




WHITE PAPER

Understanding Agentic Security: Beyond the Hype



What's in a Name?	2
Agentic AI	3
Agentic Security	4
Technology Comparison	5
Benefits of Agentic Security	7
Risks of Agentic Security	9
How to Compare AI Agents	10
Why Embed Security	11
Conclusion	12



What's in a Name?

Have you tried the latest AI SOC Analyst yet? What about AI-driven hyperautomation? Or plain old SOAR that can make API calls out to LLMs? Are multi-agent swarms really multi-agent if they're using the same underlying language model?

We get it. It can be incredibly frustrating to listen to the noise (not to mention the over-inflated promises) around AI for Security Operations and suss out what's really happening "under the hood" and what truly matters.

We've written this report to help cut through the alphabet soup. We want to provide a frame of reference that we hope is useful for the cybersecurity community. Our objective is to define key terms in plain English and to pinpoint where and how the benefits of AI will go beyond previous security automation technologies. At Embed, we focus on alert triage and investigation, so that's the lens we'll use here versus other functions that security teams perform.



Agentic AI

Given all the hype around AI agents, you might be surprised to learn that agents have been around for decades, going back to the 1960s. The reason they're making headlines now is that they can work alongside large language models (LLMs) to tackle a **wide variety of tasks**. Previously, AI agents were designed for specific jobs like playing chess, but they wouldn't operate beyond that particular problem. Today, because LLMs inherently capture a broad base of knowledge, AI agents leveraging them can naturally take on more challenges.

So what exactly is Agentic AI?

Agentic AI systems combine autonomous perception, reasoning, and action to navigate complex environments and solve multi-step problems.

There are a few key attributes that make up Agentic AI systems:

1. **Perception** - agents can gather data from the environment using a variety of means (e.g., sensors, databases, APIs, etc.) and understand useful features of the data.
2. **Reasoning** - agents can evaluate options to determine how to solve complex problems.
3. **Action** - agents can take actions using tools (e.g., APIs and other interfaces) which have an effect within the environment.
4. **Learning** - agents can adapt over time and improve their performance.

Given the generality of the agentic approach, it's no surprise it can be applied to many different problems.



Agentic Security

When we say Agentic Security, we mean **applying Agentic AI in the cybersecurity domain**. There are many important cybersecurity problems that can be addressed with AI agents. For example, patch management and pen testing are two such complex problems where agent-based automation could be effective.

Our focus at Embed Security is on **alert triage and investigation**, a significant problem every organization faces. Given a set of alerts from products like EDR, cloud, and email security systems, security teams must quickly determine whether each alert is real and actionable. While most alerts turn out to be false positives, it only takes one missed threat to compromise an organization.

Investigating alerts is a natural fit for the agentic approach because it mirrors how human analysts work with perception, reasoning and action.



The learning aspect is particularly valuable - agents can improve their investigation process based on feedback from experts. The multi-step nature of alert investigation, where each step informs what to do next, aligns perfectly with how agents operate. For example, an agent might start with an alert about suspicious network traffic, decide to investigate the source IP's history, discover it's associated with previous incidents, and then automatically pivot to examining other systems that communicated with that IP - all while documenting its findings and reasoning along the way.

Technology Comparison

Ok, now that we've talked about agents, let's take a step back and clarify where exactly agents will improve the alert investigation process. To do that, it's helpful to break down the different types of automation approaches and how they map to the important aspects of alert investigation (for more about the history of security automation, check out our previous blog post). Let's consider four types of automation:

1 Rule-based playbooks

These are the traditional SOAR playbooks that follow pre-defined, if/then logic. They're great for standardizing known processes and handling predictable scenarios, but they're rigid and can't adapt to novel situations.

2 AI-assisted analysis

These leverage machine learning or AI models for specific tasks like pattern recognition or anomaly detection. While these algorithms provide insight, they typically have limited context awareness and focus on single-step decisions rather than end-to end processes.

3 AI Copilots

These interaction assistance tools help analysts by suggesting next steps or gathering relevant data. They don't work autonomously though, requiring constant human guidance and decision-making.

4 Agentic AI

This combines the process automation capabilities of playbooks with the intelligence and adaptability of AI, while adding crucial capabilities like autonomous planning and dynamic decision-making.

Let's look at how these approaches handle alert investigation aspects.

Remediation Actions	Performs low-risk responses based on exact matches to predefined conditions. Requires explicit rules for each possible scenario.	Can make limited, recommended responses based on well-defined threat scenarios.	Suggests possible contextual responses.	Can propose complex, contextual response sequences.
Reasoning	Relies on simple if/then logic with pre-programmed decision trees.	Makes isolated predictions based on trained patterns. Lacks ability to connect insights across multiple steps.	Provides analysis suggestions and helps analysts connect dots.	Makes contextualized decisions incorporating multiple factors, including past experiences.
Investigative Actions	Limited to simple, pre-configured data gathering steps in a fixed sequence.	Can help parse and understand data but requires human direction.	Can help gather relevant data based on analyst input.	Autonomously executes multi-step investigations, choosing appropriate tools and pivoting based on findings.
Planning	Follows predefined, static workflows. Cannot adjust plans based on new information uncovered during execution.	Typically not used for planning.	Can help analysts develop and adapt investigation plans but requires human guidance.	Can dynamically create and adjust investigation plans based on findings.
	Rule-based Playbooks	AI-assisted Analysis	AI Copilot	Agentic AI

Table 1 - Alert Investigation Aspects vs Security Automation Technologies

Agentic AI represents a significant leap forward in handling the complexity of modern alert investigation. While rule-based systems automate simple tasks and AI models can spot patterns, only agents can truly mirror the way skilled analysts work - adapting their approach based on what they discover, making nuanced decisions, and executing appropriate responses. While copilots assist security analysts by collaborating on alert investigations, AI agents can autonomously triage and investigate high volumes of security alerts independently, only escalating significant findings that require human expertise.

This doesn't mean agents will replace human analysts (see our [previous blog post](#)). Instead, they serve as force multipliers, handling routine investigations autonomously while escalating complex cases to human experts with rich context and preliminary analysis. This allows security teams to focus their expertise where it matters most - on complex threats and strategic security improvements.

Benefits of Agentic Security

The adoption of Agentic AI for alert triage and investigation brings several key benefits to security operations teams. Let's explore the advantages that make this technology so important:

Enhanced Threat Detection Accuracy

Agentic AI significantly improves the identification of true threats by conducting thorough, consistent investigations of every alert using security best practices. Unlike traditional approaches, agents can simultaneously analyze multiple data sources, maintaining complex context throughout the investigation process. Where human analysts might check 2-3 data sources due to time constraints, AI agents can methodically examine dozens, cross-reference findings, and identify subtle connections that indicate real threats. This comprehensive approach can lead to more accurate threat identification and fewer missed indicators of compromise.

Reduced Analyst Workload and Burnout

Security teams today face an overwhelming volume of alerts, leading to analyst burnout and high turnover. Agentic AI addresses this challenge by providing 24/7 autonomous alert triage, effectively eliminating security noise from false positives. When escalation is necessary, agents provide rich context to human analysts, allowing them to focus their expertise on strategic work and complex investigations rather than routine tasks.

Standardized, Consistent Processing

One of the most significant advantages of Agentic Security is its consistency. Unlike human analysts who may approach investigations differently or skip steps when overwhelmed, agents maintain the same high standard of investigation for every alert. This standardization ensures complete audit trails, enables accurate metrics tracking, and helps organizations maintain compliance requirements across all shifts.

Faster Response Times at Lower Cost

The autonomous operation and consistent processing of Agentic Security leads to dramatic improvements in operational efficiency. Organizations typically see 50-90% reductions in investigation time depending on the alert type, while handling ten times more alerts with the same team size. These benefits compound over time as agents learn from experience, continuously improving their performance while maintaining consistent coverage. The result is a more effective, efficient, and sustainable security operation that can better protect the organization without expanding headcount or burning out human analysts.



Risks of Agentic Security

While the benefits of Agentic AI are compelling, we must also understand the risks. Let's examine some concerns when implementing AI agents in security operations

Domain-Specific Accuracy

Security investigation requires domain expertise and precise knowledge of technical concepts. AI agents must demonstrate consistent accuracy in security-specific tasks, not just general language capabilities. Misinterpretation of security concepts or failure to recognize subtle indicators of compromise could lead to missed threats or false escalations. Agents should be trained on cybersecurity concepts and investigation processes.

Transparency

The opaque nature of some AI systems poses risks in security operations, where decisions need to be auditable and defensible. Security teams must be able to understand agent reasoning. This includes visibility into what data sources were accessed, which investigation steps were taken, and how conclusions were reached.

Tool Access and Potential Abuse

AI agents require access to various security tools and systems to conduct investigations effectively. Organizations need to ensure agents have appropriate (least) privileges, action limitations and rate limiting to prevent overloading a resource. Audit logs should capture all agent actions. If an agent can use a tool to access information that might be compromised, guardrails should be in place.

If you want to dive deeper into risks of AI agents, we recommend reading the recently published [OWASP guide on threats to Agentic AI](#).

How to Compare AI Agents

As more vendors enter the market claiming to offer AI agents for security, it's important to understand how to evaluate these solutions effectively. Two fundamental aspects should guide your assessment: accuracy and explainability.

While vendors can showcase impressive demos, real-world accuracy in cybersecurity requires (1) domain-specific training that incorporates cybersecurity investigation knowledge, (2) performance metrics assessing false positive and negative rates, and (3) agent-based learning to understand how the system can be expected to improve over time. Note that learning is actually a challenging problem in cybersecurity. Expert labels for alerts are hard to come by, and even if you had them, the value decreases over time as attackers and defenders co-evolve.

Equally important is trust. Security teams need to establish and then maintain trust in Agentic AI decisions. Embed achieves this through transparency in describing the evidence gathered and questions answered during the investigation. The agent's reasoning should align with analyst intuition. This transparency should extend beyond simple decision explanations to include the full investigative workflow. That means security analysts must be able to see how the agent arrived at its findings, including what data sources were used.



Why Embed Security

While many companies are rushing to add AI capabilities to their security products, Embed Security was purpose-built from the ground up for Agentic Security. Our founding team brings together decades of real-world experience in both cybersecurity and artificial intelligence – a rare combination that gives us unique insight into what actually works in production environments.

Our founders, Seth Summersett and Jeff Johns, have spent their careers at the forefront of cybersecurity and AI innovation. With a combined 40 years of experience at organizations like NSA, Mandiant, FireEye, Meta, and Google, they've seen firsthand the evolution of both threats and defensive capabilities. Jeff brings over two decades of expertise in AI/ML, having developed and deployed award-winning machine learning models that protect millions of assets globally. He also has done novel research in the field of reinforcement learning during his PhD.

This deep experience matters because building effective security agents requires more than just connecting to an LLM API. It demands several skills:

- Deep cybersecurity expertise to train agents on real investigation processes,
- Production AI/ML experience to build reliable, scalable systems,
- Understanding of SOC workflows to integrate seamlessly with analysts.



Conclusion

While AI has been gradually making its way into security tools over the past decade, the emergence of Agentic AI represents a fundamental shift in how we can approach alert investigation. This isn't just another incremental improvement in automation - it's a qualitative leap forward in our ability to handle the complexity and scale of security operations.

However, we need to maintain perspective. Despite the transformative potential of AI agents, they won't replace human analysts or solve all security challenges overnight. What they will do is dramatically improve our ability to handle the growing volume of alerts while maintaining high investigation quality, allowing human analysts to focus their expertise where it matters most.

As you evaluate AI solutions for your security operations, we encourage you to look beyond the marketing hype and focus on what matters: proven accuracy in real-world environments, transparent reasoning that builds trust, and the ability to integrate with your existing processes. Look for partners who understand both the technical capabilities and limitations of AI, as well as the practical realities of security operations.

At Embed Security, we're committed to bringing the benefits of Agentic AI to security teams in a way that's practical, proven, and built on real-world experience. We believe the future of security operations lies in the thoughtful combination of human expertise and AI capabilities, and we're excited to be at the forefront of this transformation.

If you're interested in learning more about how Agentic AI can improve your security operations, we'd love to continue the conversation.

Reach out to us [here](#) to discuss your needs and see our technology in action.