SACR

JULY 2026

THE AI SOC TECHNOSCOPE SERIES: THE AI SOC MARKET, 2026

SEAN SOSNOWSKI

7AI

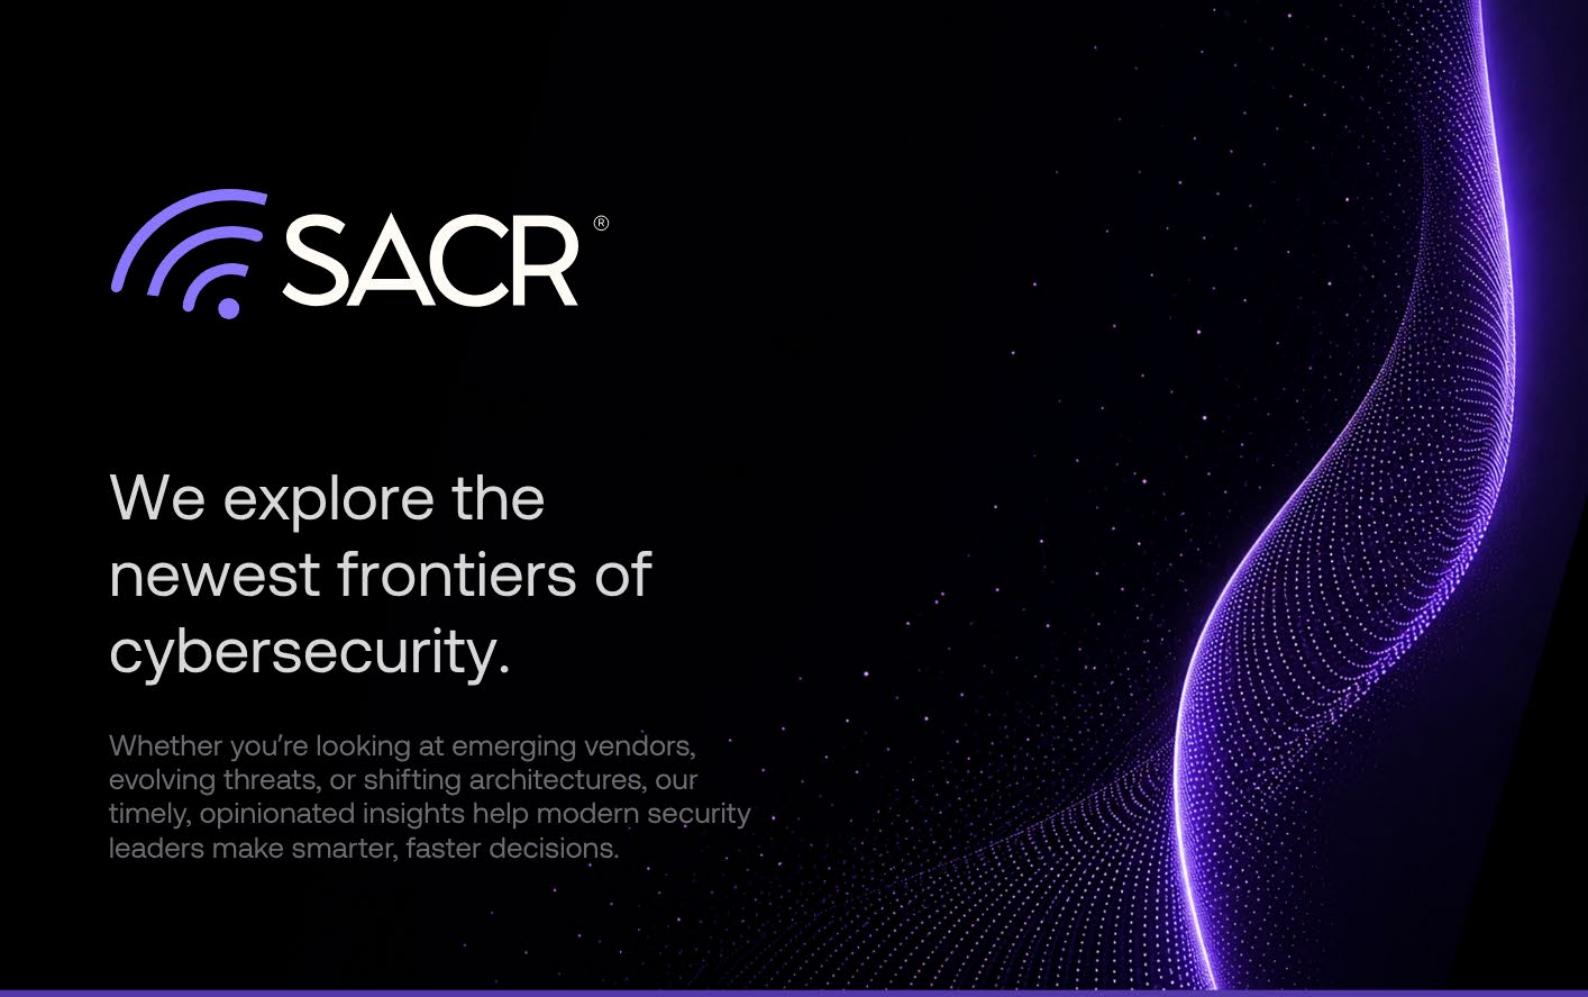

We explore the newest frontiers of cybersecurity.

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 Us

Software Analyst Cyber 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.

Acknowledgements

Practitioners and CISOs,

We're excited to share our latest research on the AI SOC market. This report explores how organizations can evaluate AI SOC platforms and build a more trusted approach to security operations.

The Author

Sean Sosnowski serves as the Research Director for Security Operations and Cloud Security at SACR, where he leads research on SOC strategy and operations, detection engineering, and the evolving role of automation and agentic AI in security workflows. Drawing on a decade of intelligence experience in the U.S. Marine Corps he has authored several analytic reports on emerging technologies and threats regarding sensitive national security and military operations.

Table of Contents

This is SACR's third consecutive year of original research on the AI SOC market, following previous episodes: The Path Toward AI-Augmented SOCs (2024) and SACR AI SOC Market Landscape for 2025. This year's research inaugurates a new franchise, the AI SOC Technoscope Series, which opened with Part 1, Building the Trusted SOC.

Part 1 established that a Trusted SOC is not purchased as a platform or achieved through autonomous actions by AI agents. Security teams/leaders build it by granting AI agents authority for specific response actions only when the evidence, governance, reliability, and verified outcomes support that authority.

This Part 2 applies that model to the market itself. We evaluated 18 vendors across three operating environments and testing scenarios.

1. The regulated SIEM-centric enterprise SOC
2. Hybrid mid-market SOC
3. Cloud native data-lake SOC

The result shows which products are positioned to support which authority portfolios, and where.

Caveats:

The AI SOC category is broad and still consolidating. Peer trackers count close to 140 vendors that use, market, or support AI within security operations, and SACR treats more than 60 of those as pure-play AI SOC platforms worth direct comparison. This report evaluates 18 of them in depth. This report is purely independent, with no vendor relationship influencing any rankings. Everything was weighted

That narrower set was not a matter of convenience. To qualify for ranking, a product had to materially shape the path from case understanding to response, through decisioning, action recommendation, governed execution, or verification, backed by enough generally available capability and production evidence to support a defensible comparison rather than a restated pitch deck. Inclusion in the ranked 18 is a statement about evaluability, not a statement about quality. A vendor's absence from the ranked set is not a judgment against it. It typically means the evidence available to us during the research period did not yet support a comparison at the same depth as the vendors we ranked.

This analysis builds on the same primary research described in Part 1: structured interviews with [20] security leaders and practitioners, and briefings or live demonstrations across the broader vendor landscape.

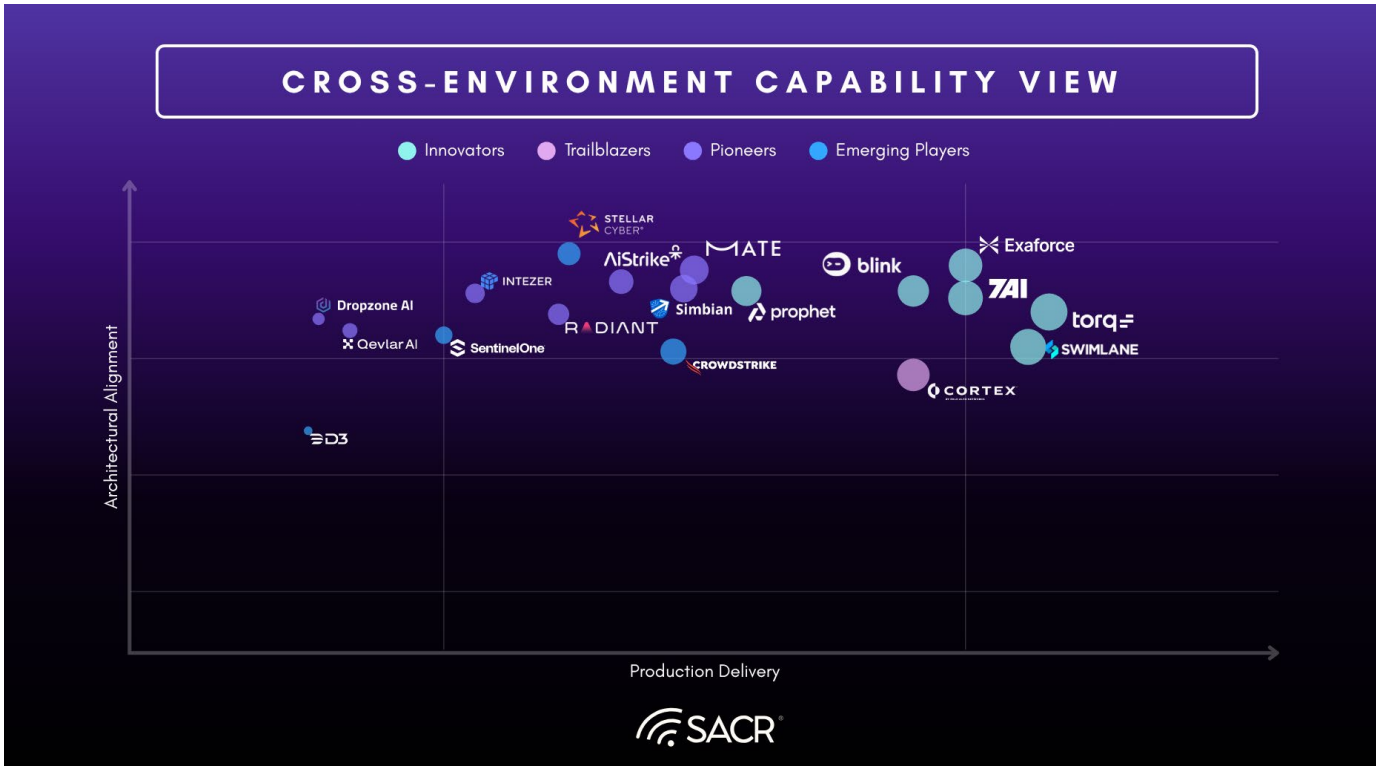

Key Actionable Summary

The first report identified the action gap between incident understanding and trusted response. Our market analysis finds that this gap remains the dividing line between broad capability and category leadership. Investigation, summarization, and enrichment have become common across the market. The systems that carry evidence into decisioning, governed action, execution, verification, and proof remain far less consistent. Products that appear similar in a demonstration can create substantially different operating models in production.

- **Evaluation model:** We evaluated 18 vendors across Architectural Alignment and Production Delivery. Architectural Alignment measures how completely a product connects the lifecycle from evidence through verified outcome. Production Delivery measures how reliably it provides that capability in real operating environments today.
- **Market finding:** Leadership is architecture-neutral. AI-native, SOAR-derived, and platform-consolidated products can all support the Trusted SOC. Their strengths differ according to lifecycle ownership, governance, production maturity, integration requirements, and operational burden.
- **Top-line results:** No single vendor can be considered a universal AI SOC leader. Suitability materially changes based on evidence architecture, control-plane ownership, governance intensity, operating capacity, and the response actions organizations intend to delegate in their own environments.
- **Environmental finding:** Vendor fit changes across the regulated SIEM-centric enterprise SOC, hybrid mid-market SOC, and cloud native data-lake SOC. Each environment places different demands on governance, integration, data access, execution, and operational overhead.
- **Buyer takeaway:** Use the Trusted SOC as the destination and these rankings as a guide to the most credible path for your environment. Evaluate how much of the lifecycle the product can own today, what authority it can safely exercise, how it proves the outcome, and what your organization must supply around it.

What Buyers Are Purchasing

AI SOC purchases today take one of three primary forms. Each places control of evidence, case state, decisioning, and response in a different part of the security stack. The practical choice is where the buyer wants this operating layer to reside, which systems should remain authoritative, and how much integration, maintenance, and governance burden the organization is prepared to own.

Path 1: Add an AI-Native Layer Over the Existing Stack

Organizations choosing an independent AI-native platform generally want to preserve their current SIEM, EDR, identity, cloud, and workflow tools while adding a system that can assemble evidence, investigate incidents, develop cases, and coordinate response across them. This path offers cross-stack flexibility, modern reasoning, and faster deployment without requiring broad platform consolidation. Its effectiveness depends on integration depth, API access, credential design,

customer context, and the platform's ability to verify actions executed through external control systems.

This path is a natural fit for hybrid mid-market and cloud native teams that operate fragmented environments, need greater analyst leverage, or want to improve investigation and response without replacing the surrounding stack. Buyers assume responsibility for deciding which systems remain authoritative and how the AI-native layer receives enough context and permission to act safely.

Path 2: Modernize SOAR and Automation with AI

Organizations with established orchestration, case management, and response workflows may extend those systems with AI-assisted investigation and decisioning. This path builds on existing strengths in connector breadth, approvals, deterministic execution, auditability, retry logic, and failure handling. AI can improve prioritization, evidence gathering, case development, playbook selection, and the handling of incidents that require interpretation rather than a fixed sequence of steps.

This path is well suited to regulated enterprises and operationally mature SOCs that already rely on formal workflows and controlled execution. Buyers gain a familiar governance foundation and may face greater workflow administration, integration maintenance, and process complexity. They should determine how deeply AI reasoning is connected to the execution layer and whether the system can preserve evidence continuity from investigation through verified closure.

Path 3: Consolidate Into a Broader Security Platform

Organizations already standardized on a major security platform may extend that platform across telemetry, detection, investigation, case management, and response. Native access to endpoint, network, identity, cloud, and threat-intelligence data can reduce integration friction and connect decisions directly to enforcement. This path can provide strong production scale, established governance, and a more unified operating experience across the platform's installed ecosystem.

This path is most natural for large or regulated enterprises that prioritize consolidation, native control, and operational consistency. Buyers should evaluate the breadth of third-party support, the amount of additional platform adoption required, and the degree to which the system can maintain case and evidence continuity outside its native ecosystem.

Choosing the Operating Model

Each path can support progress toward a Trusted SOC. The purchasing decision should begin with the authority portfolio the organization wants the product to assume: which evidence it can access, which investigations it can conduct, which actions it can recommend or execute, where approval is required, which credentials it can use, and how outcomes will be independently verified.

Sourcing and delivery choices cut across all three paths. Organizations may operate the product directly, deploy it with vendor assistance, use it through a co-managed or managed service, or combine vendor capabilities with internally built agents and automation. These choices shift the staffing, maintenance, accountability, and proof burden between the customer, the software provider, and the service operator.

This makes environment-specific evaluation essential. The buyer is selecting a product to assume defined responsibilities inside a particular security operating model, with evidence that it can perform those responsibilities safely, reliably, and at the maturity required by the organization.

Research Methodology

We developed this report through primary and secondary research conducted across the AI SOC market. We collected evidence through security-practitioner interviews, vendor briefings, live product demonstrations, targeted surveys, product documentation, and supporting market research. Practitioner research established the operating problems, architectural constraints, adoption requirements, and purchasing considerations shaping real deployments. Vendor research provided evidence on product architecture, integrations, investigation workflow, decision systems, response capability, governance, deployment maturity, and customer use.

We assessed capabilities according to what vendors could demonstrate or support with available evidence. Greater weight was given to generally available product functionality, live workflow, documented integrations, production use, and customer evidence. Beta, private-preview, and roadmap capabilities were separated from mature functionality and did not receive equivalent credit. Product positioning and general statements about automation were not treated as evidence of response authority or production maturity.

The market is developing quickly, and the available evidence remains uneven. Vendor participation, access to live deployments, customer references, and measurable production outcomes varied across the evaluated companies. The findings represent SACR's assessment of the best evidence available during the research period and should be understood as a current view of a rapidly changing market.

Evaluated Vendors

The AI SOC market now includes over 100 companies that use, market, or support AI in security operations. The ranked cohort was limited to products whose relevant AI SOC offering was publicly identifiable by December 31, 2025, which materially addressed the path from case-level understanding to governed response, which fit one of the three customer-operated product paths evaluated, and had sufficient verifiable evidence of current capability and production maturity to support a defensible comparison. This report is not meant to serve as a census of every vendor in the market, it is a comparative analysis of platforms that buyers can reasonably evaluate as a path to operating a Trusted SOC.

To select the ranked cohort we began with the broader AI SOC ecosystem, including AI-native SOC platforms, SOAR and automation providers, security-data and analytics platforms, large security suites, builder tools, and service-led offerings. We then applied a scope screen to narrow the list.

To be included in the ranked set, a product had to materially shape the path from case-level understanding to response: through AI-assisted decisioning, action recommendation or preparation, approval and policy controls, execution, verification, or a demonstrable combination of those functions. Products whose primary role is telemetry collection, data management, investigation support, or managed service delivery remain relevant to the broader market, but are not automatically comparable to response-layer platforms.

We evaluated products rather than entire companies. For broad platform vendors, only the AI SOC, SOAR, or response capabilities relevant to this

analysis were assessed. A vendor's larger security portfolio, revenue, market capitalization, or general brand presence did not determine inclusion or placement.

The evaluated set also required enough evidence to support a defensible comparison. SACR considered live workflows, product documentation, integration depth, generally available functionality, deployment maturity, customer evidence, and the distinction between production capabilities and roadmap claims.

Finally, the 18 vendors were selected to represent the principal architectural and operating choices available to buyers: AI-native SOC platforms, SOAR-derived platforms with AI, and broader security providers with identifiable AI SOC capabilities. The set is broad enough to compare those approaches across regulated SIEM-centric enterprise, hybrid mid-market, and cloud native data-lake environments without treating every adjacent product as a like-for-like competitor.

Inclusion in this report does not mean that these are the only relevant vendors in the market, neither does absence constitute a judgment that another product lacks value. It means we had a sufficient basis to evaluate these products against the specific case-to-response criteria and operating environments used in this analysis.

AI SOC Market Rankings

Cross-Environment Capability View

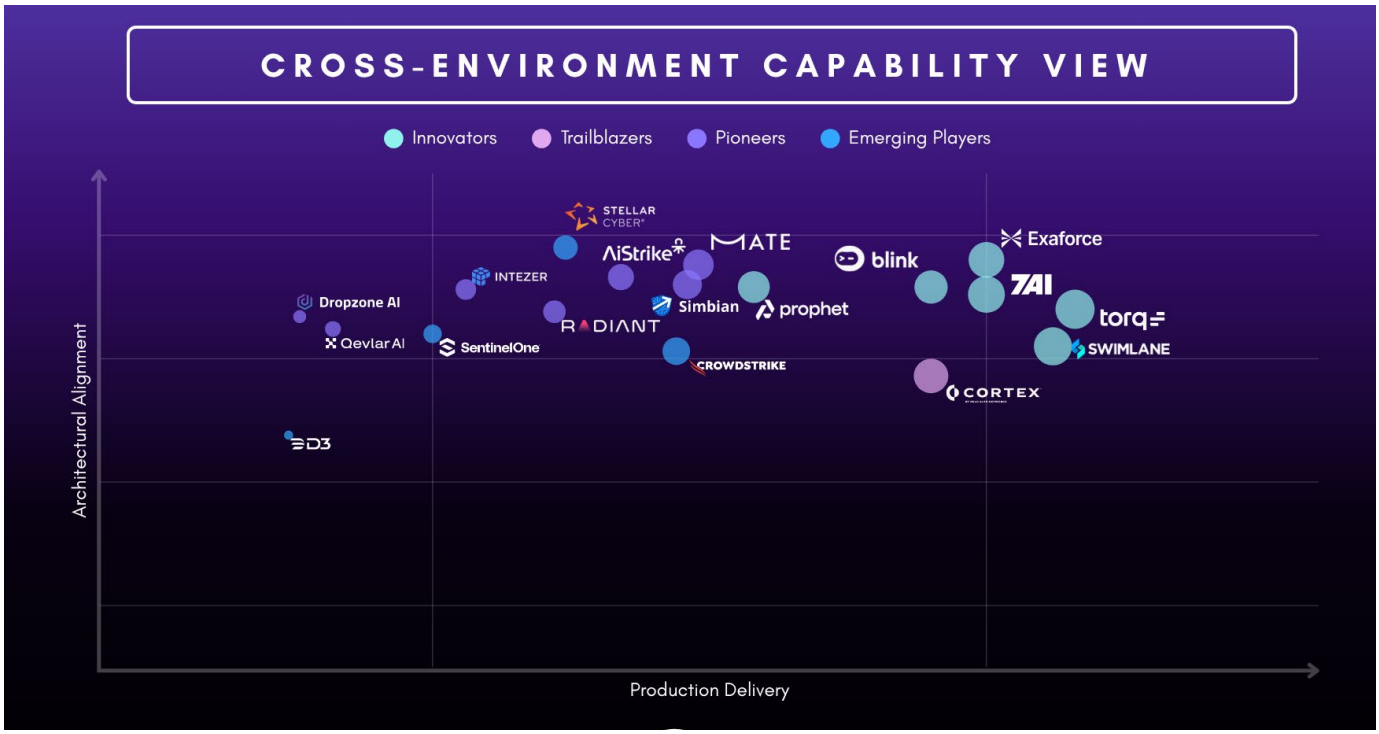

Bubble size directly correlates to the vendor's final combined average score

The comparative average consolidates performance across three operating environments. Vendor placement reflects two separate questions. Architectural Alignment measures how completely the product connects the AI SOC lifecycle from evidence through verified outcome. Production Delivery measures how reliably the product delivers those outcomes in real-world environments today. Together, the axes distinguish architectural completeness from demonstrated operational maturity.

Architectural Alignment does not reward a particular product heritage. AI-native, SOAR-derived, and platform-consolidated products can all score highly when they connect the complete lifecycle as a coherent system. Production Delivery is based on generally available capability, operational evidence, governance, scale, failure handling, and verified customer use.

Architectural Alignment

Architectural Alignment measures how completely the product's core architecture supports the AI SOC lifecycle:

- Evidence assembly
- Investigation
- Decisioning
- Bounded authority
- Response execution
- Outcome verification
- Continuous improvement

A high score means these capabilities operate as one connected evidence-to-outcome system. A lower score means ownership is distributed across the product and adjacent systems, requiring another platform, automation layer, investigation tool, or human team to complete one or more stages of the lifecycle.

Architectural Alignment measures the connectedness and ownership of the lifecycle. It is not a measure of overall product quality, company maturity, or architectural heritage.

Production Delivery

Production Delivery measures whether the product can reliably deliver AI SOC outcomes in real-world environments today. It considers:

- Production maturity
- Integration depth
- Deployment burden
- Governed execution
- Operational scale
- Failure handling
- Case continuity
- Evidence that response actions achieve the intended result.

A high score means the product can move supported incidents from investigation through controlled action and verified closure, with credible customer evidence. A lower score indicates narrower execution coverage, limited production proof, greater operational burden, or dependence on another system or human team to complete the workflow.

Innovators Category

Innovators combine strong Architectural Alignment with credible Production Delivery. They connect evidence, investigation, decisioning, bounded authority, response execution, and outcome verification within a continuous workflow, with sufficient production evidence to demonstrate that the model functions in real customer environments.

Vendors reach this position through different routes. Torq and BlinkOps emphasize automation and integration fabrics that make controlled action dependable across a wide stack. 7AI and Prophet Security emphasize AI-native case progression and graduated authority. Exaforce is differentiated by its data architecture and cloud native evidence assembly. *Swimlane reaches the category through a SOAR-derived architecture that combines federated evidence access, production AI reasoning, mature governance, and dependable cross-stack execution.* Their shared characteristic is demonstrated strength across both Architectural Alignment and Production Delivery.

Trailblazers Category

Trailblazers demonstrate strong Production Delivery with moderate Architectural Alignment. Their strengths include established integrations, durable workflows, governance, auditability, operational scale, and credible production use. Ownership of the complete AI SOC lifecycle is more distributed across the surrounding platform or automation architecture.

Palo Alto Networks reaches this position through platform consolidation, native telemetry and enforcement, and mature enterprise response controls. For buyers whose environment align with this operating model, it may provide the strongest practical fit.

Pioneers Category

Pioneers demonstrate strong Architectural Alignment with developing Production Delivery. Their products are closely organized around a connected AI SOC lifecycle, while the production record remains younger or less complete across execution breadth, deployment duration, governance depth, verification, integrations, or enterprise scale.

Mate Security and Simbian emphasize organizational context, agentic investigation, and staged authority. AIStrike connects detection improvement with investigation and response. Radiant Security and Dropzone AI emphasize analyst leverage and rapid case development. Intezer and Qevlar provide strong evidence assembly and reasoning with narrower governed-remediation depth. Broader verified execution and stronger production evidence can move these vendors toward the Innovator category.

Emerging Players Category

Emerging Players have relevant AI SOC capabilities and strong adjacent control planes, while complete lifecycle ownership and Production Delivery remain developing. Their capabilities may be concentrated in endpoint or XDR response, investigation, workflow automation, case management, or another portion of the AI SOC lifecycle.

CrowdStrike and SentinelOne bring mature native telemetry and enforcement platforms. Stellar Cyber brings an established Open XDR and case-management foundation. D3 Security brings deep SOAR, orchestration, and governance experience. Their placement reflects the current connectedness and demonstrated delivery of their evaluated AI SOC capabilities. It does not describe the maturity, market position, or overall quality of the companies.

Environment Specific Ranks

Each environment applies the same Trusted SOC destination under different operating conditions. What changes is the evidence, governance, integration, and operational burden a product must satisfy before it can credibly support response authority.

The depicted environments are meant to serve as examples of different types of SOCs. It is not an exhaustive list that captures every nuance an environment can have. Large enterprises may have fragmented security stacks, and a cloud native organization may be subject to strenuous industry regulation. We encourage buyers to leverage our analysis as a baseline of what a vendor can offer, and apply it to their own unique operating environment and regulations.

Environment 1: Regulated SIEM-Centric SOC

A mature SOC operating in a highly regulated industry where the SIEM remains the primary system of record for detection, alert correlation, investigation workflow, compliance reporting, and escalation. The mature SOC has established processes, tiered analysts, defined incident response procedures, and strict requirements for auditability, approval, and evidence retention.

Generic architecture

- Enterprise SIEM as the central detection and correlation layer
- Mature telemetry coverage across endpoint, identity, cloud, network, email, SaaS, and threat intelligence sources
- SOAR, ITSM, ticketing, and case management workflows
- Response integrations across EDR, IAM, email security, firewall, cloud controls, PAM, and vulnerability management
- Formal governance through RBAC, policy libraries, audit logs, approval gates, and evidence retention
- 24/7 or follow-the-sun SOC coverage with clear analyst tiers and incident command structure
- High regulatory pressure from financial services, healthcare, government, critical infrastructure, or public company obligations

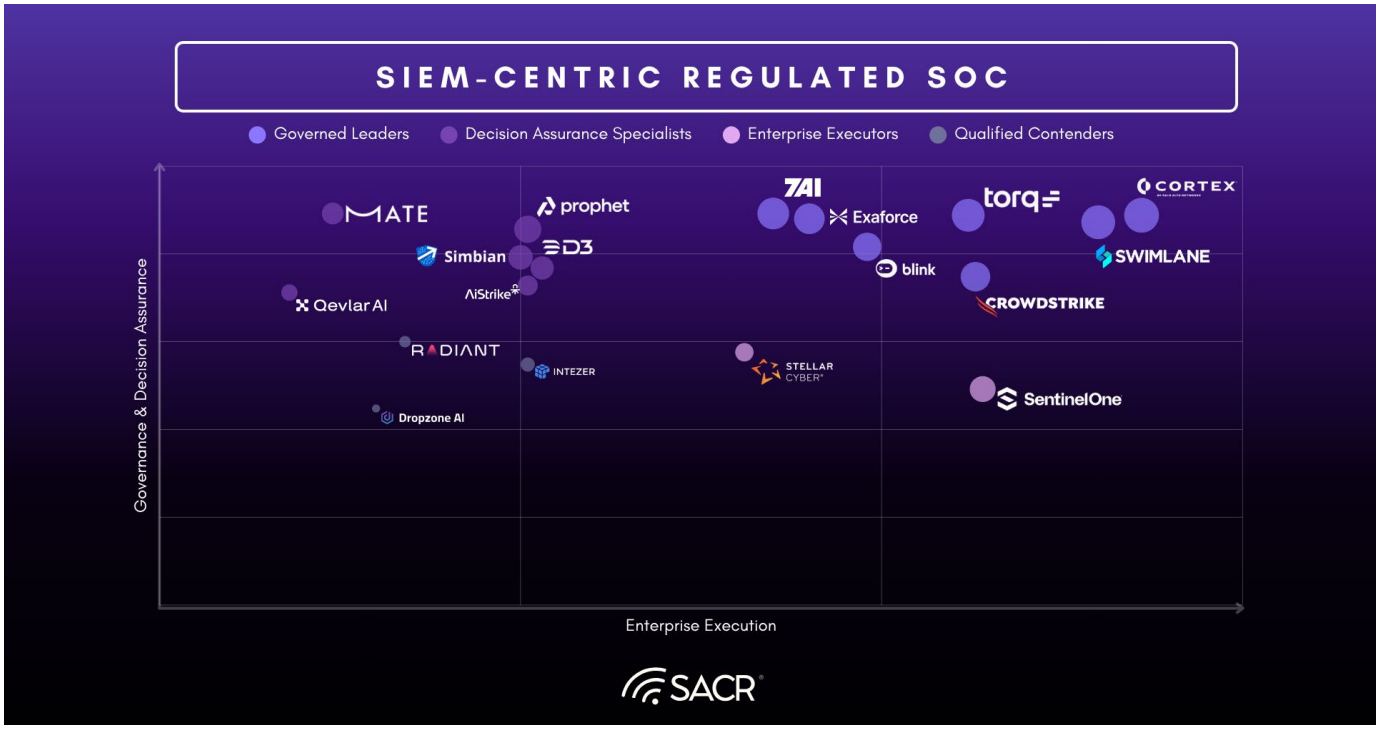

Bubble size directly correlates to the vendor's final score in this environment.

The regulated map rewards vendors that can prove both authority and accountability. Palo Alto Networks and Swimlane rise because mature platform or SOAR controls are highly valuable in this environment. Torq, 7AI, and BlinkOps remain strong because they combine modern reasoning with bounded execution. Vendors that are excellent at investigation but cannot yet show formal governance, durable case continuity, and verified post-action state are pushed toward the specialist or contender categories.

Enterprise Execution Axis

The horizontal axis measures whether the vendor can operate reliably as part of a large enterprise response system. It incorporates SIEM and workflow integration, breadth of response orchestration, deployment maturity, scale, failure handling, case continuity, and the practical ability to execute across a complex control stack.

A high score means the product can move beyond analysis into repeatable production action across enterprise systems. A lower score generally reflects narrower action breadth, less mature deployment evidence, a younger integration catalog, greater dependence on a particular architecture, or more limited proof that the product can carry enterprise operating load.

Governance and Decision Assurance Axis

The vertical axis measures whether the platform can make and govern consequential decisions in a way that a regulated enterprise can defend. It includes decision quality, evidence traceability, approval routing, policy controls, auditability, chain of custody, blast-radius limits, source-state verification, and outcome proof.

A high score means the platform can explain why an action was selected, constrain who or what can authorize it, preserve the decision record, and show what happened afterward. A lower score does not necessarily mean weak security efficacy; it often means the reasoning-to-action chain is fragmented, approval controls are incomplete, rollback is inconsistent, or post-action verification is not sufficiently mature.

Governed Leaders Category

This is the strongest position for a regulated enterprise buyer. Vendors in this quadrant combine substantial execution authority with the controls required to use that authority safely. They can support durable cases, policy-aware decisions, approvals, audit trails, and production response across meaningful parts of the enterprise stack.

The vendors are not identical. Some reach this quadrant through mature platform-native enforcement, some through SOAR and orchestration depth, and others through AI-native reasoning coupled to bounded action tools. Their common characteristic is that neither action breadth nor governance is an obvious disqualifier for a regulated deployment.

A regulated buyer should treat this quadrant as the primary shortlist, then separate vendors by operating model. Palo Alto Networks and CrowdStrike are strongest when platform consolidation is acceptable. Swimlane is strongest when mature SOAR governance and deployment optionality are central. Torq and BlinkOps are attractive when the buyer wants a more independent orchestration layer. 7AI is strongest when reasoning-led case progression and graduated autonomy matter more than long enterprise tenure.

Decision Assurance Specialists Category

These vendors show credible reasoning, transparency, policy control, or evidence assurance, but do not yet demonstrate the same breadth, scale, integration maturity, or production operating depth as the Governed Leaders. They are not weak products, their placement means that the buyer can more readily trust how they reach a decision than assume they can execute every enterprise action at scale.

This category contains several different product shapes. AI-native platforms such as Prophet, Simbian, and Mate score well because their control design is explicit. D3 Security reaches the quadrant through mature SOAR governance.

Exaforce and Qevlar bring strong evidence models. AIStrike benefits from risk-tiered authority and detection feedback. Their common limitation is not necessarily reasoning quality; it is the maturity or breadth of enterprise execution proof.

This category is useful for enterprises that already possess a strong execution layer or are willing to keep consequential action under existing SOAR, platform, or human control. These products can materially improve decision quality and case confidence even when they are not selected as the universal remediation fabric. Buyers should test integration burden, approval handoffs, action coverage, and whether the product can preserve one continuous evidence record when execution moves to another system.

Enterprise Executors Category

Enterprise Executors can perform substantial security action and have credible deployment scale, but the full reasoning-to-case-to-action-to-proof chain is less mature than their enforcement capability. These vendors often have strong control-plane assets, installed bases, and deterministic response functions. Their lower assurance position reflects fragmentation, limited dynamic reasoning, connector-dependent verification, or incomplete evidence that the newest agentic layer operates with the same governance maturity as the underlying platform.

These vendors are credible when native platform action is more important than introducing a separate response layer. The diligence focus should be on whether the product can preserve one case record, explain why a dynamic action was selected, enforce approval and credential boundaries, and independently verify the resulting state. Buyers should not assume that mature deterministic response automatically proves mature agentic response.

Qualified Contenders Category

Qualified Contenders provide meaningful operational value but do not yet meet the regulated scenario's highest bar on both dimensions. The placement often reflects a product whose strength is investigation, triage, or practical analyst acceleration rather than enterprise-wide governed remediation.

This category is not defined as an exclusion category. It identifies products that may be valuable in a narrower role, but which require additional controls, an external execution system, more human approval, or more evidence before they should become the primary autonomous response layer for a regulated enterprise.

Environment 2: Hybrid SOC

An organization with a mixed SIEM and data-lake architecture, a lean security team, uneven telemetry maturity, and a practical need to reduce analyst workload without rebuilding the SOC. The SIEM still exists, but it is no longer the only place where security data lives. The data lake or security data platform is used for broader telemetry, historical search, enrichment, and cost control.

Generic architecture

- SIEM for high-value alerts, mature detections, and compliance-relevant logs
- Data lake, security data platform, or warehouse for broader telemetry and lower-cost retention
- Telemetry from endpoint, identity, cloud, email, SaaS, network, vulnerability, and application sources
- Mixed workflow across ticketing, case management, Slack or Teams, SOAR-lite, and manual analyst processes
- Response integrations across EDR, IAM, cloud console, email security, firewall, SaaS admin tools, and ticketing
- Small internal SOC, hybrid SecOps/IT team, or co-managed MDR support
- Moderate compliance pressure, but less formal governance burden than a regulated enterprise SOC

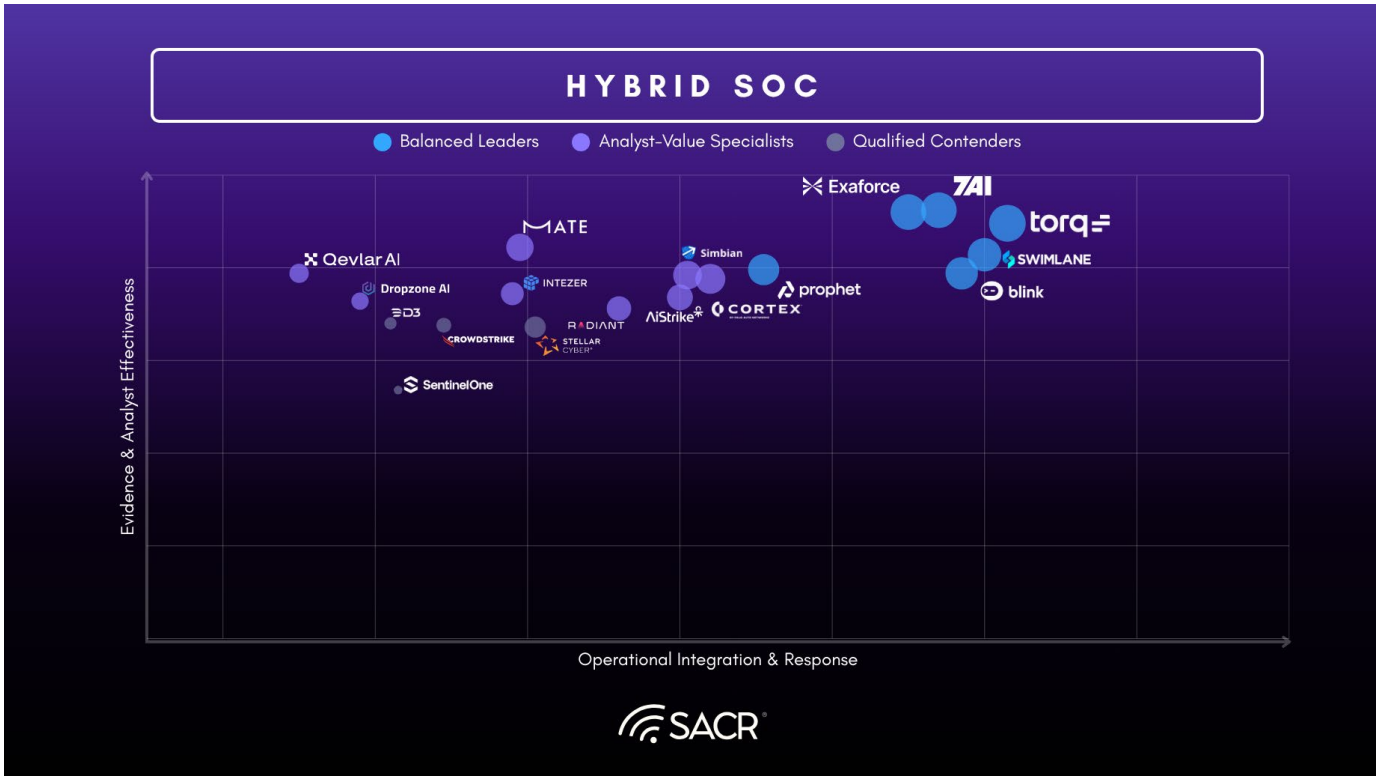

Bubble size directly correlates to the vendor's final score in this environment.

The hybrid map rewards vendors that can create a coherent case and act across a mixed stack without demanding a major transformation program. Torq and 7AI lead, with Swimlane close behind after demonstrating production operation across hybrid and federated environments. BlinkOps and Prophet Security also clear both thresholds, with different implementation, reasoning, and maturity tradeoffs. Platform-native vendors fall when their strengths require the buyer to consolidate around their ecosystem, while investigation-first vendors fall when they cannot carry the workflow through governed closure.

Operational Integration and Response Axis

The horizontal axis measures how easily the platform can connect to a mixed environment, assemble actions across systems, and deliver useful response without forcing a major platform replacement. It incorporates integration flexibility, response breadth, deployment simplicity, time-to-value, cross-stack orchestration, and the practical burden placed on a smaller team.

A high score means the vendor can sit above a fragmented stack and turn alerts or cases into a coordinated action. A lower score may reflect platform dependence, heavier implementation, narrower response authority, a requirement for substantial workflow engineering, or insufficient evidence that the product can work efficiently in a hybrid mid-market environment.

Evidence and Analyst Effectiveness Axis

The vertical axis measures whether the product improves the analyst's ability to reach a defensible decision quickly. It includes cross-system evidence assembly, case quality, transparency, contextual reasoning, productivity gains, useful recommendation, and the ability to reduce console switching and repetitive manual work.

A high score means the product creates a coherent case rather than compiling information to a queue or interface. A lower score may reflect fragmented case objects, weak cross-source correlation, limited reasoning transparency, or an operating model that still leaves substantial review and coordination work for a human analyst.

Balanced Leaders Category

Balanced leaders are the strongest fit for a hybrid environment because they combine practical cross-stack action with strong case assembly and analyst leverage. They do not require the buyer to choose between better reasoning and useful response. The differences among them are primarily in operating model, implementation weight, and maturity.

Torq, 7AI, and Swimlane form the leading hybrid shortlist. Torq and 7AI emphasize speed, flexibility, and reasoning-led operation, while Swimlane offers greater governance and automation maturity with additional administrative weight. BlinkOps remains attractive where workflow flexibility and rapid automation dominate. Prophet Security is attractive when backtesting and controlled progression into action are especially important to buyers.

Analyst-Value Specialists Category

This is the largest category in the hybrid scenario. The vendors ranked here generally improve investigation, evidence quality, or analyst productivity, but have a visible constraint in cross-stack deployment, response breadth, time-to-value, or operating-model neutrality.

The category captures several distinct reasons for falling below the Operational Integration and Response Axis threshold. Palo Alto Networks is highly capable but heavier and more platform consolidation-oriented than an independent mid-market layer. Exaforce is the most powerful when its data architecture can be used to the fullest extent. Mate, Simbian, and AIStrike are promising but have thinner production proof or integration maturity. Radiant, Intezer, Qevlar, and Dropzone provide strong analyst value but narrower governed response.

These vendors can be excellent purchases when the buyer's primary problem is investigation quality or analyst capacity rather than universal automated remediation. The question of diligence buyers should ask is where the workflow stops. Buyers should identify whether the platform directly executes, provides recommendations, triggers existing SOAR playbooks, or relies on a managed service. That handoff determines whether analyst value turns into measurable time-to-closure improvement.

Qualified Contenders Category

The vendors in this category are not low in quality compared to the others. They are less aligned to the hybrid mid-market operating model. They each carry a form of platform weight, workflow dependence, or product fragmentation that reduces both immediate analyst leverage and cross-stack simplicity for a lean team.

Environment 3: Cloud Native Data-Lake SOC

A cloud native organization that does not operate a traditional SIEM and relies on a data lake, warehouse, cloud native log store, or security data platform as the main evidence layer. The company has a small security team, an engineering-heavy operating model, and limited formal regulatory burden. Security work often happens through Slack, Jira, GitHub, cloud consoles, identity tools, and infrastructure workflows rather than a traditional SOC queue.

Generic architecture

- Data lake, warehouse, cloud native log store, or security data platform as the main evidence layer
- Telemetry from cloud logs, identity logs, endpoint tools, SaaS platforms, application telemetry, CI/CD systems, and infrastructure events
- Detection from native cloud rules, open-source detections, EDR alerts, CSPM/CNAPP findings, identity alerts, and custom queries
- Workflow through Slack or Teams, Jira or Linear, GitHub or GitLab, incident channels, and lightweight ticketing
- Response through cloud IAM, IdP, EDR, email security, infrastructure-as-code, secrets management, SaaS admin tools, and cloud consoles

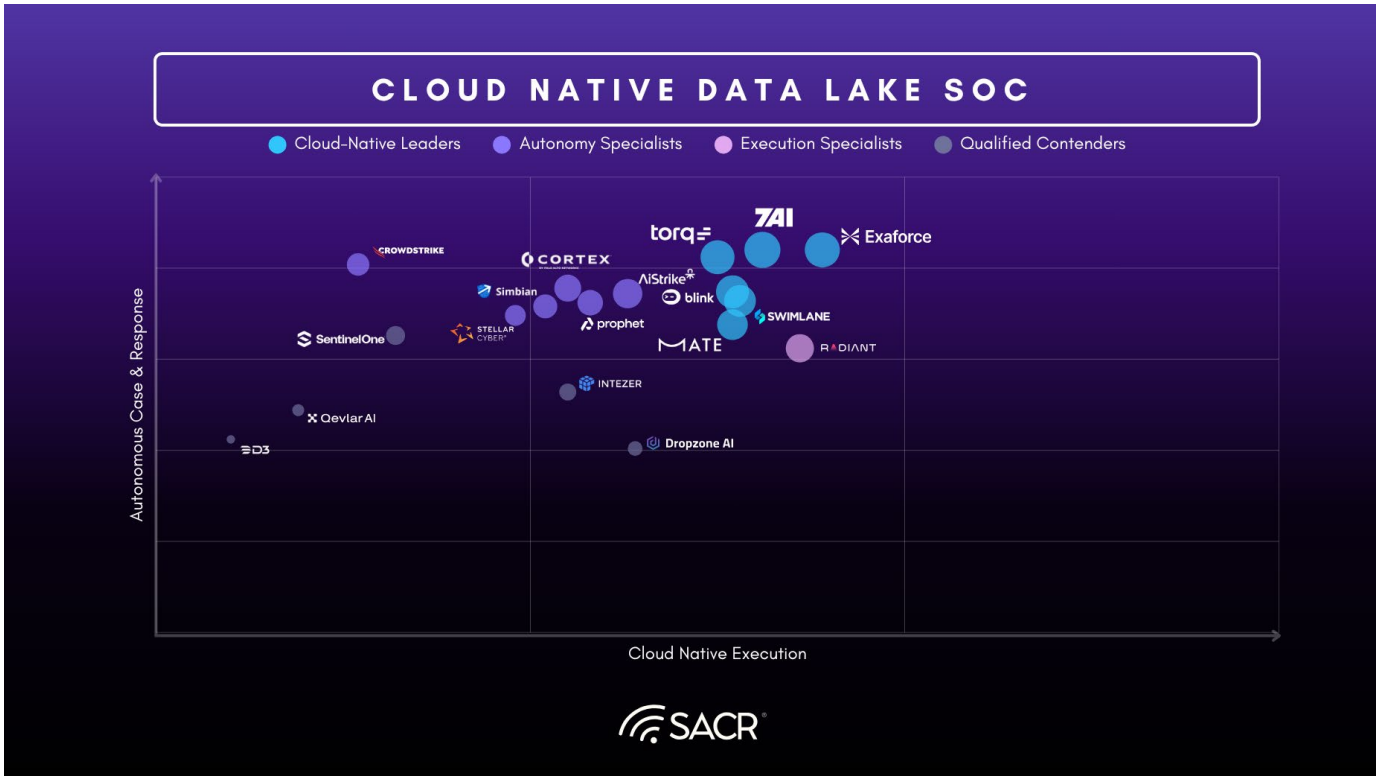

Bubble size directly correlates to the vendor's final score in this environment.

- Small security team with engineering-led response
- Low to moderate regulatory pressure, but meaningful customer trust, uptime, and breach-risk pressure

The cloud native map rewards products that can work directly with modern data and engineering context without sacrificing case construction or governed response. 7AI and Exaforce lead through evidence and data architecture, while Torq and BlinkOps lead through API-driven action and workflow flexibility. Swimlane clears both leadership thresholds by combining direct and federated telemetry access with mature case management and governed execution. Platform and other SOAR vendors remain credible, but their administrative and consolidation burden pushes them into the Autonomy Specialist category. Investigation first vendors remain contenders until they can prove broader governed execution and outcome verification.

Cloud Native Execution Axis

The horizontal axis measures the vendor's ability to operate in a no-SIEM or data lake centric architecture. It includes direct data lake access, raw cloud and identity context, engineering-native workflow integration, deployment speed, low administrative overhead, and the ability to function without requiring a traditional SOC process.

A high score means the product can become part of an engineering-led security operating model rather than merely connect to one. A lower score often reflects enterprise platform weight, SOAR-centric administration, dependence on upstream alerts, a requirement for a conventional SIEM or case process, or insufficient evidence that the product can work effectively against federated cloud native data.

Autonomous Case and Response Axis

The vertical axis measures whether the product can autonomously assemble a case, reach a defensible conclusion, and move into useful governed response. It includes no-SIEM case creation, evidence continuity, dynamic reasoning, response usefulness, authority controls, auditability, and outcome verification.

A high score means the platform can do more than query cloud data. It can turn that data into

an actionable case and support or execute remediation. A lower score generally means the product is strong at data access or investigation but has narrower response authority, less mature governance, fragmented case continuity, or weaker post-action proof.

Cloud Native Leaders Category

Cloud native Leaders can operate without a traditional SIEM while still providing meaningful case reasoning and response. They are not simply cloud-hosted products. Their architectures support direct or federated access to cloud native data, engineering workflows, and dynamic evidence assembly, while preserving enough authority and governance to move toward closure.

This category is the primary shortlist for teams building a modern no-SIEM operating model. 7AI offers the strongest combined case and autonomy proposition. Exaforce is strongest when data architecture and raw cloud context dominate. Torq and BlinkOps are strongest when API-driven action and workflow flexibility dominate. Mate is differentiated by context quality and control design but requires more production diligence, while Swimlane is differentiated by its ability to apply mature response governance and case continuity to federated, and no-SIEM environments. Although its administrative model remains heavier than the leanest AI-native products.

Autonomy Specialists Category

Autonomy Specialists can provide strong reasoning, case management, governed action, or platform-native response, but are less naturally aligned to a lightweight data-lake-centric operating model. The limitation is usually architectural weight rather than a lack of security capability.

Some vendors in this category can technically replace a SIEM or provide their own data layer, but doing so may require broader platform adoption, higher cost, or more administration than a small engineering-led team wants. Others are strong AI-native products whose integration or production proof is not yet sufficient to clear the cloud-execution threshold.

These vendors are viable when the team is willing to adopt a broader platform or already owns the relevant ecosystem. Palo Alto Networks and CrowdStrike can be excellent choices when consolidation is intentional. Prophet, Simbian, and AIStrike fit teams willing to accept younger production evidence for a more modern reasoning model.

Execution Specialists Category

Radiant Security is the only vendor in this category. It performs well on fast cloud-oriented deployment, native log handling, practical triage, and useful response, but its autonomy-governance model remains less mature than the cloud native Leaders.

Radiant's placement shows the difference between being operationally convenient in a cloud environment and serving as a fully governed autonomous response control plane. The product can create real analyst leverage and support one-click or selected zero-click actions, but deeper risk-tier policy, multi-stage approval logic, and independent verification of final risk reduction remain less complete.

Radiant can be a strong fit for a small team that prioritizes speed, broad triage, and practical containment over immediate delegation of high-impact actions. The buyer should define which actions can run unattended, how policies change by asset criticality and confidence, how failures are handled, and how the platform proves that the intended risk was actually reduced.

Qualified Contenders

These vendors provide useful capabilities but are less aligned to the complete no-SIEM operating model. Their strengths tend to be concentrated in endpoint response, investigation, forensic analysis, or governed SOAR workflows rather than in the combination of cloud native data access, lightweight deployment, autonomous case construction, and broad response.

Vendor Analysis Findings

The broader market findings describe the common direction of AI SOC development. The comparative vendor analysis reveals how differently products reach that destination and why those differences matter across operating environments. Architecture, integration model, governance heritage, deployment burden, and control-plane ownership consistently influenced vendor placement, often making a product highly suitable for one SOC model and less natural for another.

Architecture has a significant influence on vendor placement and adoption. Palo Alto Networks, CrowdStrike, and SentinelOne benefit in regulated environments where buyers want a unified platform to provide telemetry, decisioning, policy, and enforcement. That same platform dependence creates more friction in hybrid environments where buyers prioritize an independent layer across a fragmented stack. SOAR heritage can introduce administrative and process weight, but it does not inherently limit hybrid or cloud native applicability. Swimlane demonstrates that a SOAR-derived platform can support those environments when it provides direct telemetry access, federated evidence assembly, independent case construction, and production AI reasoning. D3 Security retains more of the traditional SOAR tradeoff because comparable portability and production evidence for its newer AI layer remain less developed.

Modern reasoning and investigation capabilities also need to be separated from enterprise response maturity. Vendors such as 7AI, Torq, BlinkOps, Prophet, Simbian, and Mate demonstrate that AI-native systems can assemble strong cases and support graduated autonomy, but enterprise readiness still depends on production duration, integration breadth, rollback, source-state verification, and deployment evidence. Dropzone, Qevlar, Intezer, and Radiant similarly create substantial analyst value through investigation and evidence assembly, while narrower authority and less mature governance limit their ability to serve as the primary remediation layer. Buyers therefore need to determine whether they are purchasing an autonomous investigator, a governed response operator, or a platform capable of performing both roles.

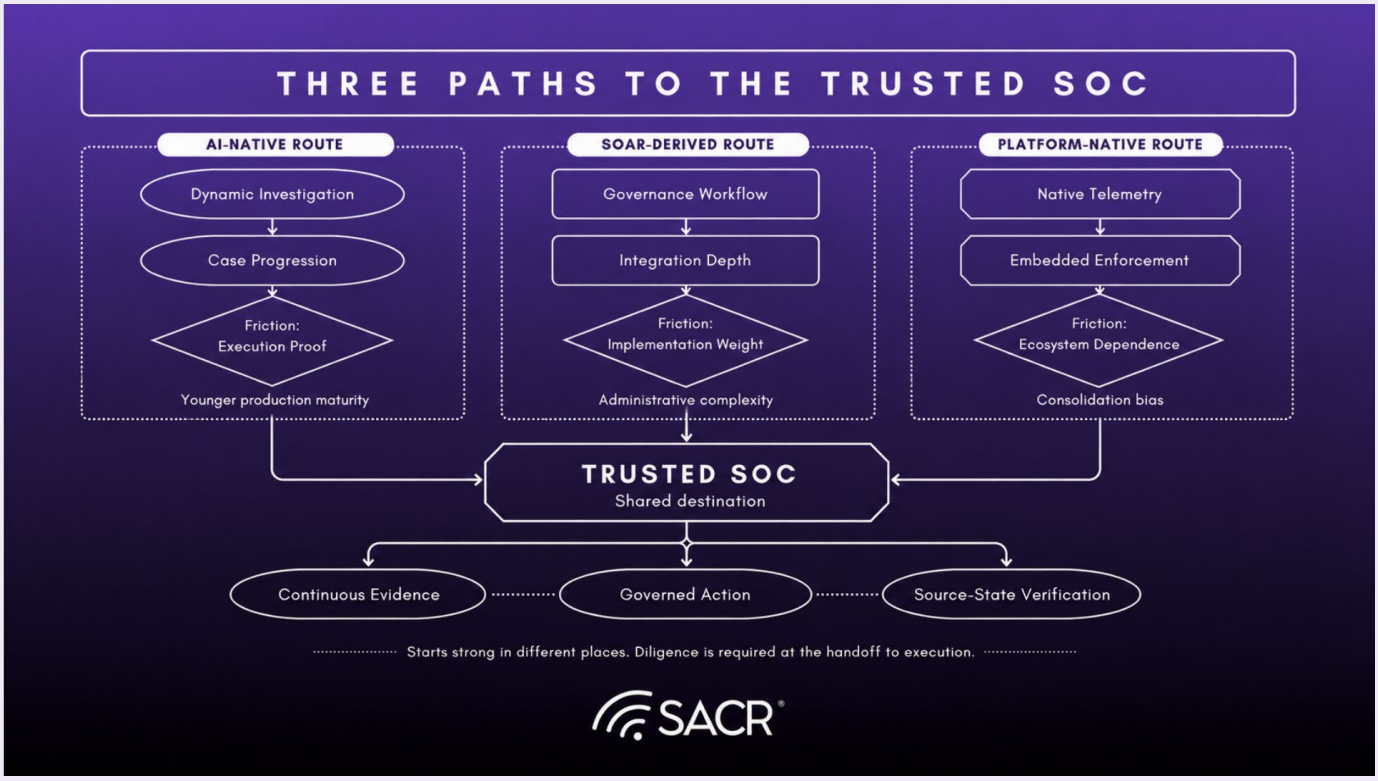

Outcome verification remains the clearest dividing line across the market. Many vendors can demonstrate that an action was initiated or that an API request was accepted. Far fewer can re-query the source system, prove that the intended state changed, confirm persistence, and connect the result to a measurable reduction in risk. This gap explains why products with strong investigation and action capabilities can remain outside the leading categories. Automated response ultimately requires evidence the execution action worked.

Viewed through the TSRO lens, the market is not divided simply between autonomous and non-autonomous products, or between trusted and untrusted vendors. Vendors differ in the type of authority they can support, the environments in which that authority is credible, and the evidence available to sustain it. AI-native vendors often lead in dynamic investigation and case progression, SOAR-derived vendors often lead in orchestration and governance, and platform vendors often lead where native enforcement and consolidation are acceptable. The strategic question is which architecture can govern the buyer's required authority portfolio without breaking evidence continuity, accountability, or operational fit. The Trusted SOC is built through that portfolio, not purchased as a designated platform.

Vendor Profiles

The following profiles explain the product architecture, operating strengths, tradeoffs, and market implications behind each vendor's placement. The evaluation applies only to the AI SOC capabilities and product scope examined in this report. It should not be interpreted as a judgment of the vendor's overall corporate maturity, financial position, installed base, or broader security portfolio.

Each profile assesses the role a product can play within a Trusted SOC authority portfolio: the actions it can support, the controls surrounding those actions, and the operating environments in which that authority is most credible.

Vendor placement may change across environments because the rankings reward different forms of operational fit. A platform may benefit from native telemetry and enforcement in a consolidated enterprise while imposing greater friction in a heterogeneous or engineering-led environment. A younger AI-native product may provide stronger investigation and analyst experience while carrying less production evidence, narrower execution breadth, or less mature governance.

7AI - Innovator

Vendor Overview

7AI is an AI SOC platform and managed service provider focused on automating investigation, decision support, and response workflows across customer environments. The same technology stack supports multiple consumption models: customers can operate the platform with their own SOC team, consume 7AI as a managed service through PLAID ELITE, or use the platform with MSSP and MDR partners.

The company's positioning is strongest around the final mile of SOC work. 7AI's argument is that the hardest operational question often comes after the case is assembled: what happens next, who approves it, and how much authority the system receives. The platform moves from evidence to action through customer-defined authority levels, including recommendation, approval-required execution, supervised execution, and policy-gated automation for lower-blast-radius actions. Action plans are generated from assembled case evidence and organizational context, while confidence and policy controls determine whether an action is recommended, queued for approval, or executed. Failed or higher-impact actions can follow defined retry, escalation, and action-specific recovery paths.

Product and Architecture

7AI's architecture spans detection optimization, agentic investigation, threat hunting, organizational context, workflow automation, and optional managed service delivery. Detection Optimization analyzes rules across connected sources, identifies noisy or incomplete coverage, recommends tuning, and maps detection coverage to MITRE ATT&CK. The platform connects to security tools, consumes alerts and evidence, runs investigations through task-specific agents, and exposes agent reasoning. The product shows the mission assigned to each agent, the inputs used, tools called, requests made, responses received, and conclusions reached.

Enterprise Insights is the most important context feature. It allows customer-specific knowledge, including user roles, accepted tools, disallowed software, business processes, and known behavior patterns, to influence investigation and response decisions. That matters because the same signal can have different meaning across organizations or business units.

Threat Hunting supports analyst-directed, natural-language hunts and threat-intelligence-driven hunts across connected customer telemetry. Analysts can start with a hypothesis, behavior, or ATT&CK technique, while incoming intelligence can initiate hunts as new indicators become available. Separately, the Skills framework lets customers create, test, and reuse investigation, hunting, and automation routines tailored to their environment.

The integration story is broad, with the platform designed to connect across security and IT systems. Integration depth determines response value because actions depend on API quality, permissions, telemetry freshness, identity mapping, and available write scope. After an action executes, 7AI can re-query the source system to confirm the resulting state and record submission, execution, and verified completion separately in the case. Federated search is an important direction because it lets the platform search connected tools without requiring all data to be copied into a central SIEM.

Market Context and Positioning

7AI receives an Innovator placement because its architecture connects unified case construction, transparent reasoning, graduated authority, and conclusion-driven response, while named production outcomes support strong Production Delivery despite a younger response catalog than the mature orchestration leaders. 7AI fits the AI SOC market as a hybrid platform and service model for organizations that want investigation automation tied to action. Its competitors include AI SOC point solutions, MDR modernization providers, SOAR and automation platforms adding AI, and

security platform vendors embedding autonomous response into their own control planes.

7AI's market position combines cross-stack neutrality, organizational context, transparent agent reasoning, and flexible delivery across software, managed service, and partner-led models. Its architecture begins with agentic investigation and extends into governed response, while detection optimization and threat hunting create feedback loops that can improve future case quality. Hunting findings can inform detection coverage, and investigation or response outcomes can support subsequent rule tuning. The value of this model depends on integration depth, customer-specific context, and how consistently those feedback loops operate in production.

Narrative Implication

7AI sharpens an important point for the AI SOC report: the final mile of SOC automation is organizational judgment encoded into a repeatable workflow. The platform is designed around the idea that the SOC needs a system that can learn customer context, gather evidence, explain reasoning, suggest or execute next steps, verify the resulting state, and let humans decide where automation should stop.

The product's publication story is delivery flexibility. 7AI can be consumed as software, service, or partner-enabled capability, making it relevant to buyers that want agentic SOC outcomes but differ in how much operational responsibility they want to retain internally.

Conclusion

Building the Trusted SOC established that AI SOC maturity is determined at the level of individual response actions through the Trusted Security Response Operations model. This market analysis applies that operating model to the products buyers can evaluate today. It shows where vendors can support the Trusted SOC, how reliably they can deliver capability in production, and which operating environments strengthen their fit.

The two ranking dimensions reflect that connection. Architectural Alignment measures how completely a product preserves the trust model from evidence through verified outcome. Production Delivery measures whether the product has the integrations, governance, and controls to sustain the operation model in practice. Together they indicate how well each vendor can support an earned response authority.

The rankings also show why a Trusted SOC will take different forms across organizations. AI-native vendors can strengthen a heterogeneous environment, while SOAR-derived vendors can provide a mature deterministic foundation to a regulated environment. Broader security platforms can connect native telemetry and enforcement with a lower integration friction. Critical upstream enablers improve evidence and detection foundations required for trustworthy decisions, while managed providers can operate portions of this lifecycle for customers that need outcomes instead of self-management. Each path supplies different parts of the Trusted SOC and places different responsibilities on the customer.

The authority portfolio should guide vendor selection. Buyers should define which actions software may recommend, prepare, execute, or verify, then use this analysis and their own diligence to select the architecture and delivery model best suited to those responsibilities in their environment. A Trusted SOC develops as the organization expands that authority only when evidence, controlled execution, and verified outcomes support it.

Software Analyst®
Cyber Research

business

personal

softwareanalyst.io

SACR®

Trusted research. Sharp insights. Real conversation.

CISO	VENDOR
SECURITY TEAMS	INVESTORS