ARCTIC WOLF

Stopping AI Enabled Attackers at Machine Speed: A Buyer's Guide to Modern MDR

Cutting Through AI Hype to Achieve Real Security Outcomes

01
Executive Summary

Attackers are moving faster, automating more of their activity, and using AI to compress the time between initial access and impact. According to the latest M-trends report, initial access to handoff, where threat group executes a high-impact operation like ransomware, happens in only 22 seconds.[1]

Security leaders are trying to stay ahead of an escalating AI threat landscape and facing a growing effectiveness gap between their investments in security technology and services. They invested in security tools but are struggling to reduce risk and respond faster.

To stop AI-enabled attacks, the security operations center (SOC) model must change. The traditional SOC model does not work for our AI future. Modern threats don't wait in queue lines, so why does your SOC work in tiers? Within the hypercharged AI-driven threat landscape, traditional SOC models are breaking down and becoming obsolete.

Managed detection and response (MDR) solutions are meant to close that effectiveness gap by delivering security outcomes through a modern security operations center without requiring organizations to build and staff one themselves. Modern MDR are built for AI led investigations and feature an agentic SOC operating model and humans in the loop to validate important decisions. To move as fast as AI-enabled attackers, detection, investigation, and response must happen in parallel, not in queues. Finally, Modern MDR should provide security posture reviews to continually improve between incidents through proactive hardening and guided remediation.

This guide is designed to help security leaders cut through the AI noise. We will provide an explanation about what modern, AI led MDR must deliver and how to evaluate providers based on outcomes rather than promises. Functionally, AI-led MDR should enable faster response, clearer accountability, and measurable risk reduction in the face of an AI accelerated threat landscape.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

02
A CISO's Perspective on AI in Security Operations

AI is now part of every cybersecurity conversation. Appearing in board meetings, vendor briefings, budget discussions, and incident reviews. AI is pervasive for security leaders. It's not a question about whether AI belongs in security operations, but how to use it effectively.

The market is flooded with AI claims, overlapping terminology, and promises of autonomy. Security leaders are being asked to chart an AI strategy while separating real operational value from marketing noise. Meanwhile, attackers are already using AI to automate reconnaissance to find vulnerabilities and chart attack paths. Attacks are accelerating to allow adversaries to move laterally faster than traditional security operations were designed to handle.

AI works best when it is used as a force multiplier inside a redesigned security operations model, not bolted onto legacy workflows. When AI handles speed, scale, and parallel investigation, security professionals can focus on judgment, escalation, and accountability. The balance of humans and AI working in parallel allows organizations to defend against AI enabled attackers at machine speed.

As a CISO, it is not your responsibility to chase new technology claims. Rather, you should ensure that your chosen AI tools within security operations actually help you detect, defend, and respond faster to keep adversaries at bay and organizational risk low.

"The real shift isn't AI itself, it's how to best utilize AI as a tool within security operations. To keep pace with attackers, the SOC must operate differently, and that's what modern, AI-led MDR is designed to deliver."



ADAM MARRÈ
Chief Information Security Officer
Arctic Wolf®

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

03
The Role of AI in Security Operations

AI has the potential to improve security operations, but only when it is applied to the right problems, like any tool. For example, AI can enable machine speed security analysis, allowing investigations to run in parallel, and compressing the time between detection and response. The SOC's ability to act quickly matters to level the playing field as attackers increasingly automate reconnaissance, exploitation, and lateral movement.

The mistake many organizations make is assuming that simply adding AI to existing workflows is enough. It isn't. When AI is layered onto legacy alert-centric SOC workflows, it accelerates activity but also increases bottlenecks, not security outcomes. For example, faster alerts do not guarantee faster containment, and automated triage does not automatically produce better decisions. Placing AI within the legacy tiered SOC model (i.e., tier 1, tier 2, tier 3, tier 4) results in AI amplifying existing bottlenecks instead of removing them.

Another challenge of adding AI into security operations is whether AI is making trustworthy decisions. Decisions made at machine speed must still be explainable, auditable, and governed within a SOC's workflow. Without clear oversight, AI can alert false positives, miss detections, and create unintended actions. Cybersecurity is too important to accept uncertainty.

The more effective role of AI in security operations is as a force multiplier inside a redesigned operating model. AI should own speed and scale by correlating telemetry, advancing investigations, and prioritizing action, while human experts retain responsibility for judgment, escalation, and accountability. When security operations are rearchitected around this division of labor, AI improves outcomes instead of introducing new risk.

The balance of AI with humans in the loop allows organizations to defend against AI enabled attackers while maintaining trust to gain positive security outcomes.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

04
Understanding the MDR Market

Managed detection and response solutions are widely described as a turnkey, remotely delivered security operations capability that provides continuous detection, investigation, and response, including active threat containment rather than simple alerting.[2]

Many services marketed as MDR do roughly the same things. They simply alert on detections and escalate those alerts to the customer. However, to improve security posture and lower risk, more must be accomplished than detection and response. Many MDR vendors do not offer proactive hardening without added cost beyond a base offering. Additionally, other MDR services lack clarity for whether risk is truly being reduced. The reason many MDR services fail to deliver measurable risk reduction is because they aren't focused on proactively improving security through hardening, so they remain a reactive measure that doesn't increase security posture.

In your consideration of MDR, it's important to understand the roles and responsibilities between the customer and the vendor. Such as, who owns the response? How are security outcomes measured? How is AI governed? Can the service consistently reduce risk over time?

The rise of an AI led SOC within MDR solutions makes the differences between providers more apparent. While autonomy and speed are appealing, trust becomes an important deciding factor for security leaders. AI that lacks transparency, validation, or real world operating history creates more questions than answers. Effective AI led MDR must deliver immediate value without introducing new complexity or shifting operational risk back to the customer.

Signs of Legacy MDR

The following patterns indicate an MDR service that has not evolved with the threat landscape:

- Automates Tier 1 or Tier 2 workflows without redesigning the SOC operating model
- Lacks clear AI governance, oversight, or validation
- Provides limited visibility beyond endpoint centric telemetry
- Forces vendor lock in through proprietary ecosystems
- Offers little or no proactive security improvement between incidents
- Relies heavily on the customer to investigate and take action to contain incident

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

05

What Modern AI-led MDR Must Deliver

Modern MDR must deliver outcomes. In an AI accelerated threat landscape, the value of MDR is measured by how quickly threats are contained, how consistently risk is reduced, and how clearly progress can be demonstrated over time. Look for the following:



End to end detection, investigation, and response that is AI led and human validated

AI should advance investigations at machine speed and operate in parallel, while experienced security professionals retain oversight, judgment, and accountability. Alert forwarding alone does not meet this bar. MDR must actively disrupt threats, not simply notify customers when something happens



Reduce noise and deliver clear response actions

High volumes of raw alerts slow down teams. Modern MDR replaces alert floods with business relevant incidents that clearly indicate what matters and what to do next. Faster detection only creates value when it leads to faster, more confident action.



Real response authority, not just escalation

Effective MDR services take preapproved response actions on behalf of the organization and remotely contain threats and guide remediation. When response depends entirely on a customer's action, dwell time increases and security outcomes suffer.



Visibility across the full attack surface

Modern attacks move fluidly across endpoints, identity systems, email, cloud services, networks, and applications. MDR must correlate telemetry across these domains to detect multistage attacks and eliminate blind spots attackers will exploit.



Governed, trustworthy AI that doesn't increase risk

Decisions made at machine speed require transparency, validation, and human oversight. AI that cannot explain how conclusions are reached or how failures are handled introduces new risks.



Continuous security posture improvement

Modern MDR must continuously improve an organization's security. Proactive hardening, exposure reduction, and ongoing guidance are essential. CISOs should experience measurable reductions in attack frequency and impact that collaborate with their MDR provider, not just faster response to the same incidents.

These capabilities define the difference between AI led MDR that delivers lasting value and services instead of simply automating pieces of an outdated security operations model.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

06

AI-Led MDR Evaluation Framework

Evaluating MDR providers requires looking past feature lists. The right questions focus on how the service operates day to day, how decisions are made, and whether security outcomes improve over time. This framework is designed to help security leader separate modern, AI led MDR services from legacy offerings that rely on alerting and escalation and AI that will amplify bottlenecks.

- *How does your offering deliver security operations? What am I responsible for and what are you responsible to execute?*
- *How is AI used? Where are humans involved and where is AI making decisions?*
- *How are AI decisions validated, and where does human oversight apply?*
- *How are AI models trained? With real world operational data or synthetic data?*
- *What telemetry does the MDR solution ingest and do I have a choice to use any tool I choose? Do you cover [insert most important data sources]?*
- *Does the MDR solution limit the volume of data ingested or make unlimited data ingestion cost prohibitive?*
- *Does the service reduce alert volume, or simply forward alerts?*
- *How is risk reduction or security posture improvement measured over time?*
- *What is the model for continuous hardening and proactive guidance?*
- *What financial or operational accountability exists if outcomes fall short?*

These questions surface whether MDR is truly managed, or whether responsibility still sits largely with the customer, how much the vendor may limit your choice of what telemetry it will use for security operations, whether unlimited ingestion is cost prohibitive and how comfortable you are with the way AI is used.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

06 AI-Led MDR Evaluation Framework continued

EVALUATION AREA	WHY IT MATTERS	BEST-PRACTICE BENCHMARK
Security Operations Model and AI	The legacy-tiered approach to SOC workflows can't keep pace with AI-accelerated attacks, turning speed gaps into business risk	AI-led SOC operating model where investigations run in parallel, AI drives execution at machine speed, and humans oversee decisions and security outcomes
AI governance and validation	Unchecked AI introduces new operational risk	Understand AI training, validation, oversight, and accountability
Outcome reporting and accountability	CISOs must demonstrate value to executives and boards	Clear metrics tied to risk reduction and ROI
Proactive posture improvement	Reactive response alone does not reduce future risk	Continuous posture reviews, hardening guidance, configuration improvements, and measurable risk reduction
Telemetry coverage and control	Gaps in visibility create blind spots attackers exploit, control determines adaptability over time	Ingests customer chosen telemetry across endpoint, identity, email, cloud, network, and applications without forcing vendor lock-in
Detection and investigation	Alert volume does not equal security outcomes; speed, accuracy, and context reduce dwell time	AI accelerated investigations that reduce alert noise and deliver business relevant incidents, validated by experienced security professionals
Response authority and execution	Delays increase business risk and escalation only models slow containment	Provider takes pre approved response actions remotely and delivers clear, prioritized remediation guidance
Tailored Security Outcomes	Security is not a one-size fits all and every organization is different	Customer context is accounted for across security operations, so the outcomes are tailored for your business based on the provider's knowledge of your business

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

07
Effective AI-led MDR and Arctic Wolf

Effective MDR is no longer defined by faster alerting or incremental automation. It is defined by whether security operations can keep pace with attackers that already operate at machine speed. Keeping pace requires a different operating model, and Arctic Wolf MDR delivers it.[3]

A proven, trustworthy AI powered Agentic SOC provides the critical foundation for AI-led MDR. AI only delivers value when it is trained, tested, and validated in real operating conditions. Arctic Wolf's Aurora® AI is continuously stress tested inside one of the world's largest commercial security operations centers, across thousands of active customer environments. Before deployment, AI capabilities are designed to outperform human only benchmarks, demonstrating faster investigations, higher detection fidelity, and more consistent outcomes. Human oversight, reinforced learning, and decision validation remain built in, not bolted on.

The ability to trust AI is inseparable from how the SOC is designed. To keep up with attackers, the SOC must eliminate bottlenecks, accelerate response, and improve consistency without sacrificing accountability. Many MDR providers still automate Tier 1 and Tier 2 tasks inside legacy SOC workflows, preserving queues and human bottlenecks. Arctic Wolf redesigned the SOC from the ground up for AI led investigations. Instead of sequential handoffs, AI agents operate in parallel across detection, correlation, and investigation, while more than a thousand security professionals focus on validation, escalation, and refinement. For example, Arctic Wolf's Swarm of Experts integrates over 300 specialized, artificial intelligence agents into a fully automated Agentic SOC.



©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

07

Effective AI-led MDR and Arctic Wolf continued

Visibility is another defining requirement. Effective MDR must work across the customer's entire environment, not force consolidation into a closed ecosystem. Arctic Wolf's Aurora Superintelligence Platform ingests telemetry from customer-chosen tools across endpoint, identity, network, cloud, email, and applications. This open approach avoids vendor lock-in, preserves existing investments, and reduces operational burden by centralizing data, detections, and insights into a unified foundation, especially if unlimited ingestion is offered without extra cost.

Outcomes matter more than activity. Arctic Wolf MDR is designed to deliver measurable reductions in attack frequency and impact, supported by real world detection performance, full attack surface visibility, and continuous validation. Between incidents, security posture improves through proactive hardening, structured guidance, and Security Posture in depth Reviews (SPiDRs) delivered by a dedicated Concierge Security® Team. The goal is not just a faster response, but a safer environment over time.

Accountability completes the picture. Security leaders are expected to defend security decisions and explain business risks. Arctic Wolf backs its MDR outcomes with financial accountability, including an industry leading Security Operations Warranty. Arctic Wolf's accountability shifts meaningful operational and financial risk away from the customer's organization. Customers receive predictable costs, flexible telemetry ingestion, and Agentic SOC coverage equivalent to multiple full time security professionals.

The result is confidence: AI is working safely, effectively, and fast enough to stymie attackers. Organizations gain clear visibility and control across their entire attack surface with Arctic Wolf's Agentic SOC working to defend. Our MDR solution goes beyond reacting to threats to continuously reduce risk. Arctic Wolf is a security operations partner that stands behind our work with security outcomes and risk transfer.

©2026 Arctic Wolf Networks, Inc. All rights reserved. | Public

08
Conclusion

Security operations is changing. It has to change. Attackers are already operating at machine speed, using automation and AI to compress timelines and increase impact. Many security teams, however, are still relying on operating models built for a slower, more manual threat landscape.

AI is a force multiplier but only when applied to decrease bottlenecks. Faster alerting alone does not improve security outcomes. Automating pieces of legacy workflows does not remove bottlenecks. And AI without transparency or oversight introduces new risks instead of reducing it.

Effective MDR is defined by creating a partnership with a customer to take on work of a full SOC, is architected to use AI effectively, and is accountable for their security outcomes. Detection, investigation, and response must operate as a single, integrated capability. AI must accelerate scale and speed, while humans retain judgment and control. Security posture must improve between incidents, not just during them. And outcomes must be measurable, repeatable, and defensible at the executive level.

To learn more about Arctic Wolf's AI-led MDR, please visit our **Aurora Managed Detection and Response webpage** for demos, datasheets and how our MDR solution works.

ABOUT ARCTIC WOLF

Arctic Wolf is the cybersecurity and AI company that ends cyber risk by transforming it into business resilience. Powered by the Aurora® Superintelligence Platform, Arctic Wolf delivers modern security operations built on proprietary AI and decades of real-world expertise. By combining AI-driven automation with expert validated precision, Arctic Wolf helps organizations confidently manage cyber risk — so organizations can operate with control and the freedom to innovate.

To learn more about Arctic Wolf, visit **arcticwolf.com**.

FOOTNOTES

1. Mandiant, M Trends 2026 Report, "Initial access to handoff shrinks to 22 seconds," 2026.
2. Gartner, Market Guide for Managed Detection and Response, October 2025.
3. Arctic Wolf, Essential Guide to the Agentic SOC, 2026

©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, and Arctic Wolf Concierge Security Team 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.

This Market Guide may include forward looking statements. These reflect our current views and are subject to change. They are not guarantees, and actual results may vary. This document is for informational purposes only and reflects current views on security operations and AI. It is not a guarantee of performance. Actual results may vary based on customer environment, configuration, and implementation. AI outputs and outcomes may vary and should be validated within each organization's environment.

AW_G_MODERN MDR_0726