Qevlar AI

for SOC leaders

AI SOC Solutions Comparison

Buyer's Guide for choosing the right AI SOC

May 2026

Executive Summary

Choosing the right AI SOC solution is one of the most consequential investments a security operations leader will make in 2026.

The market now spans three distinct solution categories, all used across Fortune 500 companies and MSSPs:

1. Building your own AI SOC Solution
2. Standalone AI SOC Agents or Platforms
3. AI capabilities embedded within existing detection tools (SIEM, XDR, SOAR, etc.)

The differences between them go far beyond feature lists. We explore the pros and cons of each category and compare AI SOC Agents head-to-head with AI capabilities in incumbent platforms.

This guide offers a practical approach: the comparison criteria in this report are informed by Gartner's "Solution Criteria for Detection and Response AI SOC Agents" (February 2026) and real-world evaluations conducted by MSSPs and global enterprises throughout 2025 and 2026, in which Qevlar AI was part of as an evaluated solution.

You can apply these evaluation criteria to your process by narrowing down the comparison list to what matters most to your SOC or expanding it to include extra requirements.

The Three Options: Understanding the Landscape

Option 1

Build Your Own AI SOC Solution

Some mature security teams choose to build their own AI SOC capabilities using LLMs, internal data pipelines, and custom investigation workflows. This approach offers maximum flexibility and control, and can be tailored precisely to your environment. In practice, most teams significantly underestimate the complexity involved.

Product challenges

A strong engineering team can build basic alert enrichment or triage logic.

What is much harder is building a full investigation and response system that consistently delivers accurate, explainable, and reproducible results in production.

Questions you need to answer before building an AI SOC your team can trust:

- How will we establish investigation logic, including knowing which signals matter, how to correlate them, and how to dynamically expand an investigation as new indicators emerge?
- How can we ensure determinism at scale to avoid drift, hallucinations, and inconsistent verdicts across identical alerts?

> *Qevlar AI ran an experiment by sending alerts 100 times to an LLM to investigate. Alerts were given different severity ratings and threat classifications for the exact same inputs. Study results*

- How will we mai ntain accuracy by reducing false positives without increasing false negatives?
- How can we achieve explainability to produce outputs analysts can trust, validate, and defend to customers or auditors?
- What is our strategy for reliably querying and correlating data across multiple tools at scale?
- How will we handle continuous improvement by incorporating analyst feedback, organisational context, adapting to new threats, and maintaining performance over time?

Economics challenges

Even if you can build it, the economics rarely work in your favour. Maintaining and evolving a production-grade AI SOC system is a continuous, multi-year engineering commitment requiring dedicated ownership and deep domain expertise.

Teams that have started down this path consistently report that the effort to reach production-grade reliability — and maintain it — far exceeds initial estimates.

Option 2

AI Embedded in Detection Tools

Embedded AI features are designed to work within a single vendor's ecosystem, representing a vertical adoption approach.

This model can work well for organisations built on a specific consolidated vendor stack — particularly where the AI features cover the alert types that dominate the team's queue.

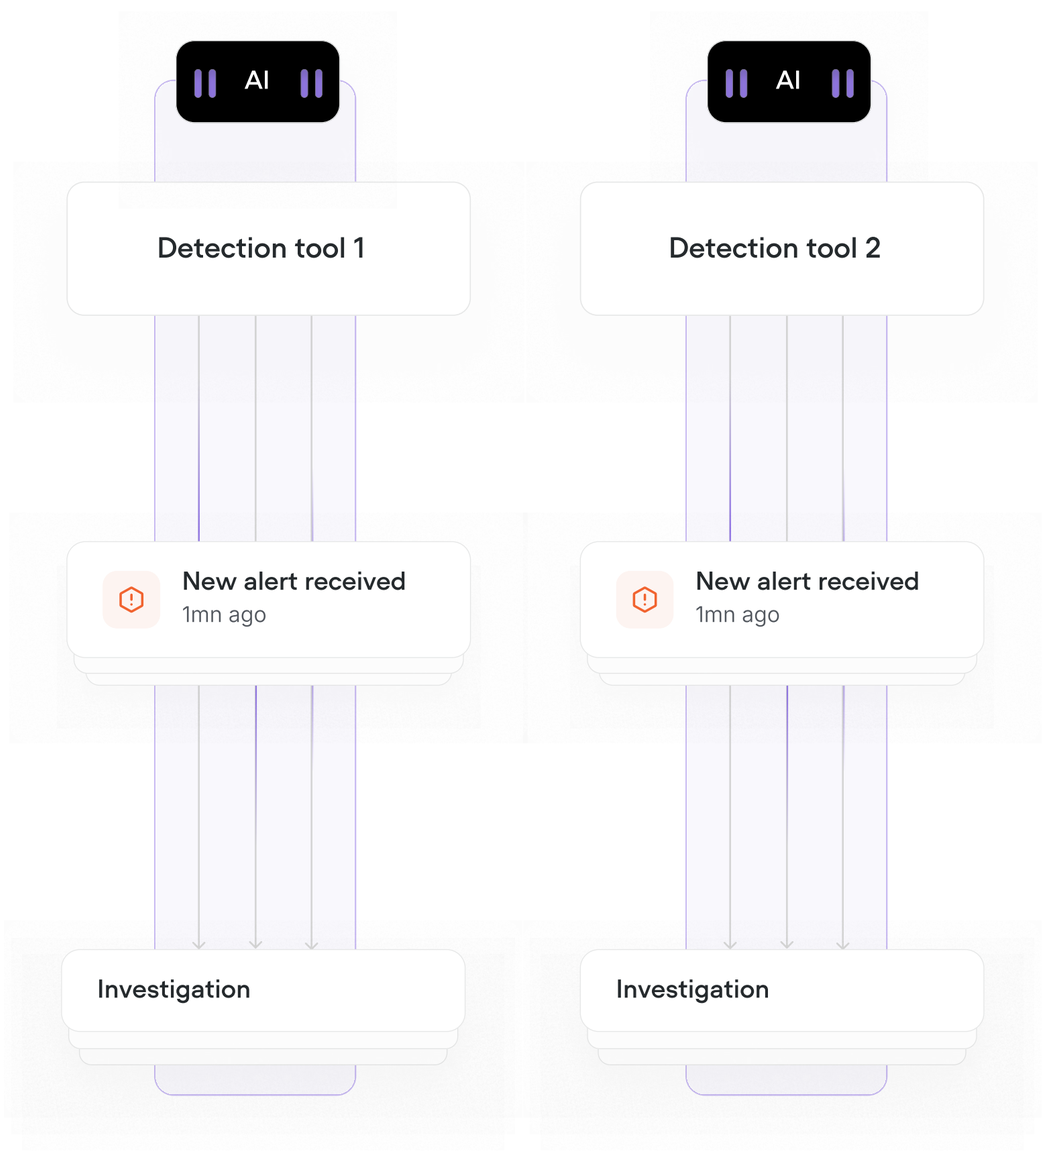

Be cautious to choose this approach if any of the following apply to your environment:

- You operate a heterogeneous, best-of-breed security stack;
- you are considering changing or expanding your toolset;
- you want to investigate and respond at the correlated incident level rather than alert-by-alert;
- you require full auditability and transparency of every AI reasoning step.

Option 3

Standalone AI SOC Platforms

For organisations with a heterogeneous security stack who want to automate operations across teams and tools, purpose-built AI SOC platforms are the better fit.

They work horizontally across your entire security stack as a vendor-agnostic layer, correlating signals, data and alerts across tools.

Be cautious when comparing vendors within this category: the market is crowded, and not every solution is proven in production.

The criteria in this comparison guide are specifically designed to help you separate solutions that perform under real SOC conditions from those that do not.

Evaluation criteria for AI SOC Platforms and AI in detection tools

We have accumulated and structured comparison criteria provided by Gartner's "Solution Criteria for Detection and Response AI SOC Agents" (February 2026) and real-world evaluations conducted by MSSPs and global enterprises throughout 2025 and 2026.

Each row in the comparison describes a capability dimension, explains what it means in practice for your team, and compares how each solution category addresses it.

The Qevlar AI column is highlighted as a named example of a leading purpose-built AI SOC platform; other vendors in this category may differ on individual criteria. All other vendor references are anonymized.

Disclaimer:

The AI SOC market is evolving rapidly. Some data points may change as vendors release new capabilities. Always validate the current product state hands-on during evaluation.

Criteria	What It Means for Your SOC	Qevlar AI	AI SOC Agents	AI in Detection Tools
1 · Stack Fit				
Stack Independence	Can the AI SOC solution work alongside all your existing tools?	• A **Vendor-Agnostic** Connects to your existing SIEM, EDR, email, identity, network, or cloud tools.	• A **Vendor-Agnostic** Typically, AI SOC agents are vendor-agnostic and connect via REST APIs to any tool.	• C **Vendor-Locked** AI features are tied to the vendor's own product ecosystem. Investigations are limited to data sources within that platform.
Hosting and depoyment model	Does the solution support SaaS, on-premises, or Bring Your Own Cloud (BYOC) deployment to meet your data sovereignty and security requirements?	• A **SaaS + BYOC** Deploy on Qevlar's SaaS or in your own cloud (BYOC)	• B **Mostly SaaS** Most AI SOC agents offer SaaS-only deployment. On-premises and private cloud options are limited, which can be a blocker for regulated industries.	• B **Cloud-Only AI Features** The underlying detection platform may support on-premises deployment, but AI/LLM features typically require cloud connectivity.
2 · Investigations				
Deep Investigation Capabilities	Can the AI go beyond log enrichment to extract additional related artifacts and pivot across your entire connected stack?	• A **Advanced** Investigations expand beyond the alert boundary: multi-source pivoting, detection of related IOCs, uncovering authentication anomalies, correlation of alerts into incidents, and revealing the full attack scope and impact	• B **Limited** Most AI SOC agents investigate each alert using only the artifacts provided and data retrieved from the detection tool. Pivoting across connected tools to understand the broader threat scope is limited.	• C **Enrichment Only** Limited to log lookups and TI queries within the vendor's own telemetry. Deep investigation is outside the design scope of embedded AI features.
Cross-Alert Incident Correlation	Does the AI automatically link related alerts into incidents, revealing the broader attack story rather than treating each alert in isolation?	• A **Yes** Automatically correlates investigations across alerts into incidents. Detects related IOCs across the environment, revealing attack patterns that might be missed by isolated alert triage or through low-severity signals.	• B **Rarely Available** Most AI SOC agents investigate each alert independently. Alert grouping exists in a few tools.	• B **Available through static rules** Each alert is analysed within its own context. Platform-level incident grouping where available is rule-based rather than AI-driven correlation.
AI Architecture & Hallucination Prevention	What prevents the AI from producing inaccurate, or inconsistent conclusions? Is there architectural determinism, or is the output purely LLM-dependent?	• A **Graph Orchestration** The graph orchestrator plans investigations and adapts steps based on findings and context, while LLMs handle only narrow tasks. Same inputs always produce the same results.	• C **LLM-Dependent (Most)** Most agents use a single LLM or multi-agent pipeline without deterministic guardrails. Under edge-case inputs, results can be random or inconsistent.	• C **LLM-dependent (Most)** Results can be random or inconsistent

Criteria	What It Means for Your SOC	Qevlar AI	AI SOC Agents	AI in Detection Tools
Transparency in Reasoning	Can analysts see exactly why the AI reached its conclusion — every step taken, every observable queried, every source cited?	• A **Transparent** Displays every step taken, every observable analyzed, and every source queried. Analysts get complete traceability from raw alert to final verdict	• A **Transparent** Typically, display every investigation step, artifact analysed, and every data source queried.	• C **Black-box** Most embedded AI tools provide a verdict or recommendation without deep reasoning. Analysts cannot see how the conclusion was reached, making it difficult to trust, challenge, or learn from AI outputs.
3 · Adapting to your SOC				
Organizational Context	Can the AI ingest your SOPs, playbooks, CMDB data, and business context to make decisions tuned to your specific environment?	• A **Yes, with pre-deployment testing** Business-specific context is automatically applied across investigations with no training period. Context comes from user input, analyst feedback, internal policies, and AI-generated insights. Analysts can replay past investigations to test new context before deployment and see how context shaped decisions.	• B **Yes, without testing** Most agents offer some form of context onboarding. Approaches range from file ingestion (SOPs, runbooks) to structured CMDB integration. Some require training periods to apply context reliably. The ability to test new context items before applying them to future investigations is not available.	• C **Stack-Bounded** Leverages only the context available within the vendor's own security stack. External SOPs, custom runbooks, or asset classifications from other systems are typically not factored into AI analysis.
Historical Context	Does the AI factor in previous alerts, incidents, and investigation outcomes for the same user, asset, or attack pattern — to improve current decision-making?	• A **Yes** Factors in past investigation outcomes and evidence and pulls past tickets directly from ITSM for additional context.	• B **Few Vendors** Some agents check the last few related alerts for the same entity when analysing a new case. Deep historical correlation across investigation history is offered by only a small subset of the market.	• C **Platform History Only** The AI can reference the historical data held within the vendor's own platform. Cross-tool historical context — past alerts from a different SIEM or EDR — is not accessible.
Analyst Feedback Loop	Can analysts correct or validate AI verdicts, and does the system learn from that feedback to improve future investigations for your environment?	• A **Yes** Analysts can provide verdict feedback to refine conclusion logic. Feedback is automatically turned into context for future investigations, which can be validated before deployment, with expected impact assessed.	• A **Yes, with limitations** Most mature pure-play agents support analyst feedback on verdicts. Mechanisms range from thumbs-up/down ratings to detailed text corrections. No clear evidence of the feedback being converted into reusable context.	• A **Yes** Most AI capabilities support analyst feedback on verdicts

Criteria	What It Means for Your SOC	Qevlar AI	AI SOC Agents	AI in Detection Tools
Customizable Conclusion Logic	Can your team configure thresholds, escalation rules, and dismissal logic that shape how the AI concludes an investigation — rather than accepting vendor defaults?	• A **Yes** A no-code rule builder for outcome classification and response actions, configurable to your processes and risk tolerance.	• A **Offered by Some Vendors** Custom strategies and configurable conclusion rules are available in more mature pure-play agents. This capability allows the AI to reflect the SOC's own risk tolerance and escalation policies	• C **Not Available** Investigation conclusions follow vendor-defined logic with no mechanism for customers to adjust AI decision thresholds, dismissal criteria, or escalation rules
Context the AI Builds Itself	Can the AI propose new context items itself, by learning from your environment?	• A **Yes** Suggests new context items based on recurring patterns surfaced across investigations, and routes them to your team for review and approval before they're applied.	• B **Few Vendors** Available in a small subset of players; most rely on context added manually by the SOC team.	• C **Not Available** Auto-generated context is not part of embedded AI feature sets today.
4 · Response & Automation				
Remediation Guidance & Response Actions	Does the platform suggest or execute remediation steps — from recommended next actions to automated containment (isolate endpoint, revoke session, block IP)?	• A **Guided + Automated** Provides investigation and remediation guidance with recommended next steps in line with your internal policies. Supports automated containment actions that can be executed at the incident level to cover the full threat scope.	• A **Guided + Automated** Most agents provide remediation guidance and recommended next steps. Automated containment actions (quarantine, user suspension) are available in a growing subset of vendors	• B **SOAR-Routed** Response actions are mainly routed through the vendor's SOAR module or require analyst prompting
Automation Workflows	Can investigation outcomes trigger automated downstream workflows — ticket creation, escalation, team notifications — without manual analyst steps?	• A **Yes** A no-code workflow builder for outcomes classification and response actions, configurable to your processes and risk tolerance.	• A **Offered by Some Vendors** Configurable automation workflows are available in some pure-play agents. Custom response playbooks triggered by AI verdicts vary in scope and ease of configuration	• B **Platform Automation** Automation relies on the underlying platform's orchestration engine or building own automation agents

Criteria	What It Means for Your SOC	Qevlar AI	AI SOC Agents	AI in Detection Tools
5 · Broader SOC Workflow Enablement				
Detection Engineering	Does the AI help your detection engineers identify coverage gaps, tune existing rules, and create or improve detections more efficiently?	• A **Supported** Identifies noisy rules, suggests tuning and adding new ones to increase coverage across SIEM/EDR/ cloud stack	• B **Limited** Some agents identify noisy rules. Detection tuning suggestions are uncommon in this category.	• A **Supported** Detection-native platforms often provide AI-assisted rule creation and tuning as a built-in capability. You can also create your custom agent for detection engineering
Threat Hunting	Can the AI assist threat hunters with hypothesis generation, automated query creation, and proactive cross-source hunting beyond incoming alert queues?	• B **Emerging Capability** Autonomously runs hunts triggered by behavioral anomalies. IOC and vulnerability exploitation hunting is on the roadmap.	• B **Emerging Capability** AI-assisted threat hunting with hypothesis creation and query generation is available in a few pure-play agents.	• B **Emerging Capability** Some vendors offer AI agents for threat hunting focused on natural language queries and responses. This remains analyst-driven and does not provide autonomous, cross-source threat hunting.
Vulnerability Management	Does AI help connect security incidents with vulnerability management and prioritize risk?	• A **Yes** Connects vulnerabilities with real-world threats, linking CVEs to active exploitation and security incidents. It adds asset context and prioritizes risk so teams can act faster.	• C **Not available**	• B **Separate capability** Some vendors offer standalone AI agents for vulnerability prioritization, though these are usually not integrated with SOC agents or incident workflows.
6 · Enterprise & MSSP Readiness				
Multi-Tenancy & MSSP Support	Does the platform support true multi-tenancy — with per-tenant configuration, client-specific organisational context and dedicated MSSP reporting?	• A **Yes — Built for Both Enterprises and MSSPs**	• B **Varies — Enterprise-First** Some AI SOC agents serve only on enterprise customers	• A **Yes**
Proven in Production	Has the solution been validated in real production environments — not just in POC — by enterprises and MSSPs who trust it to automate parts of their SOC operations at scale?	• A **Deployed across 1,500+ Companies** Field-tested and trusted in production by Fortune Global 500 companies and leading MSSPs across 10 countries.	• B **Varies** The pure-play AI SOC agent market spans vendors at very different maturity levels. Production validation should be rigorously verified.	• B **Proven platforms, AI still maturing** The underlying detection platforms have large, mature enterprise deployments. However, the specific AI triage and investigation features are often significantly newer than the core platform and may not have the same production track record.

Ready to see AI SOC in action?

Qevlar is the AI SOC platform for self-improving defense. It investigates every alert across your entire security stack and turns each outcome into intelligence that strengthens response, detection engineering, threat hunting, and vulnerability management.

Every case compounds into knowledge the next one can use,
so your defenses get stronger the longer it runs.

Contextual intelligence

Compounding learning from every case

Org. context · Behavioral baselines · Historic cases · Cross-team intel · Compounding learning

#1 **Investigation & Response**

- Autonomous investigations at both alert and incident levels
- One-click containment
- Customizable workflows
- Transparent, evidence-backed report in a console of your choice

#2 **Detection Engineering**

- False-positive tuning suggestions
- Auto-tuning of noisy rules
- New detection rules

#3 **Threat Hunting**

- Behavioral anomaly hunting built in investigations
- Hunting for CVEs & IOCs

#4 **Vulnerability Management**

- Prioritization with sprint-ready remediation
- CVE-aware incident prioritization
- Proactive vulnerability exploitation hunting

Live in production in 1,500 organizations globally

Cyberdefense · Atos · sodexo · Fortune Global 500 · ECI · MediaMarkt · i-TRACING · stoïk · GlobalConnect · nomios · almond · PERFORMANTA

See how it works →