The AI SOC Buyer’s Guide 2026

How to Choose the Right AI Vendor for Your SOC

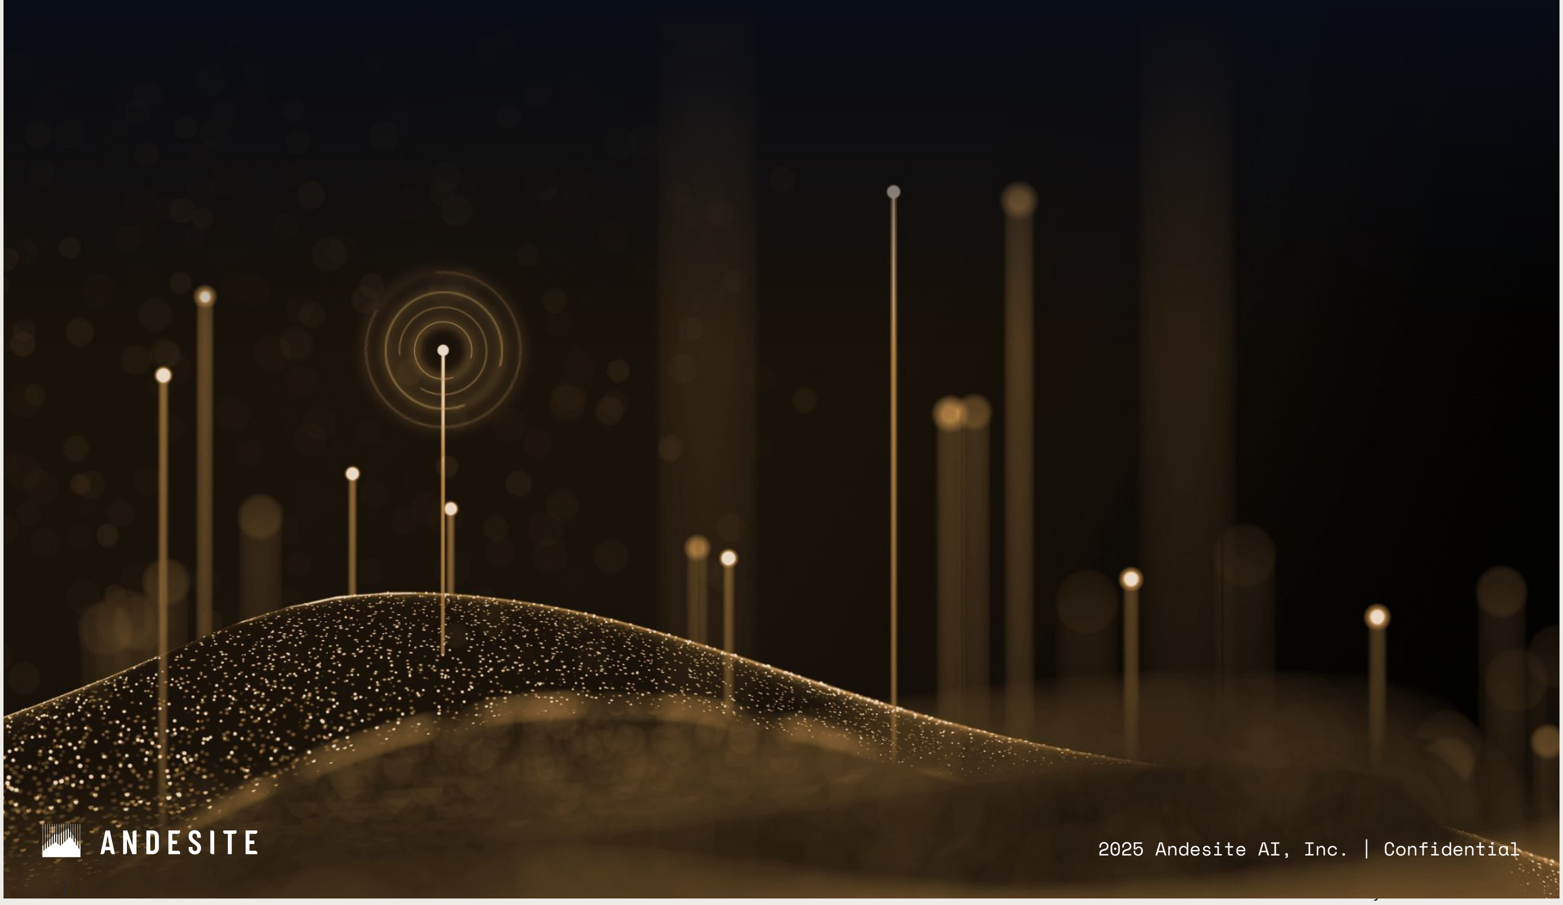

ANDESITE

Contents

Executive Summary

The emerging AI SOC field is fast-evolving and crowded, with a myriad of new offerings that boast similar features regularly going to market. In an increasingly complex and sophisticated threat landscape, it's important to arm those tasked with keeping our world safe with a better system to do their job.

The goal should be to empower them to focus on threat prevention and risk reduction, rather than being constantly kept on their back foot. It's time to flip from a reactive stance to a proactive strategy.

In finding the right AI SOC solution for your organization, you face the dual challenge of addressing today's SOC pain points while building the cyber defense of tomorrow. Emerging solutions can seem all too similar and it's difficult to identify the features and products that will give your cybersecurity team a meaningful advantage. This complicates the process for decision makers, who have to sit through multiple rounds of sound-alike pitches about products that may not be easy to differentiate. In addition, each SOC and organization's needs, use cases, team configurations, tools, and sources are different—and can't always be solved by out-of-the-box solutions.

We've created this guide to help CISOs and other decision makers ask the right questions and identify the best AI SOC solution for your unique needs, not just for today but far into the future.

Get started →

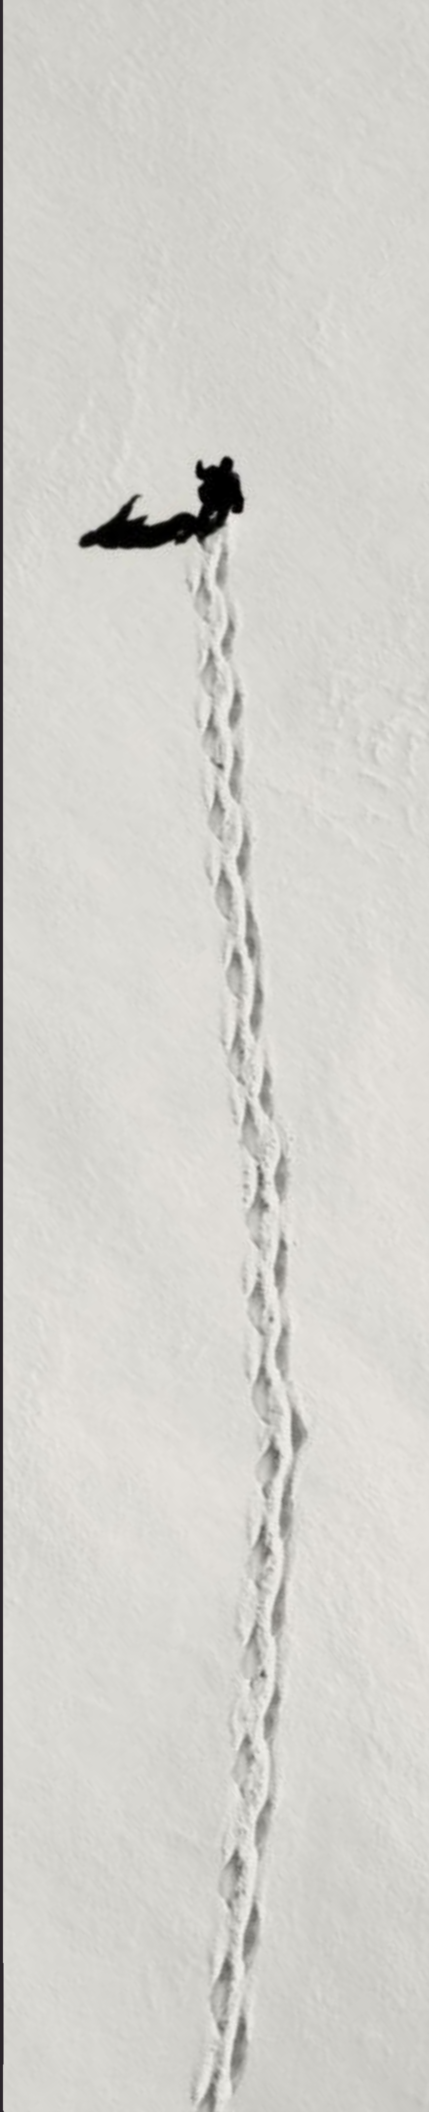

How to Choose the Right AI Vendor for Your SOC and Organization

The following questions will help you understand how the various AI SOC vendors you're considering measure up on key factors and functions, and whether or not they align with your specific SOC and organizational needs.

1 2 3 4 5 6 7

Does it put humans at the helm or is it fully autonomous?

Cybersecurity teams are overworked and overwhelmed by the proliferation of increasingly sophisticated threats and an ever-expanding, yet siloed, arsenal of cyber defense tools and data feeds.

A collaborative strategy keeps your team's hard-earned expertise anchored front and center.

AI-enabled SOC solutions are designed to support and alleviate analyst burnout by accelerating alert management, investigation, and enrichment. The goal of these solutions is to make both their work and their lives better by allowing AI to take care of the menial, time-consuming tasks, like managing alerts and generating report summaries.

AI SOC solutions generally fall under two categories—fully autonomous, where the AI handles actions from start to finish, and those that keep humans in the loop for a collaborative approach. Many companies are ready to go all in with AI, but it's important to understand the limitations that come with an autonomous system. While both approaches save time, a collaborative strategy keeps your team's hard-earned expertise anchored front and center.

The advantages of autonomous AI

With autonomous solutions the machine does everything, working completely on its own. The focus is on menial tasks—AI can prioritize threats, clean data, and perform repetitive, volume-based tasks quickly. It easily handles the tedious work that makes humans unhappy, drastically simplifying high-volume alert management.

This works for the types of tasks AI excels at, such as processing large amounts of data and doing repetitive work at much higher speed and accuracy. Automation brings support and innovation into the SOC, freeing cybersecurity teams up from volume-heavy tasks to make their work more effective and fulfilling.

When it comes to complexity, autonomous AI can be risky.

Threat investigations can range broadly from routinely simple to incredibly complex, and when it comes to complexity, autonomous AI can be risky. The power of AI lies in its ability to operate at scale. While autonomous AI is helpful for managing and triaging high-volume alerts, human expertise is essential when things get more complicated.

AI is limited by the quality and scope of the data it's trained on. When faced with situations outside of that training data, AI lacks the ability to truly "understand" and falters as a result. Just as it's beneficial to focus AI on what it excels at, you also want to keep humans focused on what they do best.

The benefits of keeping humans at the helm

The Human-AI SOC works alongside your team by automating mundane tasks while keeping people in charge of decision-making.

Rather than replacing human talent and intelligence, the Human-AI SOC works alongside your team by automating mundane tasks while keeping people in charge of decision-making.

While autonomous AI can identify correlations across data sources, it doesn't have the experience, intuition, or flexible reasoning needed to make sound judgment calls. That is usually a job for senior team members whose reflexes combine conceptual understanding, deep expertise, and even intuition. Without human oversight, real threats can slip through the cracks.

AI can help by providing actionable insights that cybersecurity teams can use to make critical decisions, enabling prevention rather than reaction. Still, people have a depth of understanding that AI can't mimic. We draw on things like organizational and emotional context, past experiences, and lateral thinking. Effective decisions are based on a mix of quantitative data and qualitative judgment that's uniquely human, and shouldn't be trusted to AI.

At every point, humans should provide input and review what the AI is doing. This matters, because humans can see relationships and deeper connections that may not mean anything to the AI. If you're dealing with complexity that requires deep analysis, you need a solution that keeps humans not just in the mix, but at the helm.

Does the product adapt to your security use cases, ecosystem, and workflows?

A flexible architecture that adapts to your security ecosystem and workflows is likely the winning option, but it can mean greater complexity and additional set-up time.

Ideally, an AI SOC solution should:

- Be tailored to your risk profile
- Accelerate time to detect, investigate, and respond
- Help reduce your overall risk exposure

To ensure the solution meets your precise requirements, you need built-to-fit integrations, bespoke connectors, and an architecture that adapts to your workflows with the ability to scale and change as your needs evolve. Your solution should integrate with the platforms and tools your team relies on, like SIEM, SOAR, CAASM, CTI, email and identity security solutions, and more.

The goal should be to connect all of your data sources seamlessly and empower your team to access their security tools and insights in the same space.

The right solution can adapt to exactly what you need and be able to address complex use cases. This calls for a customized set-up that reflects the specifics of your organization. You want a solution that can serve your ecosystem as it is, while also offering the ability to expand with your needs. The goal is a sustainable AI-powered approach that can adapt to your business priorities, no matter how much things change.

If opting for an out-of-the-box solution, the advantage is readiness—a turnkey product will start working in hours. But it's important to remember that this quick deployment may come at the expense of customization, which can be problematic for complex ecosystems.

What you don't want is to tack just one more tool onto the SOC maze or insert additional steps into the process. Your ecosystem, team, and demands are constantly evolving. To keep up, you should invest in a more scalable solution that's flexible and adaptable. The goal should be to connect all of your data sources seamlessly and empower your team to access their security tools and insights in the same space.

Does it provide actionable, timely insights?

The AI SOC category's defining features are:

- Alert triage
- Investigation
- Enrichment

It's safe to assume that all vendors deserving consideration deliver on these three functions. While many offerings compete on speed, what truly differentiates products are factors like quality, reliability, and effectiveness.

The ability to deliver timely information and insights that SOC teams can use to prevent attacks and minimize damages further sets some vendors apart in this increasingly crowded field.

Your team needs to be able to quickly assess risk exposure and adjudicate threat level to accelerate time to detect, investigate, and respond. What they don't need is another new tool. Rather, the AI should enable a decision layer that operates above, beyond, across, around, and with existing data sources, tools, and platforms. A robust AI SOC solution should survey the incessant wave of alerts, analyze structured and unstructured data, automate investigations, and provide enrichment to deliver relevant, contextualized, and actionable insights—quickly.

This allows your experts to focus on the meaningful, analyze the evidence, manage the case, and deliver effective responses fast. Analysts should be able to use the tool to rapidly assess risk exposure and adjudicate threat level, significantly accelerating the response cycle. The greatest measure of success is how efficiently and effectively your AI SOC can separate the signal from the noise.

The greatest measure of success is how efficiently and effectively your AI SOC can separate the signal from the noise.

The right tool should connect data silos, tools, and platforms across your security ecosystem to deliver contextual information and offer verifiable evidence. In doing so, the AI provides a decision layer that your analysts can use to reliably identify the insights that matter to your specific risk profile. For example, if there's a threat to patient data at healthcare companies, but your organization is in retail, it's not likely to be an issue for you. When threat hunting, the AI SOC tool should identify the blast radius with contextual awareness.

Next, consider how the AI prioritizes and groups alerts. This is important for evaluating the fit, along with the quality and depth of the enrichment and investigation process, and the ability to combine multiple sources in the scope of a single investigation. What autonomous solutions do is go alert by alert—but what's really needed is prioritization.

Look for a solution that groups alerts from different sources and prioritizes threats. Not only does this save time, it also makes it easier to understand the threat scope and differentiate potential impact. These capabilities empower your teams to focus on the most immediate and meaningful threats—with the context and correlations they need right at their fingertips.

When evaluating vendors, ask:

? Can you get a combined summary and report from multiple sources?

? Can you interrogate the conclusions and assess the risk in the specific context of your organization?

To enable a more robust response, you must be able to look across data sources and correlate related information. You also need a space where analysts can seamlessly investigate and collaborate across all of your security platforms. The goal is to improve analyst productivity across workflows, tools, data sources, and platforms.

How does it process threat intelligence and other unstructured data and enrichment sources?

Next, consider whether or not the product processes threat intelligence and other unstructured data. And if so, how—can you initiate an investigation from a threat intel report?

You want to avoid extracting or ingesting data, which slows things down and introduces risk. Instead, look for a solution that can analyze structured or unstructured data where it is.

Cybersecurity teams need to be able to access, connect, and use all available data sources and security tools together. You want to avoid extracting or ingesting data, which slows things down and introduces risk. Instead, look for a solution that can analyze structured or unstructured data where it is. A robust AI SOC gives analysts the power to access, visualize, and work with all the elements and data feeds in the existing security ecosystem. This reduces inefficiencies across data sources, tools, and platforms.

It's also important to choose a solution that can use contextual awareness to connect the dots across multiple alert and threat intel sources, grouping them into a consolidated view. Contextual awareness is crucial to correlating information and unlocking valuable insights. Ideally, you should be able to initiate an investigation by simply uploading a document or inputting a URL to prompt the AI to evaluate the risk relative to your organization.

Another important consideration is how long it takes to process a report, abstract insights, and incorporate them into an investigation. Time is of the essence, and automation should make analysis of threat intelligence documents available within minutes.

The solution should be able to review and gather insights quickly, without your team having to do independent review.

For threat intelligence reports provided by external parties, the solution should be able to review and gather insights quickly, without your team having to do independent review. This includes unstructured data such as PDFs and URLs. You should be able to upload the source, pull together insights, and correlate that information with your priority alerts. The ability to launch an investigation from the summary of a document and cross reference with your alerts means you'll have a better understanding of how the threat will impact your organization specifically.

Finally, ask if it can consolidate threat intelligence reports and alerts within the scope of a single investigation. You want an AI SOC that can survey alerts and automate investigations and enrichment so you know the information your analysts are working with is complete, relevant, contextualized, and actionable.

With the right AI SOC, you can empower your team to focus on the meaningful and deliver effective responses quickly. It should seamlessly integrate across your enterprise environment and proactively combine data sources to surface and prioritize key indicators for analyst attention. This will reduce the time between receiving threat intel reports to extracting key insights and initiating an investigation to minutes.

How does the product adapt to your specific organization context and risk profile?

Every organization's risk profile is different. A wide range of factors contribute, including your industry, regulatory requirements, and audiences served.

For threats that are in the system, if you don't have contextual awareness, you don't know if it's attacking a specific type of user.

The level of risk presented by a specific threat is not so much tied to potential damage it can do in theory, but to the actual danger it presents to your organization's realities. So, it may be more (or less) concerning for you operationally based on what type of customers you serve, the kind of information you're handling, and other business factors.

Industry-specific attacks demand contextual awareness. For threats that are in the system, if you don't have contextual awareness, you don't know if it's attacking a specific type of user. Without that insight, inefficient attempts at remediation may cause disruption. Contextual awareness lets you focus on the real threat, understanding exactly where within your organization remediation is needed.

The goal is for AI to proactively combine organizational data sources to unlock insights that are relevant to your specific risk profile and to provide contextual relevance and information to the SOC team assessing threats. This allows you to prioritize and deliver timely investigations and early detection for the threats that are most pressing for your business.

Is it audit-ready?

For highly regulated industries, compliance is a key consideration. Checking each AI SOC vendor's trust center is an important first step in deciding which one might align with your needs. Are they compliance-high? Do they meet your requirements? Are they aligned with the frameworks needed to protect your organization?

Depending on your industry and sector, you need a CISA Secure by Design product aligned with:

- FedRAMP High
- SOC 2 Type 2
- NIST CSF
- NIST 800-53 (High)
- ISO 27001
- ISO 27701
- ISO 42001
- HITRUST e1 + AI Security
- PCI DSS
- HIPAA
- AI RMF

This is crucial for financial services, healthcare, technology, critical infrastructure, and the public sector.

Regulatory requirements make auditability a necessity. While AI inputs and outputs are clear, it can be difficult to understand how the machine arrives at certain conclusions—a challenge known as the "black box" problem. Hallucinations are extremely risky for cybersecurity and must be guarded against.

For AI-driven investigations, you have to be able to review the entire process and trace it back to verified sources and insights. Just as humans must be able to show how they arrived at a certain decision or conclusion, AI's "thought" process must also be documented for accountability purposes, and to ensure that the system is operating as intended. That's why your AI SOC solution must produce a clean evidence trail. Being able to document, review, and audit the process is table stakes for any AI-powered SOC.

Is the AI safe and secure?

Finally, it's important to opt for a product that has security, compliance, and AI safety at its core. To more effectively reduce risk and exposure so you can respond to increasingly sophisticated threats and attacks, safety must be foundational to any reliable AI SOC vendor and product.

As a baseline, the vendor should never use your data to train its AI.

As a baseline, the vendor should never use your data to train its AI. You shouldn't have to tune the AI or request this as an extra feature. It should be a guarantee. While contextual organizational awareness is key to effective investigations, it's highly risky if it becomes accessible to anyone else.

End-to-end encryption at rest, in-transit, and in storage should be a requirement. In the same vein, if the vendor provides single tenancy, there's no potential for your data to cross paths with other customers'. If your security requires an on-prem setup, you'll need a vendor that supports air-gapped, self-managed deployments.

It's important to understand what monitoring controls are in place, how the product is deployed, and if the provider conducts regular red teaming testing.

Ask how the vendor AI architecture protects your network, applications, and data. It's important to understand what monitoring controls are in place, how the product is deployed (self-managed, on-prem, cloud, SaaS, or hybrid), and if the provider conducts regular red teaming testing.

What about access and identity? Single sign on (SSO) and multi-factor authentication (MFA) are great starting points. Best practices are to control access with an identity provider (IDP), a common access card (CAC), or personal identity verification (PIV).

Just like user access is role-based, the AI should also have controlled access. Ideally, the product should be able to have the same access as the users.

To avoid expensive data migration and extraction delays, ask if the vendor requires ETL or if it ensures that only the minimal data needed for the task at hand is temporarily stored in the deployment environment. This reduces complexity, minimizes exposure, and improves security.

Use the checklist to compare AI SOC products →

The AI SOC product checklist

With the relentless proliferation of increasingly sophisticated threats, you need to empower your cybersecurity teams, not complicate their tech stack even further.

That's why Andesite has redesigned the architecture of the cyber defense ecosystem with a human-AI collaboration layer that operates above, beyond, across, around, and with your existing data sources, tools, and platforms.

We call it the decision layer, and it's where analysts can access, visualize, and work with all the elements and data feeds in your security ecosystem as it is.

Analyze structured and unstructured data with no ETL, survey the incessant wave of alerts, and automate investigations and enrichment to deliver relevant, contextualized, and actionable insights.

Here's a checklist that you can use to compare AI SOC products—and if your SecOps team's needs check the boxes below, Andesite may be the right solution for your organization.

We are looking for an AI SOC tool that:

- □ Keeps humans at the helm instead of being fully autonomous
- □ Has flexible architecture that adapts to your ecosystem
- □ Provides insights
- □ Uses organizational context to assess threats within specific risk profiles
- □ Prioritizes and groups alerts
- □ Processes threat intel and other unstructured enrichment sources from uploads and URLs in minutes
- □ Combines several alert and threat intel sources within the scope of an investigation
- □ Doesn't require ETL or data migration
- □ Doesn't use customer data to train the AI
- □ It is audit-ready
- □ Meets compliance-high standards, controls, and frameworks
- □ The evidence trail of AI-driven investigations be traced back to verified sources
- □ Ensures end-to-end encryption
- □ Offers access and identity safeguards and security
- □ The architecture is safe and secure and can adapt to your security ecosystem

Ready to learn more about the Human + AI SOC?

Book a Demo →

ANDESITE

Learn more at andesite.ai, visit our Trust Center, and reach out at contactus@andesite.ai