



IDC MarketScape

IDC MarketScape: Worldwide Standalone AI SOC Platforms 2026 Vendor Assessment


Christopher Kissel Monika Soltysik

THIS IDC EXCERPT FEATURES NIKSUN AS A LEADER

IDC MARKETSCAPE FIGURE

FIGURE 1

IDC MarketScape worldwide standalone AI SOC platforms vendor assessment



Source: IDC, 2026

Please see the *Appendix* for detailed methodology, market definition, and scoring criteria.

ABOUT THIS EXCERPT

The content for this excerpt was taken directly from IDC MarketScape: Worldwide Standalone AI SOC Platforms 2026 Vendor Assessment (Doc # US54127026).

IDC OPINION

Security operations centers (SOCs) are being asked to do more with a workforce that is not growing at the same rate as alert volume. IDC's own survey work finds that 63% of organizations experience alert overload frequently or worse, and that false positives and process limitations, not staffing gaps, are the most-cited reasons analysts cannot investigate every alert. The result is a SOC where mean-time-to-detect and mean-time-to-respond stay stubbornly high even as tool spend increases, because the bottleneck has shifted from data collection to correlation, triage, and decision-making.

Standalone AI SOC platforms are built to close that gap. By applying agentic reasoning to triage, investigation, and response, these tools compress work that once required a human analyst reading through multiple consoles into a single, largely automated workflow, while preserving rollback and human checkpoints where buyers still demand them. This is not full autonomy; it is a shift in where analyst judgment gets applied, from repetitive triage toward oversight and exception handling.

The remainder of this study examines how the vendors in this assessment are executing on that promise, and where vendor claims and buyer expectations still diverge, starting with the inclusion criteria and market dynamics that define this category.

IDC MARKETSCAPE VENDOR INCLUSION CRITERIA

The AI-driven security operations center can be roughly described as a security analytics layer delivered as an appliance or software as a service (SaaS) that drives alerts, investigation, and remediation to SOC analysts. Responses are often autonomous depending on how the client asks. A more formal definition can be found in the *Market definition* section.

The technology and minimal financial requirements to qualify for this study include:

- A company must have a combination of \$20 million in revenue or in seed capital as of March 1, 2026 (this was revised from December 31, 2025, as the study evolved).
- The solution must also be generally available (GA) for deployment as of April 1, 2026.

- For formal review and scoring, capabilities that were GA as of August 7, 2026 (the end of the Black Hat conference) were considered.
- These solutions are deployed on the data/analytics layer of SOC visibility tools.
- Agents are very likely to be a part of the solution, but it is not a requirement. Agents are something of a double-edge. While they may add visibility and help create a preferable workflow, they must be secured, and the solution must be resilient.
- This study is about the analytics layer of the SOC. In this case, the exclusion criteria matter for this study.
- In this study, IDC does not include companies that leverage endpoint to become extended detection and response (XDR) solutions.
- There are managed detection and response (MDR) and managed security service providers (MSSP) that have developed in-house, standalone AI SOC capabilities, and in some cases may even be licensing the technology. These are covered with IDC's Managed Security Service and not considered here.
- Additionally, in this study, IDC does not include companies that are next-gen SIEM or SIEM (i.e., vendors that have their own data lake as a source). Actually, we note that a few vendors developed this capability during the developmental window of this study and were excluded.

ADVICE FOR TECHNOLOGY BUYERS

Enterprise buyers are moving faster than the market can standardize how standalone AI SOC platforms should be evaluated. IDC's global survey work and vendor evidence gathered for this study point to a set of practical considerations that cut across all platforms in this assessment, regardless of architecture or pricing model. Some of these observations reflect what buyers themselves report wanting; others reflect gaps and inconsistencies IDC observed while comparing vendor claims against reference accounts and demonstrated capability. Together, they offer a starting checklist for structuring a technology evaluation, a request for proposal, or a proof of concept, rather than an exhaustive buying guide.

- **Never lose the fundamentals.** No matter how sophisticated AI technologies and platforms become, a failure to execute on the most elemental cybersecurity fundamentals will cause the greatest undoing. Machines, identities, and agents that are not properly credentialled, especially in the age of Mythos, will very likely be exploited. Taking monthly inventories of hardware, software, and applications is a tedious but essential task. But if there are unaccounted applications, end users, personas, etc., investigating the reason for this anomaly will save headaches down the line. Consider preventive and proactive measures as a continuous loop that ties in with detection and response, because smart thinking ahead of time will reduce the blast radius and help refine and improve the quality of alerts.

- **You own the metrics; the metrics do not own you.** This is your SOC; thus, IDC encourages increasing use of automation, but the degree to which there is a human-in-the-loop is up to the technology buyer. The best description of what a CISO does is that he or she buys down risk: whether that means more headcount, more software, great internal guard rails, or living with a certain amount of risk, are all contingencies that must be accounted for, understanding that real-world cybersecurity is gated by fiscal constraints. One of the great things about agentic AI is that it will enable cybersecurity vendors to add greater customization to the tools and services they provide.
- **Design for reversibility, not just accuracy.** IDC's *Global AI-Driven SOC Survey* found that rollback capability is simultaneously the leading trust driver in autonomous decisions (55%) and the leading priority once an action has been taken (25%), ranking ahead of accuracy, explainability, and even incident prevention. That is a meaningful signal: buyers are not asking vendors to be perfect; they are asking to be able to undo a mistake quickly. Technology buyers should treat rollback and recovery mechanics as a first-order evaluation criterion, not a footnote to a containment features list. Ask vendors to demonstrate an actual rollback, including how state is captured before an action executes, how long recovery takes, and whether rollback works across third-party tools the platform does not directly own. Vendors that can only describe rollback in the abstract should be treated with more caution than the marketing language suggests.
- **Press vendors on closed-loop verification, because most have not built it.** 91% of survey respondents said verifying an autonomous action occurred and that risk actually decreased is the ultimate goal of automation. Only 26% currently rely on automated evidence collection to confirm that outcome; the rest depend on follow-up telemetry checks, manual spot-checking, or a simple API success response. That gap between expectation and practice is one of the widest in the survey. Buyers should ask vendors specifically how a claimed remediation is verified after the fact, whether that verification is automated or manual, and what happens when the verification step disagrees with the initial action. A vendor that can only confirm an action was attempted, rather than that it succeeded, has not solved the problem buyers say matters most.
- **Expect to extend the SIEM and existing tooling, not replace it.** Roughly 74% of surveyed organizations own a SIEM and expect a standalone AI SOC tool to work with it, and buyers are effectively split between full platform replacement (38%) and layering an autonomous capability on top of existing tools (40%). API-based integration into current workflows was the single most-preferred deployment model, ahead of migrating to a new investigation console. The practical implication for buyers is to weigh integration depth, API maturity, and coexistence with current SIEM and case-management investments as heavily as any autonomy claim. A platform

that requires wholesale replacement of an already functioning stack should be expected to justify that cost against a genuinely superior outcome, not simply a newer interface.

- **Treat vendor-reported outcome metrics as a starting point for questions, not a conclusion.** Nearly every vendor profiled in this study cites impressive self-reported figures: false-positive reduction rates above 90%, alert investigation times compressed from hours to minutes, and hundreds of millions of findings triaged, with only a small fraction requiring action. These figures are rarely independently audited, and methodology varies from vendor to vendor; a reduction percentage means little without the underlying baseline and sample size. Buyers should ask for the denominator behind any headline statistic, request the same metric calculated against their own environments during a proof of concept, and talk to reference customers about how the vendor's numbers were derived. IDC could not independently verify most of the outcome claims in this study, and buyers evaluating the same vendors should expect to perform that verification themselves.
- **Determine who actually owns detection engineering.** Vendors in this category take meaningfully different approaches to where detection logic originates. Some platforms generate and continuously tune their own detection content, incident scoring, and entity models as a core part of the product. Others are designed to ingest whatever detections already fire in a customer's SIEM and add investigation, correlation, and response on top of them, leaving detection engineering largely in the customer's hands. Neither model is inherently superior, but they carry different long-term staffing implications. A platform that depends on customer-authored detection content requires an organization to retain detection engineering skills in-house indefinitely; a platform that natively owns detection shifts that burden to the vendor, along with the risk of vendor lock-in. Buyers should explicitly ask which model a given vendor follows before assuming that autonomy extends to detection creation.
- **Interrogate guardrail architecture as a specific mechanism, not a marketing claim.** Nearly every vendor in this study describes its platform as having guardrails, but the actual mechanics vary widely. Some vendors rely on policy-based action limits and cost or execution caps; others use named, editable rules that flag specific users or assets for elevated approval; still others combine model-attack protections, such as prompt-injection defenses, with separate throttling controls for alert-volume spikes. Approval workflows also differ: some platforms require human sign-off for any high-impact action unless it has been explicitly pre-authorized. In contrast, others default to autonomous execution within a defined risk tier. Buyers should ask each finalist to walk through one specific, high-impact scenario end to end, showing exactly where a human would be inserted into the loop and why, rather than accepting the word guardrails as a sufficient answer on its own.

- **Test plug-and-play claims against realistic implementation effort.** Vendors routinely describe their platforms as easy to deploy, citing SaaS delivery, prebuilt integrations, or native AI assistance for configuration. Reference interviews conducted for this study told a more nuanced story: one long-tenured customer described its platform candidly as not a plug-and-play product, noting that meaningful value required building custom agents, integrations, and workflows tied to the organization's specific tools and use cases. That is not necessarily a criticism of the platform, since the same flexibility that requires configuration work is often what enables deep customization later. It does mean buyers should discount vendor-stated time-to-value figures and instead ask reference customers directly how long it took them to reach production use, and what internal resources were required.

VENDOR SUMMARY PROFILES

This section briefly explains IDC's key observations resulting in a vendor's position in the IDC MarketScape. While every vendor is evaluated against each criterion outlined in the Appendix, the description here summarizes each vendor's strengths, challenges, and opportunities.

NIKSUN

NIKSUN is positioned in the Leaders category in the 2026 IDC MarketScape for standalone AI SOC.

NIKSUN delivers its agentic AI SOC capability through Olympus, an AI engine spanning Maya AI, Spotlight AI, and Seeker AI, built on the NIKSUN Knowledge Warehouse (NKW), a single cross-correlated data lake ingesting full packets, flows, logs, endpoint activity, and metrics. Deployment spans fully managed cloud SaaS through NIKSUN Cloud Services, on-premises hardware, virtual software, and an air-gapped option, priced across four dimensions, data type, data rate, retention, and enabled use cases, rather than by ingest volume, with no separate charge for data normalization, integrations, onboarding, or enhancement requests.

The platform reached general availability March 1, 2026; NIKSUN has been self-funded since its inception.

Strengths

- **A single correlated data lake behind every other claim.** NIKSUN's Knowledge Warehouse (NKW) ingests full packets, flows, logs, endpoint activity, and metrics into one cross-correlated store rather than assembling detection from separate SIEM, NDR, EDR, and forensics products. One global manufacturer felt the cost of the opposite approach directly: during a Cobalt Strike intrusion, analysts spent the better

part of the day manually stitching together evidence across disjointed tools before establishing the scope of the compromise, by which point the attacker had already moved laterally and exfiltrated data. After migrating to Olympus, the same organization's analysts could ask NIKSUN's AI agent, Maya, what happened directly against the unified lake rather than deciding which tool to check next, then enable automated remediation through Maya's agentic action capability to respond to and resolve the incident.

- **Network-level detection that holds up when other telemetry goes quiet.**
Because Olympus captures and correlates network evidence independently of what endpoint agents report, it continues functioning as what NIKSUN calls an independent witness even when host-based tooling is absent, disabled, or bypassed. A regional healthcare provider running its entire IT and security function with two people had no dedicated security platform watching its patient portal beyond a firewall and an annual penetration test; once deployed, NIKSUN's Maya AI agent identified an active Log4Shell exploitation and data exfiltration in progress that the provider had no other means of detecting, then moved to contain it across host, network, and firewall before more records left the environment.
- **AI conclusions traceable to raw evidence, not asserted.** NIKSUN's governance model rests on grounding every AI conclusion in the underlying correlated data and making that evidence retrievable on demand, rather than asking analysts to take an agent's summary on faith. During the manufacturer's Cobalt Strike remediation, an engineer on the call asked directly how the team could trust what Maya was recommending rather than a confident-sounding guess; the analyst asked Maya for the raw evidence behind its conclusion, and the underlying packet and log data came back within seconds. That moment, specific to a live deployment rather than an abstract description of the architecture, is the kind of proof point that matters most for buyers evaluating hallucination risk in an agentic platform.
- **OT, IoT, and mobile/cellular visibility most endpoint-first platforms don't have.** NIKSUN's agentless, passive approach to operational technology and industrial control systems, combined with purpose-built mobile and cellular monitoring through NetMobility, reflects close to three decades of defense, critical-infrastructure, and carrier heritage predating most competitors' entry into these environments. That heritage is reinforced by a genuine, if partial, government-certification footprint, including DoDIN APL listing, Common Criteria certification, DISA STIG alignment, and NIST SP 800-171 and CMMC coverage, alongside a true air-gapped deployment option. For buyers with meaningful OT, IoT, or unmanaged-device footprints, this is a materially different starting point than the endpoint-centric architecture most platforms in this study are built around.
- **Network and application performance living in the same data lake as security telemetry.** Most platforms in this study treat network and application performance as

a separate discipline, one a security team has to go check in a dedicated observability tool, for example, to understand the impact of an attempted cyberattack. NIKSUN folds NPM and APM, through NetVCR, into the same NKW that powers detection, so performance degradation, application health, and security anomalies are correlated against each other rather than investigated in separate systems by separate teams. NIKSUN frames this explicitly as unifying security, performance, and compliance, SIEM, NDR, EDR, SOAR, threat intelligence, forensics, NPM, APM, and observability, into one platform. That performance context is worth weighing on its own merits: using performance signals in triage can reduce alert volume and improve alert fidelity. For organizations weighing whether to keep NetOps and SecOps on separate tooling, this convergence is a genuine point of differentiation.

Challenges

- **A scope question the vendor's own materials raise directly.** NIKSUN's submission repeatedly self-describes as a native SIEM/SOAR/XDR, a native EDR, and a native NDR, and names LogWave (SIEM/SOAR) and NetDetectorLive (NDR/XDR) among its core products, while the NIKSUN Knowledge Warehouse (NKW) itself is a dedicated, unified data lake. That is difficult to square with this study's own scope, which explicitly excludes next-gen SIEM, defined as including a dedicated data lake, and EDR products. Whether NIKSUN's core Olympus platform sits inside or outside those exclusions is a genuine open question rather than a wording nuance. (Worth noting, NIKSUN's stated view is that grounding agentic action in a unified data lake is what enables it to act without hallucinating, in contrast to platforms that reason over fragmented or unnormalized data.)
- **A differently positioned certification strategy behind an otherwise strong defense narrative.** None of the three commercial certifications this study asks about directly (i.e., MITRE ATT&CK Evaluations, FedRAMP, and SOC 2 or SOC 3) has been achieved; all three are listed as in progress. NIKSUN has instead historically focused on different, more stringent standards, including DoDIN APL listing, Common Criteria certification, and DISA STIG alignment, security accreditations required by defense, intelligence, and critical-infrastructure customers. The in-progress commercial certifications limit the platform's provable coverage for buyers that weigh those specific credentials, though their absence does not by itself indicate a security gap, since the underlying Olympus architecture already carries the certified components those standards require.
- **Headline outcome figures that recur without independent attribution.** A roughly 120-times-faster time-to-root-cause claim, cutting an estimated 600 minutes to about five, along with a 50–90% total-cost-reduction figure, are claimed by the vendor; NIKSUN cites customer confidentiality for not naming clients behind either figure. The vendor did supply the underlying methodology aside from noting that these are benefits from using a single, centralized SOC software stack.

Consider NIKSUN when

NIKSUN suits organizations that want network, endpoint, cloud, OT, IoT, and mobile/cellular telemetry correlated in one data lake rather than reconciled across point products, particularly those with meaningful operational-technology, industrial-control, or defense and government requirements, where NIKSUN's multi-decade heritage and certification footprint, DoDIN APL, Common Criteria, DISA STIG, is harder for endpoint-first competitors to match.

It is a fit for organizations that cannot staff a traditional multitool SOC. A regional healthcare provider running its entire IT and security function with two people used NIKSUN's managed cloud service to stand up detection the same day, catching an active breach that its firewall-and-annual-penetration-test posture had no way to see; a global manufacturer replacing a fragmented, multivendor stack after one costly breach reported a similar shift from reconciling tools manually to asking its AI agent directly what happened. Both point to real fit for buyers needing enterprise-grade coverage without adding headcount.

APPENDIX

Reading an IDC MarketScape graph

For the purposes of this analysis, IDC divided potential key measures for success into two primary categories: capabilities and strategies.

Positioning on the y-axis reflects the vendor's current capabilities and menu of services and how well aligned the vendor is to customer needs. The capabilities category focuses on the capabilities of the company and product today, here and now. Under this category, IDC analysts will look at how well a vendor is building and delivering capabilities that enable it to execute its chosen strategy in the market.

Positioning on the x-axis or strategies axis indicates how well the vendor's future strategy aligns with what customers will require in three to five years. The strategies category focuses on high-level decisions and underlying assumptions about offerings, customer segments, and business and go-to-market plans for the next three to five years.

The size of the individual vendor markers in the IDC MarketScape represents the market share of each individual vendor within the specific market segment being assessed. The rough sizes of the bubbles come from the end of year 2025 actual revenues as IDC could best verify and these pertain to the platform reviewed; not to overall revenues for the company.

IDC MarketScape methodology

IDC MarketScape criteria selection, weightings, and vendor scores represent well-researched IDC judgment about the market and specific vendors. IDC analysts tailor the

range of standard characteristics by which vendors are measured through structured discussions, surveys, and interviews with market leaders, participants, and end users. Market weightings are based on user interviews, buyer surveys, and the input of IDC experts in each market. IDC analysts base individual vendor scores, and ultimately, vendor positions on the IDC MarketScape, on detailed surveys and interviews with the vendors, publicly available information, and end-user experiences in an effort to provide an accurate and consistent assessment of each vendor's characteristics, behavior and capability.

When vendor sections were written, they were sent back to vendors for factual review. IDC always retains the right to accept or decline the suggested edits. As this document was in the publication cycle, IDC talked to the vendors about their score (Leader, Major Player, or Contender).

Market definition

In the study, we use the word “solution”. IDC sees XDR, SOC, SOAR, and Hyper SOAR as a merging form factor. Almost all of these solutions use of agents, but that is not a necessity for inclusion in the study.

The following are the required solution features:

- The platform must have continuous visibility of endpoints. This isn't the type of control synonymous with endpoint detection and response (EDR) but the platform must see if a device is managed or unmanaged. Additionally, the solution must be endpoint agnostic.
- The solution must produce alerts. No matter how it gathers telemetry, it has to generate alerts, prioritize alerts, and diminish false positives. The ability to create agents is preferred, but no matter the agent technology, the agents must be safely deployed.
- Security orchestration automation and response (SOAR) is necessarily included on the platform. The SOAR generates a workflow and advises or initiates ephemeral and permanent responses to incidents.
- The platform must provide contextual awareness including detection, contain, and recovery strategies. The platform must be able to describe what prompted an alert, what the blast radius is, and provide strategies on how to reduce the attack surface.
- User entity behavioral analytics (UEBA) is an essential part of the solution. UEBA is designed to be self-learning and consistently wizing. What UEBA does is establish statistical baselines for every entity within the network as well as setting the relationships between devices within the network. UEBA is one way to understand how devices are connected and provides the ability to reset devices back to the network if devices need reimage.

- Ideally, the solution should be able to integrate NDR, EDR, threat intelligence, web/email security, and SIEM. IDC looks favorably for integrations with IAM, firewall, data security, and other IT/Security platforms.
- There is not necessarily “one solution fits all.” Vendors can differentiate by offering DLP, special treatment of threat intelligence, and the inclusion of other sources of telemetry. Also, notable value-adds can include help with asset inventory, compliance, IT operations, and application monitoring.
- Extensibility to hybrid architecture is preferred. A standalone AI SOC platform should include visibility for on-premises networks as well as for the leading public cloud environments such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP), and virtual private clouds (VPCs).
- The solution should include specific insights about identities including user behavioral analytics (UBA), tracking identities through multiple surfaces, and protection of Active Directory.
- The solution should be able to continuously assess risk.
- No matter how good the solution is, it must protect its agents and its core analytics from attack. Resiliency also matters.

As the study evolved

Below will be a list of Strategies and Capabilities that were formally considered for scoring the study, and the general guidelines are a solid grounding for the study, but the reader should be aware of observations of notes:

- Agentic AI is the next logical step in IT and security platforms, but a real appreciation should be given to GenAI capabilities. Without exception, every platform measured in this study offered explicit NLP descriptions of how an alert was generated (noting that conclusions made in a demo environment are tough to super-impose into what would be a client’s specific environment).
- The ability to port agentic capabilities in different form factors. The vast majority of the time, these standalone AI SOC tools will go onto a network as SaaS. But often, different clients may want an air-gapped solution, a solution deployed as local enterprise software or have a version that is dedicated to one or a small number of machines. This type of versatility was a consideration in the assessment.
- IDC believes in the human-in-the-loop concept, but certain responses need to be immediate (isolating a suspected endpoint or adding a known bad domain/IP address to a badlist on a firewall at the time of discovery).
- It becomes clearer that the detection and response cycle begins “left of boom.” Solution providers that can provide automated red teaming, exposure management, vulnerability scanning and/or attack surface management, microsegmentation, and

business risk management are reducing the blast surface proactively and consequently getting better alert fidelity.

- At this point, we have ambivalent feelings about detection rules. The old caveat about existing detection rules is that they handle about 80% of practical alerts efficiently. This has two benefits. First, detection rules help create a consistent SOC process; and second, rather than triggering tokens to create a playbook, the action is simply taken. Conversely, there are two problems with detection rules. The first is that they require constant tinkering and updating. The second is that attackers are adding depth to their tactics. Detection rules are all but useless against a zero-day attack.
- Perhaps one of the greatest aspects of LLMs is that they can build code much more quickly than human developers. Ostensibly, that would dramatically lessen the need for formal connectors, integrations, and software development kits (SDKs). Like the comment above about detection rules, a true bidirectional connector is preferred over a simple connector, but the advantage to security vendors that have a shelf of connectors and formal integrations versus custom-created connectors is unclear (IDC still leans toward the integration model).
- The AI harness matters. Many of these platforms are designed to generate customized agents. However, the agents should be debugged before deployment. The harness could include but is not limited to checks that sensitive data does not leak, conclusions made outside of an underlying LLM model, multiple redundancies safeguarding against hallucinations, and assurance checks to make sure agents do not access parts of the network that they have no permissions to go to.
- Pricing is hard to quantify. In this evaluation, IDC provided a sample business for midsize companies and enterprises. Many of the vendors were transparent about their pricing paradigms, though it was still hard to provide even sample pricing without knowing how many routers, switches, and concentrators or what the practical heterogeneity of what client network is. When a business asks for a proof-of-concept, it is only then that the cybersecurity solutions provider can understand what happens at deployment.
- The biggest factor toward rating these companies and platforms was (naturally) how much the standalone AI SOC could make a client's SOC more effective. These factors would include manual hours saved in triage or response, risk diminishment, and added visibility. Many of the vendors though had capabilities that could eliminate the need for NPM, APM, configuration assessment tools and even SIEM replacement. Self-evidently, if a standalone AI SOC solution can either add functionality to the SOC or displace a tool, that should be considered when a client decides to commit to a tool. In this study, when formal metrics were relatively close, the ability to help with tool consolidation was a tiebreaker.

Related research

- *2026 IDC Survey: Global AI-Driven SOC Survey, 2026 — From Automation to Conditional Autonomy* (IDC #US54783026, August 2026)
- *2026 Network Detection and Response Market SWOT Analysis* (IDC #US54594024, June 2026)
- *IDC Survey: Agentic AI Use Case Adoption Survey Within the Cybersecurity Function, 2026* (IDC #US54531226, May 2026)
- *Worldwide Network Detection and Response Forecast, 2026–2030* (IDC #US52790926, April 2026)
- *Novel Threat Detection and Response Report: Presenting 20 Novel Use Cases* (IDC #US53936425, December 2025)
- *IDC MarketScape: Worldwide Extended Detection and Response Software 2025 Vendor Assessment* (IDC #US52997325, September 2025)
- *Worldwide SOAR and Firewall Automation Forecast, 2025–2029* (IDC #US53802825, September 2025)

Synopsis

This IDC MarketScape study assesses vendors offering standalone AI security operations center (SOC) platforms that met the inclusion criteria of this study, scored against strategy and capability criteria built from vendor RFI responses, product demonstrations, and customer reference interviews. The category exists because SOC alert volume keeps outpacing headcount: 63% of organizations report frequent alert overload, and false positives and process limitations, not staffing gaps, are the most-cited reason analysts cannot investigate every alert. Standalone AI SOC platforms apply agentic reasoning to triage, investigation, and response, compressing multiconsole work into a largely automated workflow. Adoption spans enterprise teams, MDR providers, MSSPs, and systems integrators, most expecting a new platform to work alongside an existing SIEM rather than replace it.

The urgency behind this category is no longer theoretical. In 2026, Anthropic disclosed that its unreleased Claude Mythos model autonomously discovered thousands of high-severity vulnerabilities and chained them into working exploits, Wiz's Project Atlas outscored Mythos on a vulnerability benchmark, and Hugging Face was breached not by a human, but by an autonomous AI agent that escaped its sandbox. Each event compresses the gap between disclosure and exploitation, signaling offensive AI is scaling as fast as defensive AI. This study weighs vendors across two axes, strategies and capabilities, pairing the IDC

MarketScape figure with vendor profiles so buyers can match corroborated maturity against their gaps rather than marketing claims.

“This is an exciting time for both cybersecurity vendors and the clients they serve. Agentic technologies are compressing detection and response from hours to minutes, and speed is no longer optional. The adversary has already added AI to its own tactics, using autonomous systems to discover vulnerabilities, chain exploits, and in at least one case, breach a target directly. Vendors that can show their agents reason transparently, act within defined guardrails, and improve with each investigation will define this category. The rest will be measured against them,” said Chris Kissel, vice president, IDC Security & Trust.

ABOUT IDC

International Data Corporation (IDC) is the premier global market intelligence, data, and events provider for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries. IDC's analysis and insight help IT professionals, business executives, and the investment community make fact-based technology decisions and achieve their key business objectives.

Global headquarters

One Beacon Street
Suite 33100
Boston, MA 02108
USA
508.872.8200
X: @IDC
blogs.idc.com
www.idc.com

Copyright and trademark notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, and web conference and conference event proceedings. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit idc.com/about/offices. Please contact IDC at customerservice@idc.com for information on additional copies, web rights, or applying the price of this document toward the purchase of an IDC service.

Copyright 2026 IDC. Reproduction is forbidden unless authorized. All rights reserved.