality with scale, and a delivery model that breaks the historical trade-off between growth and effectiveness. At that point, the distinction between traditional and AI-native MDR is no longer semantic, it is operational, economic, and competitive.



The Critical Debate: Insourcing vs. Outsourcing AI

The rise of agentic security operations revives an old question with sharper consequences: should automation be operated internally, or consumed as a managed outcome? Or both? In an AI-driven SOC, this is no longer a debate about control or customization. As decision-making shifts from humans to systems, automation concentrates risk. The question becomes simple: who owns failure.

The Agentic SOC Platform (Internal):

Agentic SOC platforms are designed to empower internal enterprise SOC teams. Their promise is compelling: offload repetitive work, accelerate investigations, and allow analysts to operate at a higher level of abstraction, supervising AI-driven workflows rather than executing every step manually. In large organizations with mature security programs, this vision aligns well with long-term goals of efficiency and analyst leverage.

In practice, however, adoption is cautious. Internal security teams remain skeptical of deploying high-autonomy systems inside their own environments. Concerns about reliability, explainability, and AI “hallucination” are not theoretical, they translate directly into operational and reputational risk. When an AI-driven system misclassifies an incident, fails to identify lateral movement, or triggers an incorrect containment action, responsibility sits squarely with the enterprise. There is no external buffer. The CISO owns the outcome.

A useful illustration of this adoption path can be seen in platforms like Torq, which initially focused on SOC automation and orchestration rather than explicit autonomy. By allowing teams to define workflows, observe execution in production, and retain approval over high-impact actions, Torq helped analysts build trust in machine-driven execution. Its later move toward agentic capabilities via the Socratise platform builds on this foundation, layering reasoning on top of workflows teams already understand. The pattern is instructive: internal SOCs adopt autonomy gradually, with visibility and control, rather than by delegating decision-making all at once.

As a result, agentic SOC platforms are most often deployed in constrained roles: triage assistance, enrichment, investigation support, and recommendation generation. Autonomy is gated. Human approval remains mandatory for impactful actions. This does not invalidate the model, but it defines its ceiling. Insourced agentic SOC platforms function as productivity multipliers, not accountability replacements. Their adoption curve is governed less by technical capability than by governance tolerance.

The AI-First MDR Service (Outsourced):

AI-first MDR services take the same underlying technologies and deploy them through a fundamentally different contract with the buyer. Rather than empowering internal teams, providers position AI as the backbone of a managed service as internal technologies, where investigation and response are delivered as outcomes, not tools. The critical distinction is that responsibility for failure is transferred, not shared.

For mid-market organizations in particular, this framing is decisive. These buyers are not seeking maximum control or architectural purity. They are seeking predictable outcomes with minimal operational burden. When AI-driven investigation and response are delivered as a service, trust shifts away from the model itself and toward the provider standing behind it. Automation becomes acceptable precisely because the consequences of failure are externalized.

Even among large enterprises with well-established SOCs, this dynamic persists. Organizations continue to outsource critical functions such as 24/7 monitoring, surge response, and specialized investigations, not because they lack tooling, but because constant coverage and rapid response remain operationally expensive and difficult to staff internally. AI does not eliminate this need. It amplifies the advantage of service-based delivery by making it cheaper, more scalable, and less dependent on human availability.

Why Enterprises and SMBs Diverge Sharply

What often gets lost in discussions about AI SOC versus AI native MDR is that these are not opposing philosophies or competing futures. They are rational responses to where accountability lives inside an organization. Much of the debate assumes a single buyer profile and a single definition of readiness, when in reality security programs evolve as risk ownership shifts inwards. The divergence in AI SOC adoption is often explained as a function of technical maturity or organizational sophistication. The true dividing line is risk ownership. Enterprises and SMBs operate under fundamentally different constraints in how failure is absorbed, governed, or otherwise punished. As a result, they are rationally converging on different models for adopting AI in security operations.

Large Enterprises

Large enterprises do not approach agentic SOC technology as simply an opportunity for replacement. They approach it as a governance problem. These organizations already run staffed SOC teams embedded within audit, compliance, insurance, and board oversight frameworks. In this environment, any system that investigates

incidents or executes response actions inherits the same accountability burden as a human analyst. Autonomous decisions must be explainable, reversible, and defensible, often months later, under regulatory or legal scrutiny.

This makes broad autonomy structurally difficult to deploy, regardless of technical capability.

As a result, enterprises adopt AI primarily as an augmentation layer. Agentic capabilities are welcomed where they compress analyst workload, alert triage, enrichment, summarization, and guided workflows, but decision authority remains human. Even when autonomous response is technically viable, it is tightly gated behind approvals and policy constraints. The cost of an incorrect containment action inside a complex production environment is disproportionate to the marginal efficiency gains autonomy might deliver.

Crucially, this does not eliminate outsourcing. Most large organizations continue to externalize a meaningful portion of SOC coverage, particularly for 24/7 monitoring, surge capacity, and edge-case response. This persistence of MDR is not an indictment of tooling, rather, it reflects the economics of human availability and the practicality of transferring certain risk-bearing functions. In enterprise environments, agentic SOC platforms are



evaluated as productivity tools, while accountability remains in-house or contractually shared with service providers.

Large Enterprises

- Already have SOCs
- Use AI to augment, not replace
- Still outsource ~20–30% of SOC coverage (overnight, edge cases)
- Prefer tools / platforms + partial services

Mid-Market & SMBs

SMBs and mid-market organizations operate under a different set of constraints. Many lack the headcount, budget, or operational depth to staff a SOC at all. But more importantly, they lack the capacity to absorb the consequences of failure. A missed incident or a delayed response can be existential. For these buyers, the appeal of an “agentic SOC platform” is limited. Tools today still require supervision, tuning, and ownership of outcomes. That is precisely what these organizations are trying to avoid.


This is why AI-native MDR resonates so strongly in the mid-market. The defining value proposition is not autonomy; it is accountability transfer. When AI-driven investigation and response are

delivered as a service, trust is mediated through the provider, not the model. Buyers are more willing to accept automation when responsibility for failure, operationally, contractually, and reputationally, sits outside their organization. In this context, AI is not replacing analysts; it is replacing the cost structure of service delivery.

Much of the public debate around whether “AI SOC is real” collapses these buyer realities into a single narrative. Enterprises questioning the safety of autonomous investigation are not invalidating AI-native MDR. SMBs adopting automation-heavy services are not endorsing fully autonomous SOC platforms. Each is responding rationally to its own risk profile and governance constraints. Treating these decisions as comparable leads to false conclusions about market readiness.

Most debates about “is AI SOC real?” fail because they collapse these categories into one.

The outcome is a durable split. Agentic SOC platforms will find adoption inside enterprises as constrained augmentation tools, bounded by governance and accountability requirements. AI-native MDR will gain traction where responsibility transfer matters more than architectural purity. This divergence is not a temporary phase. It is the natural consequence of how security risk is bought,





managed, and blamed, and it will shape the SOC and MDR markets for years to come.

SMB / Lower Mid-Market

- No SOC, or 1–2 overloaded analysts
- Historically forced into MDR
- Least likely to insource an “agentic SOC platform”
- Most likely to adopt AI-native MDRs

Buyer’s Guide to finding the right platform for you

As AI driven security operations mature, buyers quickly realize that the real decision is not whether to use AI, but to apply it within their operating model. This is where many organizations take a pause. The choice between an AI-enabled SOC internally, and an AI-augmented MDR offering is not binary. It is contextual, shaped by maturity, risk tolerance, and realities of people, process, and business constraints.

Core Questions

Before comparing vendors or delivery models, security professionals should anchor on a small set of foundational questions. These questions cut through feature lists and reveal whether a solution will actually improve outcomes.


The first question is how the system explains its reasoning. AI driven triage and recommendations must be interpretable. Practitioners need visibility into why an alert was prioritized, which signals mattered the most, and where uncertainty exists. Explainability is not a nice to have, it is the mechanism through which trust is built and maintained.

The second question is how much transparency and control the organization retains. Buyers should understand when AI is recommending, when it is acting, and when it is deferring to humans. Mature platforms allow teams to tune this balance over time as confidence grows, rather than locking them into rigid levels of autonomy.

Privacy and data handling form the third question. Security data is inherently sensitive. Organizations must understand how telemetry is stored, how long it is retained, where the data is stored, whether data is shared across tenants, and how models are

trained. Clear boundaries are essential especially when introduced to regulated environments.

The fourth question is why this approach works now. Vendors should be able to articulate what has




changed technically and operationally that enables better outcomes today than earlier automation attempts. Answers grounded in improved context ingestion, reasoning, and integration depth are far more meaningful than references to model sizing.

Finally, buyers should ask how the system improves over time. AI that does not learn from analyst feedback, environmental changes, and incident outcomes will stagnate program maturity. Continuous learning tied to real operational decisions is what differentiates an AI SOC from traditional correlation searches.

Board and Executive Talking Points on AI-Driven Security Operations

When discussing AI driven security operations with executives or boards, the conversation must move beyond tools and features. These talking points help anchor the discussion in outcomes and risk management.

- AI is about leverage, not replacement
 - This is not a workforce reduction strategy. AI reduces manual overhead so highly skilled security professionals can focus on judgement, strategy, and resilience.
- We are buying decision quality, not speed
 - Faster triage only matters if it leads to better outcomes. The value of AI lies in improved prioritization, clearer context, and more consistent decisions under pressure.
- Reduces long-term operational cost and risk
 - Organizations already spend heavily on tool sprawl and reactive responses. AI driven operations improve efficiency while keeping knowledge and control within the organization.
- We maintain accountability and governance
 - AI does not act autonomously without oversight. Humans remain accountable for containment, escalation, and business tradeoffs. Transparency and explainability are mandatory.
- This scales with the business
 - As the organization grows, attack surface and complexity grow with it. AI allows security capability to scale without linear increases in headcount or burnout.



Vendors Analyzed for AI SOC powered AI MDR Capabilities

The following vendors offer AI SOC powered MDR services and were analyzed deeply for their technical capabilities and business use cases, via product deep dives, multiple briefings, questionnaires and customer interviews.

24/7 MDR Service Delivery at Scale

AI SOC-powered MDR enables true 24/7 coverage, scaling analyst capacity and delivering consistent outcomes without human fatigue. To further evaluate innovation in this industry, we did a deep dive into the following 7 vendors with deep dive demos, questionnaires and customer interviews. Here are some highlights –

7AI: Supports 24/7 MDR service delivery with automated escalation, auto-response to high-priority incidents, and external support teams for additional context. Customers can assign cases directly to 7AI, which is building out full 24/7 service offerings for MDR replacement.

AirMDR: Delivers 24/7 MDR with AI-native analysts handling the vast majority of alerts. Only 3% of cases require human touch, allowing for massive scaling of coverage and improved SLA times. AI-driven MDR is more cost-effective and consistent than traditional human-led approaches.

AiStrike: Delivers Agentic Cyber Defense as a Service, which replaces traditional MDR by using AI-driven agents to continuously detect, investigate, and respond to threats while operating, tuning, and optimizing defenses without requiring organizations to manage another tool.

Conifers AI: Multi-tenancy and predictable pricing models make Conifers suitable for MSSPs and organizations seeking to scale MDR delivery. The platform's agentic AI continuously adapts to new environments, supporting high-volume, multi-tenant operations.



Daylight Security: Employs a follow-the-sun expert team model (US, Singapore, Tel Aviv) and claims to support higher customer-to-analyst ratios than previously possible. The hybrid automation-services model enables Daylight to deliver premium, scalable MDR with rapid investigation and zero open threats for key customers.

Exaforce: Provides both self-managed and fully managed MDR options with 24/7 coverage. Customers can automate as much as possible, or leverage Exaforce's MDR for full coverage, with the platform supporting both approaches depending on organizational needs.

Swimlane: The Turbine platform is proven at scale, supporting thousands of daily users and billions of automated actions monthly. Swimlane's architecture enables predictable, extensible, and resilient MDR service delivery, with multi-tenancy and role-based access control for large-scale operations.

Alert Triage & Investigation Automation

AISOC platforms have transformed alert triage and investigation, reducing analyst workload, improving consistency, and accelerating response.

7AI: Automates enrichment from all sources into tickets, enabling faster investigation and response without additional headcount. 7AI can handle full investigations and conclusions for selected or all use cases, supporting both in-house SOCs and MDR overlays. The platform enables automated triage, prioritization, and escalation, with support for high-value response actions and oversight by internal teams when desired.

AirMDR: Uses AI virtual analysts to automate 80–90% of alert triage, investigation, and response. Automated playbooks execute in under 5 minutes compared to over an hour for human analysts. NLP-driven capabilities allow the system to answer questions, learn facts, and document incidents with transparency, supporting comprehensive remediation and learning.

AiStrike: Automates alert triage by grouping related alerts based on MITRE framework patterns and repeatability, applying behavioral context and organizational policies to filter noise and surface high-risk threats from toxic alert combinations.

Conifers AI: Delivers automated triage and investigation through tool-using agents that operate step-by-step, rather than as monolithic models. Conifers integrates directly with SIEM, EDR, and ITSM tools (e.g., ServiceNow, Jira), embedding its findings into existing workflows. The platform's "Cognitive First Analysis" provides consistent, well-informed triage decisions and recommendations directly within analysts' current workbenches.

Daylight Security: Features a streaming detection pipeline and agentic investigation platform (AIR) that can correlate alerts across cloud, endpoint, and identity sources in under a minute. The system leverages a knowledge graph for business context and supports fully automated, agentic investigations across multiple channels (Slack, Teams, Email), significantly reducing investigation times.

Exaforce: Exabots (AI agents) perform triage, investigation, and response. The multi-model AI engine pre-processes data to build behavioral baselines and peer comparisons, enabling the system to triage alerts and conduct investigations at scale. Exaforce claims a 10x improvement in SOC productivity and efficacy, with the ability to perform investigations like a Tier 3 analyst.

Swimlane: Swimlane's Turbine platform leverages deep agents and proprietary AI to ingest, build, and execute security investigations for both known and unknown alerts. Agents can build playbooks from natural language, provide recommendations, and execute actions with full traceability and auditability. The platform's low-code interface enables rapid creation of security applications and workflows.

Threat Detection Across Multiple Signal Sources

AISOC platforms ingest and correlate data from diverse sources EDR, SIEM, cloud, identity, and network to detect sophisticated threats.

7AI: Supports automated enrichment from all sources, integrating with existing detection and response pipelines. The platform can be deployed alongside MDR providers to cover additional use cases not handled by legacy MDRs, such as cloud or identity-focused threats.

AirMDR: Delivers more than 200 out-of-the-box integrations 40–50% more than typical MDR providers. AirMDR’s virtual analysts can ingest and correlate signals across cloud, endpoint, and identity, supporting rapid and comprehensive threat detection.

AiStrike: Integrates with existing detection sources like SIEM and EDR while providing its own detection coverage for gaps such as newly discovered threats or SaaS data sources, delivering detection as code that can be deployed directly on decentralized data stores.

Conifers AI: Integrates with SIEM, EDR, cloud security, and identity platforms, creating a semantic layer for interactive data exploration and threat hunting. The platform’s continuous learning adapts to each customer’s environment for context-aware detection.

Daylight Security: The AIR platform’s streaming detection pipeline ingests and correlates signals from EC2, CloudTrail, IDP, threat intelligence, ZTNA, and EDR, supporting multi-source detection and rapid investigation. The knowledge graph further enriches context for detection.

Exaforce: Goes beyond typical SOAR integrations by ingesting events, configs, identity data, and code artifacts from platforms like AWS, GitHub, Azure, and Google Workspace. Exaforce provides out-of-the-box behavior-based detections in addition to those ingested from SIEM and EDR, supporting comprehensive multi-signal detection.

Swimlane: Offers thousands of integrations (5,000+ third-party actions) and an autoscaling automation engine, supporting high-velocity data ingestion and action execution. Swimlane’s platform is built to process billions of automated actions monthly, supporting multi-signal detection at scale.



Automated Response & Remediation

AISOC platforms orchestrate and execute containment, isolation, eradication, and remediation actions with minimal human intervention.

7AI: Supports automated remediation with configurable thresholds for auto-action vs. recommendations, based on customer risk tolerance. Can auto-respond to high-priority items, escalate via modern equivalents of call trees, and integrate with customer remediation workflows.

AirMDR: Automated playbooks for response and containment execute in under 5 minutes, with full transparency and documentation. The system supports comprehensive remediation and learning, reducing mean time to respond (MTTR) and improving incident outcomes.

AiStrike: Provides built-in SOAR capabilities and case management to execute automated response actions, while also supporting integration with external SOAR platforms like Torq for workflow automation.

Conifers AI: Provides recommended response steps based on investigation findings, integrated directly into existing case management systems.


Conifers can trigger customer-approved remediation workflows in SOAR, Sentinel, or ServiceNow, supporting both automated and human-in-the-loop response actions.

Daylight Security: The AIR platform enables agentic, context-aware response actions, including containment and remediation across multiple channels. The system's integration with business context ensures that remediation is aligned with organizational policies and priorities.

Exaforce: Exabot Respond enables teams to build response actions through natural language, reducing the need for complex SOAR engineering. Automated containment and response actions are executed in minutes, with a human-in-the-loop model for high-impact decisions. The platform also supports custom automation agents for complex remediation workflows.

Swimlane: Provides deterministic automation of response and remediation actions, with AI agents executing playbooks reliably and immediately. Human-in-the-loop approval is supported for critical actions, ensuring transparency and auditability. The platform's automation fabric guarantees reliable, large-scale execution of remediation workflows.

In-Depth Vendor Insights



AiStrike[®]

AiStrike

AiStrike was founded by members of the original Securonix founding team after observing that security operations do not scale with human effort. The platform uses a multi-agent architecture built on RAG (Retrieval-Augmented Generation) and KGAG (Knowledge-Graph Augmented Generation) to support accurate and explainable analysis. Multiple agents are designed for specific tasks such as detection, investigation, response, alert grouping and correlation, emerging threat identification, workflow coordination and response automation, all managed through a central orchestration layer. Each agent has a defined role and analyzes live telemetry and alerts, enriched entity data covering users, assets, identities, and vulnerabilities, along with threat intelligence and historical cases. Agents use built-in tools such as queries, playbooks, and system integrations to collect evidence, test assumptions, and verify results before recommending and taking actions, aligning automated reasoning with analyst workflows.

AiStrike also evaluates existing detection rules to determine how they perform in real environments, automatically tuning noisy or incorrect logic and generating new detections when coverage is missing. The platform measures detection coverage

by mapping existing rules to frameworks such as MITRE ATT&CK and identifying specific gaps. It correlates relevant threat intelligence with internal telemetry to determine which adversary techniques matter for a given environment and generates detections for those techniques. The company also provides 24x7x365 Managed Cyber Defense Service – for full managed detection and response service.

Voice of the Customer

Prior to adopting AiStrike, the customer relied on a traditional MDR provider for Tier 1 alert triage and detection engineering. This model proved increasingly inefficient and costly in a complex hybrid environment. A significant portion of alerts escalated to the internal Tier 2/3 team, creating alert fatigue and limiting the team's ability to focus on higher-value investigations. Many escalations lacked actionable context or were repeat false positives, while detection tuning and new detection development delivered limited value. The MDR service also showed integration gaps, including lack of native workflows with tools such as Slack, ServiceNow, Wiz, and other cloud security controls, reducing operational efficiency and visibility.



Company Vision

AiStrike's vision centers on eliminating the human scalability bottleneck with the help of AI for modern security operations. The founding team, which came from Securonix, observed that even with mature SIEM and SOAR platforms, organizations still struggled with four systemic problems: detection engineering that doesn't scale beyond manual rule creation, alert volumes that grow faster than teams can handle, SOAR playbooks that remain static while threats evolve, and centralized architectures that create operational inefficiency in distributed cloud environments. Their founding thesis is to replace static, human-engineered workflows with a self-improving, federated, agentic AI SOC architecture.

Architecture and Deployment Maturity

AiStrike supports SaaS, bring-your-own-cloud deployments, and hybrid models – a SaaS application with on-premises forwarders.

Data Collection and Ingestion Methods

The platform consumes detection events, environment context and Threat Intelligence data needed for investigations.

- SIEM integrations: Detection events and alerts from existing SIEM deployments feed into the investigation workflow
- EDR platforms: Endpoint detection and response alerts
- Cloud platforms: AWS, Azure, GCP audit logs
- SaaS applications: Google Workspace, Okta, and other SaaS platforms for user activity and access patterns
- Network security: Network logs and traffic analysis
- Threat intelligence: External threat intelligence sources like Intel45, NVD, VirusTotal
- ITSM platforms: JIRA and custom ticketing platforms for result delivery and case management
- SOAR platforms: Like Torq

Capabilities

AiStrike supports the full SOC lifecycle using a closed-loop learning model. The platform improves detection engineering over time, turns threat intelligence into usable detections, and automates threat hunting. It supports investigations by identifying root causes and executes response actions through integrated case management.

AI Guardrails and Explainability

As we move into the AI SOC world, it's important to note the different ways vendors address security and privacy concerns. Here's how AiStrike protects and maintains privacy of customer information.

Data Privacy and Security:

AiStrike limits AI access to only the data needed for each task, with sensitive information protected through redaction, anonymization, and tokenization, including PII and PHI. Data is processed only for the duration of the inference and is discarded immediately, with no long-term storage. The system includes protections against prompt injection, jailbreak attempts, adversarial input manipulation, and model misuse. To reduce the risk of incorrect outputs, AI responses are grounded in trusted sources such as enterprise knowledge bases, contextual data, threat intelligence, and standard operating procedures. The platform enforces strict tenant isolation to prevent any cross-customer data exposure and uses encryption in transit and at rest with TLS and AES-256. Customer data is not used to train foundation models by default. They are certified by SOC 2 Type II, ISO 27001, ISO 27701, ISO 42001, NIST 800-53. AiStrike is SOC 2 Type II certified and is pursuing FedRAMP authorization to support government sector customers.

Explainability:

AiStrike provides full transparency for every investigation with justification steps, sources and context. Every AI action is recorded with full lineage, clear traceability, and supporting evidence to explain how decisions are made. The platform maintains immutable audit logs that capture each inference, action, and automated decision, enabling forensic analysis and compliance review.

Agentic SOC Capability Matrix

Capability	Capability
Tier 1 - Detection and Triaging	
Detection	The platform integrates with existing sources of detections like SIEM, EDR and gathers security context across sources listed above. In addition to that, AiStrike also provides their own in-house detection coverage for instances where detection is missing - such as newly discovered threats or SaaS data sources. AiStrike also delivers detection as code that can be directly deployed on decentralized data stores - e.g., data lakes to enable federated detection architecture.
Triage	At the triage layer, AiStrike applies alert insights, including policies, approved software, user roles, and behavioral patterns, to identify real threats and filter out noise. The system learns from analyst feedback in natural language and uses explanations of false positives to improve future investigations. Rather than triaging each alert independently, AiStrike groups related alerts based on MITRE framework patterns and repeatability to surface high risk threats based on toxic combination of alerts and deprioritize repeat hygiene issues.
Context Building	The system continuously correlates signals from SIEM, CNAPP, identity systems, logs, security tools, and threat intelligence, and generates prioritized recommendations, including detections to add or adjust, response actions to take, and tickets to create or update, with confidence scores and supporting evidence. Depending on policy, these actions can be executed automatically or require analyst approval.
Tier 2 - Analysis	
Investigation	A workflow or orchestration agent decomposes a request, such as "investigate this alert" or "tune detections for this threat," into individual steps, assigns them to the appropriate agents, and enforces policy and guardrails.
Justification	Every AI-driven action is explained by complete lineage, end-to-end traceability, and systematic evidence capture, enabling auditability through immutable logs that record every inference, action, and automated decision for forensic and regulatory review. Each decision is further contextualized through enriched entity mapping and investigative metadata, ensuring transparency and analytical depth.
Tier 3 - Response	
Recommendations	AI recommendations are delivered with explicit confidence scoring and are designed for structured analyst review, reinforcing human oversight. All of this operates within a centralized governance framework that supports compliance reporting, policy enforcement, and continuous monitoring across the AI lifecycle.
Execution	AiStrike provides built-in automation and SOAR capabilities to take automated actions. In addition, AiStrike provides complete case management capability to track and report on incidents. AiStrike also supports integration with external SOAR platforms to automate response.
Advanced Features	AiStrike also provides detection coverage gap analysis, out of the box threat intel and enrichments even without customers having their own intel providers.
Dynamic Adaptability	
Human Feedback Loop	AiStrike uses a human feedback loop to tune detections and alerts. All recommended actions are presented to analysts who can then approve, modify or reject actions.
Environmental Changes	AiStrike uses continuously updated data sources, along with threat intelligence and detection feeds, while also adding new detection coverage to ensure organizations maintain up-to-date threat visibility and protection.
Threat Updates	In AiStrike, threat intelligence is a native, continuously active control loop with emerging TTPs and IOCs continuously ingested. Exposure is automatically assessed by correlating TTPs to alerts, vulnerabilities, identities, and raw telemetry. The system generates and executes automated threat hunting queries which enables continuous exposure management, not just reactive alert investigation.
Tier - 3 Threat Hunting	
Threat Content	AiStrike provides out of the box threat content (such as new detection rules) in addition to third party ingests. AiStrike also provides automated threat sweep capabilities that hunt for indicators of compromise (IOCs) and adversary techniques (TTPs) across raw event data. In addition, customers can use natural language input, which AiStrike translates into structured threat-hunting hypotheses and executable queries, enabling human-in-the-loop hunting while reducing the manual effort typically required to formulate and run hunts.

Analyst Take:

Here's what we see as top 3 strengths that AiStrike provides

Threat Detection Depth

AiStrike shows strong technical depth in threat detection by using a multi-agent AI approach. The platform correlates signals from multiple stages of an attack and across third-party tools such as GuardDuty, CrowdStrike, and Okta, similar to other AI SOC platforms but what differentiates them with a lot of others is that they also generate native detections to close gaps that existing tools often miss. Its threat exposure management continuously tracks emerging threats and maps them to specific assets in the customer environment, giving security teams early, context-rich insight into how new attack techniques could impact them..

Self-Improving Detection Engineering That Addresses Coverage Gaps at the Source

The platform evaluates existing detection rules to determine how they perform in real environments, automatically tunes noisy or incorrect logic, and generates new detections when coverage is missing. It measures detection coverage by mapping existing rules to frameworks like MITRE ATT&CK and identifying specific gaps. Rather than simply deploying thousands of new rules, AiStrike focuses on improving detection quality and reducing false positives at the source.

Zero-Trust AI Architecture with Explainability Built In

Their deployment flexibility along with security certifications listed can help security leaders in regulated industries have more assurance. The platform performs data preparation ahead of using it with AI. Pre-AI sanitization includes data redaction, anonymization, and tokenization to protect sensitive data. Data is processed transiently and discarded immediately after inference with no long-term storage. The platform includes adversarial hardening to defend against prompt injection, jailbreak attempts, adversarial input manipulation, and model misuse.

Areas to Watch

AiStrike is playing in multiple categories (SOC automation,, threat exposure automation) and their marketing strategy could struggle with a single clearly defined positioning around SOC automation. AiStrike's "Service-as-Software" positioning is conceptually interesting but operationally Security leaders may need clarity on whether AiStrike competes on cost, capability depth, or deployment speed, and which buyer persona and budget line item they're targeting. The current positioning risks being "a little bit of everything" without being focussed in a specific dimension that drives purchasing decisions. The platform profiles entities and flags anomalies, but there's no indication of multi-dimensional behavioral modeling that accounts for role-based patterns, time-based patterns, peer group comparisons, or resource-specific norms.

The Reality of “Agentic AI” From a Practitioners Perspective

When evaluating AI SOC and AI MDR platforms, this is what security leaders should take into account. Focusing solely on feature lists or detection speeds is not enough. Security leaders need to test these platforms with their own data, understand how the AI makes decisions and identify which vendors solve specific problems better than others.

How to Test and Compare

Effective evaluation requires realism. Buyers should test AI SOC and AI MDR offerings using their own data, even if it is only in limited scope. The focus should be on decision quality, rather than speed alone. How well does the system understand environment specific context? How clearly does it explain reasoning? Where does it fetch additional information? How gracefully does it escalate uncertainty to humans?

It is also important to understand where vendors specialize. Some excel at telemetry aggregation from security technologies, or even SIEMs. Others focus on reasoning, prioritization, and responses SOC's should take. Others integrate deeply with cloud and identity platforms, while others emphasize analyst expertise. Understanding these strengths help align solutions with actual needs, combined with program maturity, rather than chasing broad coverage.

Constraints that Shape Every Decision

AI does not remove fundamental constraints of security programs. People are still required to exercise judgement and accountability. Processes must evolve to incorporate AI driven insights without bypassing governance. Technology must be integrated cleanly to avoid fragmented context. Business priorities must remain the lens through which risk is evaluated. Organizations that succeed are those that treat AI as an operating model decision rather than a tooling upgrade. When applied thoughtfully, AI does not just make security operations faster, it makes them more deliberate, resilient, and aligned with how modern organizations actually function.

The Future of MDR

AISOC-powered MDR is rapidly becoming the reference model for scalable, consistent, and cost-effective security operations. The long-term winners in this market will not be defined by claims of full autonomy, but by their ability to deliver high-quality outcomes with clear accountability, transparent decision-making, and measurable improvements in detection and response.

For organizations evaluating AISOC-driven MDR offerings, the critical questions are no longer whether AI is used, but how it is governed: how investigations are explained, where autonomy is permitted, how uncertainty is escalated, and how responsibility for failure is contractually and operationally defined. The future of MDR will belong to providers that successfully blend AI-driven automation with human oversight, delivering security outcomes that scale without sacrificing trust.

“Ideal” Mapping AI SOC and AI MDR Across Maturity Stages

In early stage, or resource constrained organizations, the primary challenge is coverage. Teams are small, on-call rotations are thin, and building a full SOC is often very unrealistic. In this stage, AI augmented MDR offerings make the most sense. They combine AI driven analysis with human expertise, providing 24/7 coverage, and a structured escalation process without requiring significant internal investments. The value here is speed to capability, risk reduction, and lower liabilities, not customization.

As organizations move into a growth, or mid-stage phase, the limitations of outsourced models begin to surface. Internal teams want more visibility, more control, and deeper understanding of their environment. This is where you start to see the hybrid models emerge. AI powered SOC capabilities are introduced internally to handle enrichment, triage, and prioritization, while MDR may still provide after hours coverage or specialized support. The organization begins to retain more institutional knowledge while still leveraging external scale.

In mature security programs, the balance often shifts decisively toward an internal SOC model. Teams have the expertise to interpret nuanced signals and the desire to embed security deeply into engineering and operations. AI becomes a force multiplier rather than a substitute. MDRs may still play a role for surge capacity, additional services, or specialized threat intelligence, but the core decision making and learning loops live inside the organization. At this stage, the value of AI lies in amplifying expertise and improving decision quality across the entire lifecycle.

A Debate That Resolves Through Segmentation

This debate is unlikely to end in convergence. It resolves through segmentation. Agentic SOC platforms will continue to gain traction inside organizations that are willing and able to retain accountability, governance overhead, and operational risk. AI-first MDR services will expand where responsibility transfer, consistency, and operational simplicity dominate purchasing decisions.

The mistake is to treat these models as competitors for the same buyer in the same context. They are not. They represent two structurally different responses to the same underlying pressure: how to operate security functions in a world where human-led SOCs no longer scale cleanly. AI does not collapse the insource-versus-outsourcing distinction. It

sharpens it by making the cost of error more explicit.

A telling market signal reinforces this conclusion. Several vendors that began as agentic SOC platforms are increasingly offering MDR-style delivery, either directly or through partners.

Exaforce has moved beyond pure tooling into outcome-oriented services. 7AI has been deployed through managed security operations partners. And platforms like Torq, while historically rooted in orchestration, are increasingly oriented toward enabling managed service delivery alongside internal use cases. This pattern is not incidental; it reflects where adoption friction actually sits.

This shift does not contradict the platform thesis, rather, it validates it. Autonomy is easier to buy when it is bundled with accountability. Services provide a faster path to trust, clearer time-to-value, and tighter operational feedback loops, all of which are critical for agentic systems to improve in production.

At the same time, this convergence introduces real tension. Moving into MDR creates channel conflict, operational complexity, and a higher bar for reliability. Not every platform vendor will execute this extension successfully. But as a directional signal, it is unmistakable. When even software-first vendors feel compelled to wrap their technology in services, the market is sending a clear message: agentic capability alone is not enough. Outcomes, ownership, and accountability are what ultimately drive adoption.

Conclusion

AI is reshaping security operations, but it is not redefining human accountability. Across maturity stages, the adoption of agentic SOC platforms or AI-first MDR services is driven less by the raw capabilities of AI and more by who owns the risk, how governance is maintained, and what operational constraints exist. Early-stage or resourced-constrained organizations benefit most from AI-augmented MDR services that transfer accountability while providing 24/7 coverage. Growth-stage organizations often embrace hybrid models that balance internal AI SOC capabilities with selective external support, retaining knowledge

while leveraging scale. Mature security programs internalize AI to amplify expertise, accelerate investigations, and improve decision quality across the SOC lifecycle.

AI is not an elixir to solve everything, but when thoughtfully applied, it transforms the SOC from a reactive operations center into a strategic decision engine, delivering speed, scale, and insight without sacrificing accountability. The choice between AI SOC and AI MDR is less about technology and more about responsibility. Nonetheless, the right decision will shape the effectiveness, efficiency, and resilience of security operations for years to come.

Why Traditional MDR Struggles to Scale

Traditional MDR exists because running a 24/7 SOC is difficult, expensive, and talent constrained. Outsourcing detection and response has been the most practical way for many organizations to achieve coverage and reduce operational burden. The market's size confirms this demand.

At the same time, the human-led delivery model creates hard limits. MDR services still rely on analysts to triage alerts, investigate incidents, and make escalation decisions across many customer environments. While tooling and automation have improved efficiency, they have not changed the core economics. Growth still requires more people, especially for investigation and continuous coverage.

This leads to uneven outcomes. Service quality often depends on which analysts are assigned, not on the platform itself. As providers scale, maintaining consistent depth and speed becomes harder. Most organizations receive protection that works in normal conditions but degrades under pressure, high alert volume, or complex attacks.

How AI-First MDR Changes the Model

AI-first MDR emerges as a response to these limits. Instead of adding automation around analysts, investigation and initial response are handled by systems, with humans focused on oversight, uncertainty, and approval. The unit of scale shifts from analyst hours to machine-led investigations.

The key difference is how quality behaves at scale. In human-led models, quality is limited by fatigue, turnover, and staffing ratios. In AI-driven models, quality can improve as systems process more incidents and environments. Correlation improves, context deepens, and investigations become more consistent over time.

Where this model works, false positives drop, investigations are more thorough, and response is faster and more reliable. Economic benefits follow, but they are secondary. Lower cost and better margins result from better delivery, not the other way around. Where automation fails to handle real-world complexity, AI-first MDR collapses back into human-heavy workflows and loses its advantage.

Accountability Defines Adoption

As AI takes on more decision-making, the central question becomes ownership of failure. This drives a clear split in how organizations adopt AI SOC capabilities.

Internal AI SOC platforms are used to support in-house teams. Automation is gated, explainability is required, and humans retain decision authority for high-impact actions. These platforms improve productivity but do not replace accountability. Adoption is cautious and shaped by governance tolerance.

AI-first MDR services deliver outcomes, not tools. Responsibility for investigation and response is transferred to the provider. This model appeals to organizations that value simplicity, consistency, and reduced operational risk, especially where staffing and coverage are hard to maintain. The future of MDR will be defined by models that combine scalable automation with clear ownership, transparent decisions, and trust in how failures are handled.

Ultimately, the strategic value of AI lies in its ability to elevate human judgement, reduce toil, and improve consistency, not to replace the workforce. Whether deployed as a platform or as a service, AI must operate transparently, with explainable reasoning and clear escalation paths, so that organizations retain oversight while compounding operational learning over time. Leaders who understand this distinction and align AI adoption with organizational maturity, risk tolerance, and governance capacity; position their security operations to scale sustainably, respond with precision, and maintain resilience in an increasingly complex threat landscape.



business

personal



Trusted research. Sharp insights. Real conversation.

