A CISO's Guide to the New Era of Agentic AI

The AI boom of recent years sent out a shockwave that's still expanding, upending old ways of working and uncovering new opportunities. And while it's inevitable that this explosion of innovation will eventually level out, a new development in AI technology has captured the world's attention: agentic AI.

If 2023 was the year of the chatbot and 2024 was the year of the copilot, 2025 will be the year of the AI agent. This chronology shows how quickly AI has evolved, with significant advances being made in an astonishingly short time.

Is agentic AI just one more landmark along this longer road? Or is it a destination in its own right, with the potential to deliver meaningful and lasting change for business? Industry analysts such as Gartner seem to be forecasting the latter, reporting on the tangible value that agentic AI can contribute to organizations and predicting that 60% of all IT operations tools will feature this technology by 2028.

Of course, cybersecurity is also a technology-focused discipline, so the question on all CISOs' minds right now is how agentic AI—and AI agents, which are tools and applications powered by agentic AI—can potentially improve the security operations center (SOC).

To answer that question, it's essential to establish coherent definitions and parameters around agentic AI, as well as realistic expectations of what it can, and should, be tasked to do.

Mission Critical or Mission Cynical?

Most of the latest breakthroughs in AI technology have their origins in the evolving capabilities of large language models (LLMs), which builds on machine learning (ML) and natural language processing (NLP) technologies to produce AI assistants that communicate intuitively with users, exhibit increasingly astute reasoning and contextual awareness, and come up with unique responses and content.

As these capabilities have manifested in a quick succession of new tools, cybersecurity leaders and decision makers have been confronted with a plethora of potential vendors who overstate or oversimplify the ways they integrate these technologies into their products. It's one thing for a sales team to rebrand an existing offering as "AI-enabled" and put together a flashy demo, and another for engineers to build something novel that delivers sophisticated, real-world value.

Without extensive trialing and testing, it's not easy to tell the difference. And that's why, following a wave of AI hype, organizations have too often been left with solutions that fail to deliver measurable outcomes. But when they can cut through the conjecture and identify next-gen AI tools that address real needs, it can be transformative.

High Stakes for AI Selection and Deployment

Now CISOs are going through this again with the promise of AI agents, trying to read between the lines to see which offerings can genuinely move the dial on their strategic priorities, and what's just marketing. And unlike other executives, whose main preoccupations may be efficiency or user experience, CISOs also must understand how new technologies can be securely deployed within the enterprise environment.

They don't just have to be smart when vetting new tools and vendors; they also need to be fast. There's pressure from the business, competitors, and customers, along with the added pressure of knowing that any new technology that augments, automates, streamlines, and scales workflows is certainly going to be exploited by adversaries.

In cybersecurity, FOMO (fear of missing out) is about much more than wanting the latest gadget. It's about being able to foresee and forestall an ever-changing litany of adversary tactics and techniques.

Lessons Learned from the Last AI Hype Cycle

If the hype cycles around generative AI have taught senior decision makers anything, it's that the risk of disappointment and disillusionment is real. Vendors can talk about the value they provide, but in many cases their products were rushed to market and may just be a slightly upgraded version of what they offered before, presented with a fancy new wrapper.

Also, if they offer AI solutions as expensive add-ons, it means that AI isn't comprehensively integrated within their platform products, raising red flags around functionality, compatibility, security, and user experience.

To ensure a more successful AI adoption this time around, CISOs should prioritize vendors that have fully embedded their AI agents into workflows in ways that tangibly advance their key business objectives. And they should have robust safeguards and documentation to show how it ingests, retains, and destroys data so that it complies with data privacy and protection regulations.

The CISO's Strategic Priorities

Broadly speaking, agentic AI differs from the generative AI solutions that preceded it due to its enhanced capacities for decision making, problem solving, and acting autonomously within an organization's environment. But before getting too technical in defining this technology, let's consider the core needs and functions that CISOs strive to fulfill, and whether these gaps can be closed by technology.

1. The Volume of Attacks is Unprecedented

Many existing threat detection mechanisms that the SOC depends on are based on known threats. But in the age of AI, there will inevitably be a massive increase in threats that don't have a signature, and that the typical threat intelligence feed hasn't encountered before. This is already the reality, with a near-constant barrage of AI-generated attacks. The scale of this problem is beyond anything humans can handle without technological aid.

Today, it takes AI to defend against AI, so CISOs will want to evaluate the capacity for AI agents to expedite the SOC's mean time to respond/remediate (MTTR). When it comes to KPIs, MTTR is the metric that matters more than anything to cybersecurity leaders, and most SOCs will have built up an arsenal of trusted tools and workflows to support threat detection, investigation, and response (TDIR).

While existing systems such as user and entity behavior analytics (UEBA) are excellent for the "TD" part of this equation, accelerating the "IR" is where AI—and specifically AI agents—have incredible potential to ask the right questions of the right data and automatically recommend and queue up the next steps for analysts.

AI Agents: The Emerging Insider Threat

AI agents are quickly becoming part of day-to-day business operations. They can interact with enterprise systems, access sensitive data, and even take autonomous actions on behalf of your workforce. But this new efficiency introduces a new category of insider threat. Unlike external attackers, AI agents already operate with valid credentials and access rights. A misaligned, malfunctioning, hijacked, or jailbroken agent can create the same risks as a malicious insider, including data leakage, privilege misuse, inappropriate information disclosure, or unauthorized system changes, while often going unnoticed by traditional security tools.

How Exabeam Helps Secure AI Agents

Exabeam addresses this challenge by extending behavioral analytics and risk-based detection. By ingesting data from platforms like Google Cloud Agentspace, Google Cloud Model Armor, and custom agents built with Vertex AI or Agent Builder, Exabeam baselines both human and AI activity. Our Exabeam Nova AI agents score, explain, and prioritize anomalies so security teams can see when an AI agent deviates from its expected behavior and respond quickly. This approach closes a critical visibility gap, helping you secure your digital workforce.

2. The Amount of Data Collected is Immense

Modern enterprises collect and consume more data than ever before, yielding billions of log file lines each day. Each of these log files contains dozens of pieces of information that could be relevant to threat hunters and provide clues about whether or not an event should be investigated.

Distilling that data into manageable insights drives much of AI's necessity in cybersecurity, and it's been this way for years. ML forms the foundation on which most AI is built and is essentially the development of algorithms that can be trained to recognize patterns and identify anomalies in massive troves of information. Chatbots and copilots have expanded upon this using NLP and, more recently, LLMs to provide search functionality and explainers in plain language.

Now CISOs are rightly focused on how AI agents can take this a step further, empowering Tier 1 analysts to make informed decisions and take appropriate actions without having to escalate to a Tier 3 analyst every time there's a suspicious incident.

As CISOs know all too well, the average security analyst may have to assess 100 cases a day, when ideally, they should be dedicating their time and attention to the most important 15 to 20. If new AI tools are going to be integrated into the SOC, leaders and employees are no doubt hoping that it can help them triage alerts.

However, should AI act on those alerts? There can be serious upstream and downstream consequences whenever TDIR actions are taken, so security leaders will need to consider how much autonomy they want AI agents to have.

At this juncture, perhaps the best solutions will be those that automatically offer suggestions, options, and playbooks for investigation and response when a threat is detected but ultimately cede decision-making authority to the human user.

3. The Need to Nurture SOC Talent is Urgent

The talent shortage in cybersecurity is a perennial problem for CISOs. Turnover and burnout are huge problems in the industry, and Tier 3 SOC analysts are difficult to find.

Among the top causes of burnout for cybersecurity professionals are those skyrocketing rates of cybercrime, with adversaries producing more malicious code than ever by abusing generative AI. Alert fatigue takes its toll on SOC analysts, who dismiss an estimated 62% of alerts as they try to distinguish true positives from false. As a result, 84% of cybersecurity professionals say that they have experienced burnout, which is the main culprit for 50% of resignations.

Every time an analyst quits, it compromises the organization's security operations. It can take months for even a Tier 1 analyst to get a good understanding of the unique ecosystem of tools, systems, processes, and skills required by their SOC. This presents a two-pronged priority for CISOs:

1. How can workloads and workflows be streamlined for security analysts so that they have a better day-to-day experience and remain in their roles?
2. And how can resources be adopted and allocated so that the SOC can do more with less, and Tier 1 analysts can operate more like Tier 3 analysts?

Answering these questions is crucial, especially considering 67% of organizations report a moderate to critical skills gap. If deployed effectively, agentic AI could supplement many of these skills, while empowering the current workforce of SOC analysts to genuinely enjoy their jobs again.

4. The SOC Must Prove Its Value

As the scale of cybercrime overwhelms the SOC's capabilities, 77% of security leaders worry that a serious breach could spell the end of their careers. Consequently, the continuous improvement of the organization's security posture is another area of emphasis.

The work of leveling up the cybersecurity operation is never complete. CISOs want to be able to track and measure how the SOC is performing; they also want to be able to demonstrate the value and impact of their security operations on the business. The SOC delivers measurable ROI; it's a strategic asset, not an afterthought. But proving it can be a challenge.

So, what's the potential for AI agents to support in these areas? These systems have access to data, insight into trends and patterns, and the ability to articulate it in an intuitive and actionable way. Can agentic AI assume the role of consultant, and support the SOC in strategic capacities above and beyond the purposes it serves in TDIR?

A Pragmatic Definition of Agentic AI

To truly understand how agentic AI can support CISOs and SOCs in meeting their core priorities, it's important to distinguish the key attributes that differentiate agentic AI from earlier generative AI tools. There are three primary qualities that set it apart:

1. **Agentic AI is active and proactive.** These systems perceive and understand the problem you're trying to solve, then act autonomously or semi-autonomously to address it. This proactive approach enables faster responses and more dynamic solutions compared to traditional software development, which typically requires manual input and slower updates.
2. **Agentic AI has greater reasoning capabilities.** Agentic AI can dynamically comprehend context and adapt its behavior to solve complex, multi-phase tasks. Its continuous learning enhances its ability to tackle evolving security challenges with precision.
3. **Agentic AI coordinates and cooperates.** One of the most promising aspects of agentic AI is its ability to enable multi-agent workflows. Tasks are distributed across numerous, highly specialized agents, creating a flexible, modular, machine-driven workforce.

When LLM technology is used in a proactive, agentic way, it allows organizations to deploy adaptive solutions faster. Conventional methods can be slow and rigid, while agentic AI is flexible, efficient, and responsive to evolving threats and use cases, with minimal human intervention.

Though multi-agent ecosystems still require human oversight, this level of automation and coordination is already a reality. In the Exabeam New-Scale Security Operations Platform, an AI agent correlates activity in Threat Center with risk levels to provide analysts with real-time incident and investigation summaries. In Outcomes Navigator, AI helps managers and executives evaluate SOC effectiveness.

These agents, though designed for different use cases and personas, from Tier 1 analysts to CISOs, work synergistically to create a more efficient, intelligent security operations environment.

Deploying Agentic AI the Smart Way

Unlike other industries, cybersecurity already has incredibly specific and intrinsic use cases for technology like agentic AI. Harnessing machine intelligence has been integral to TDIR for more than a decade, so it makes sense that the most powerful applications of agentic AI for cybersecurity will build on these deep foundations.

Exabeam pioneered UEBA, deploying ML models that were revolutionary: They could detect complex threats, score and prioritize risk, track the lateral progression of events and attack chains without gaps, and establish baseline normal behavior. All of these functions remain relevant. While LLMs excel in certain areas, they don't have the ability to turn vast volumes of raw, complex data into intelligent detections. ML, particularly through behavioral analytics, is still critical for extracting insights from that data and identifying anomalies that traditional approaches miss.

With the introduction of Threat Center, Exabeam further evolved its integration of emerging AI. Initially an analyst workbench that provides a single location to manage cases, detections, and alerts, AI chatbots and copilots could be layered into this tool to provide threat summaries and analyst assistants in an intuitive, intentional way that delivered tangible value.

While LLMs may not give good results when having to extract insight and make decisions based on vast troves of data, they're excellent at language translation in a contained, specific context. Hence, the AI innovations from Exabeam were met with very high satisfaction ratings among customers—a notable exception from other generative AI offerings that failed to live up to the hype. These AI capabilities were also provided at no cost to the customer; they were integrated proactively into products and delivered as part of the existing license.

Agentic AI will follow a similar path. It will thrive if it's built on a strong foundation, with use cases that address real SOC needs in ways that make sense.

Agentic AI for the SOC Is Here, and It's Transformative

The capabilities of AI agents may sound futuristic, but they're here now. Exabeam Nova, the AI agent from Exabeam, supports the SOC in many key areas. The now-legacy Exabeam Copilot introduced popular threat summaries that gave analysts a quick, accessible view of incidents. Exabeam Nova builds on that foundation with full investigation summaries: comprehensive, analyst-ready reports that include threat classification, prioritization, supporting evidence, and remediation steps.

This evolution gives Tier 1 and Tier 2 analysts the context and depth they need to make informed decisions without having to escalate every case, all because Exabeam Nova provides more than 10 times the amount of data to the Large Language Model (LLM) compared to the original Copilot.

Threat explainer capabilities are enhanced so that the AI agent not only summarizes the details of an incident but elaborates on its implications and why it may be cause for concern. It's also advanced enough that it will tell the user if it lacks sufficient information to come to an accurate conclusion and even identify the information that it would require, unlike more rudimentary AI assistants that will try to answer a query no matter what.

Exabeam Nova also demonstrates the potential of multi-agent workflows for cybersecurity because different Exabeam Nova agents are embedded in different capabilities, and work together to support the SOC:

- **Advisor Agent:** Delivers daily reports on security posture, MITRE ATT&CK® coverage, and detection outcomes. It identifies weak points, recommends improvements, and generates leadership-ready summaries for CISOs and security managers. It also supports roadmap planning by simulating the impact of proposed changes on coverage and posture.
- **Search Agent:** Enables analysts to query security data using natural language (in any language), eliminating the need to write complex query syntax. This accelerates search and lowers the barrier to entry for Tier 1 analysts.
- **Visualization Agent:** Converts queries into clear visual dashboards, timelines, and trend reports. These visualizations make it easier to understand SOC performance, monitor changes over time, and communicate findings to stakeholders.
- **Threat Scoring Agent:** Applies adaptive learning and contextual analysis to prioritize alerts and highlight high-risk sessions. It calculates multi-layered risk scores using behavioral baselines, event rarity, and business context, helping analysts focus on the most critical threats.
- **Investigation Agent:** Automatically generates case titles, summaries, threat vectors, and recommended next steps. This agent eliminates manual reporting, helping analysts of all skill levels respond faster and more accurately.
- **Analyst Assistant Agent:** A real-time, context-aware companion that supports investigation workflows. It provides relevant evidence, explains anomalies, and guides analysts through complex cases, reducing escalation and burnout.

The AI agents also integrate with other platform capabilities, pulling entity information from Attack Surface Insights and enhancing Threat Center timelines with chronological context and improved detection accuracy.

All of this leads to cost and time savings, and a drastically improved experience for enterprise SOC teams, which, given the pace and complexity of today's cyberthreats, is exactly what's needed most.

A New Generation of Security Consultants for the CISO

Taking it further, Exabeam Nova can operate in an advisory capacity so CISOs can collect metrics and proof points to both continuously improve the SOC and demonstrate impact to other senior stakeholders.

Consider Outcomes Navigator. This GenAI-assisted integration, introduced in 2023, assesses an organization's security posture against the most prevalent techniques in the MITRE ATT&CK® framework, assigning a grade. It quickly proved to be one of the most popular products in the New-Scale Platform, but now the inclusion of agentic AI has advanced it even further. CISOs can track the performance and progress of the security apparatus over time while the AI acts as a consultant, automatically providing executive-level reports and summaries of how security coverage can be improved, which would have taken days for a human to research and write.

Best Practices Yield the Best Results

In this next era of AI adoption and deployment, CISOs are going to have to navigate yet another cycle of hype, vetting vendors conscientiously to ensure that their agentic AI capabilities are as good as what they see in demos, screenshots, and product roadmaps. In times like these, the smartest choice is often to go with organizations that have a long and proven track record as AI innovators and pioneers, that have been dedicated to addressing specific cybersecurity needs, and that lead the pack in industry reports like the Gartner Magic Quadrant.

Exabeam has been an AI company for over a decade, and this makes it unique in the market. The data processing and pattern recognition of ML supports the baselining of normal activity by UEBA, which in turn informs TDIR dashboards and timelines; these high-fidelity detections are then leveraged by generative AI to drive real value for satisfied customers. Agentic AI now adds a new level, sustained by all the others that predate it.

All AI data processing by Exabeam takes place securely within New-Scale Analytics, which operates on the Google Cloud Platform (GCP) and uses Google Gemini LLMs. All LLMs are pre-trained, which means they never learn from customer data, and all data retrieved by the system is ephemeral. It's encrypted using standard TLS protocols, and the entire evaluation process occurs in memory so that data is never stored, cached, or retained.

Combine that with the fact that model processing is always performed within the customer's designated geographic region whenever possible, and Exabeam ensures the data privacy, regulatory compliance, and optimized speed and scale that modern enterprises depend on.

As a trusted trailblazer in this space, Exabeam is excited to introduce the world to AI agents, and the varied ways they can revolutionize workflows and maximize value for CISOs, enterprise security teams, and the businesses they defend.

About Exabeam

Exabeam is a leader in intelligence and automation that powers security operations for the world's smartest companies. As a global cybersecurity innovator, Exabeam provides industry-proven, security-focused, and flexible solutions for faster, more accurate threat detection, investigation, and response (TDIR).

exabeam

Learn more at www.exabeam.com →