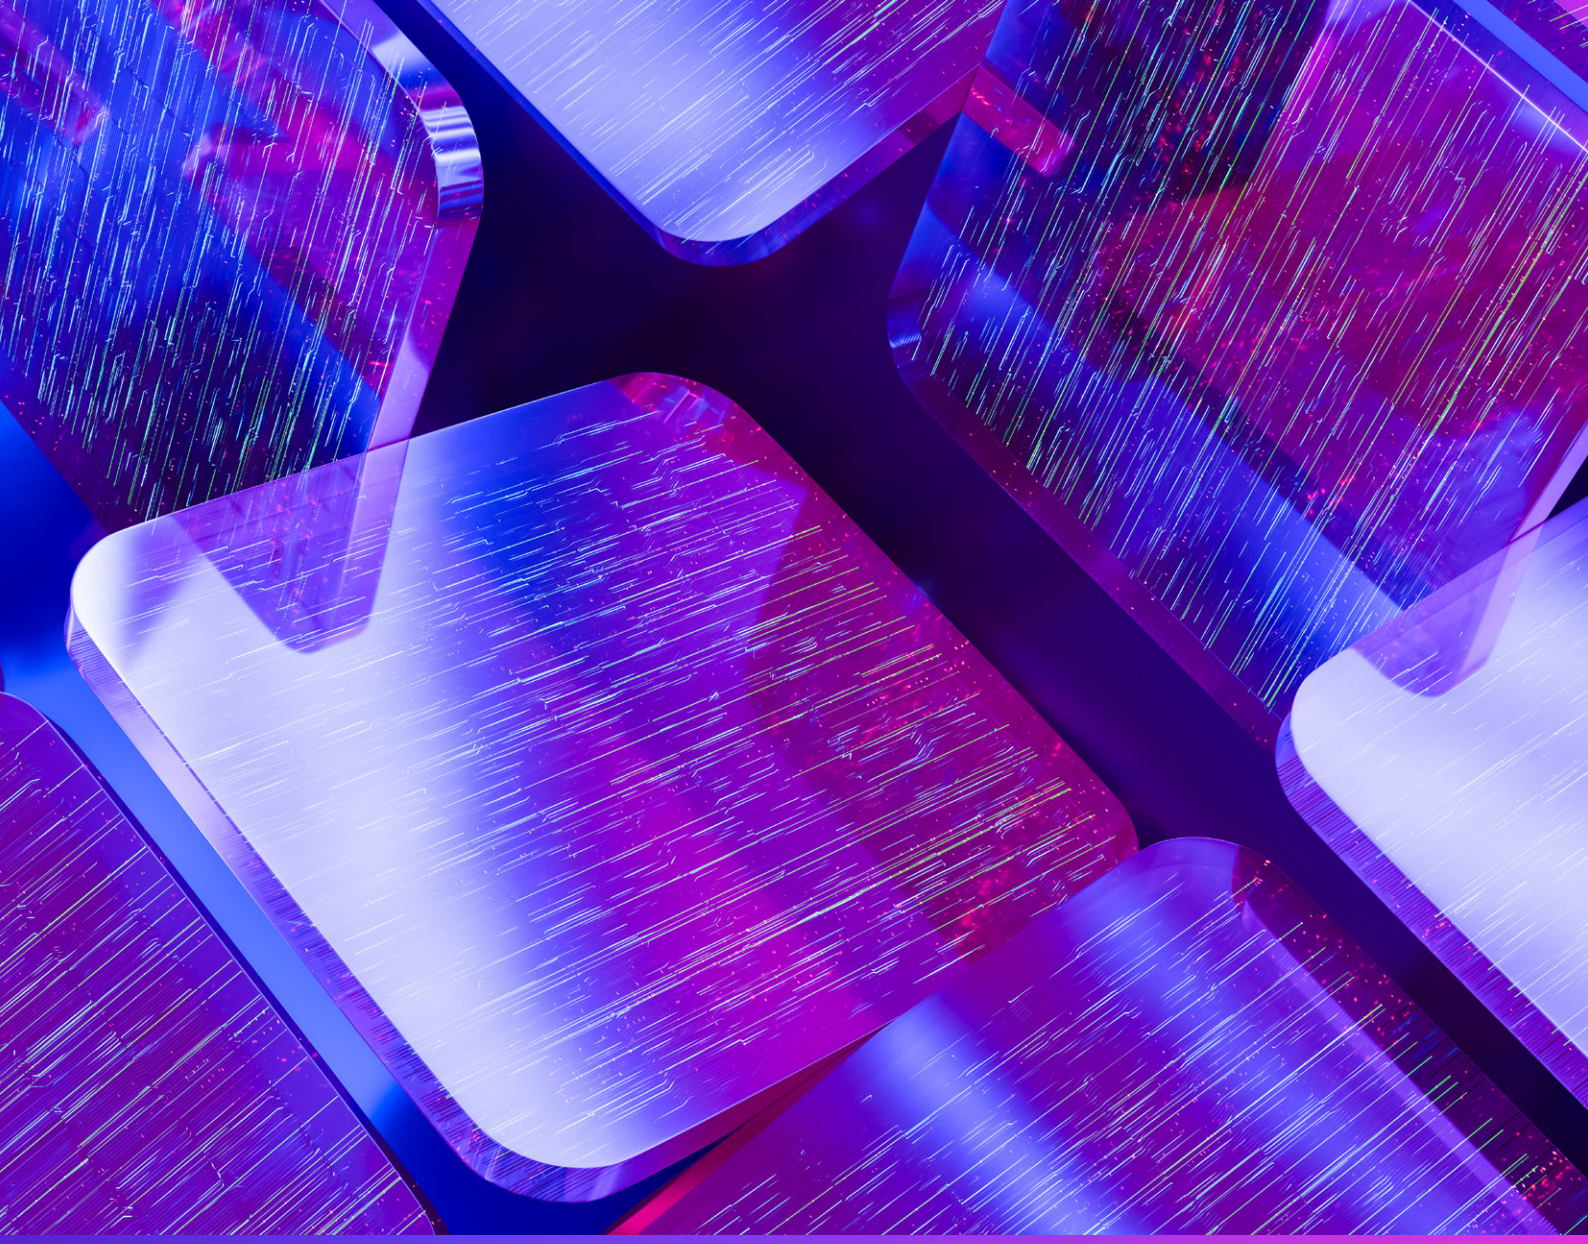

WHITE PAPER

Cutting Through the Hype What Agentic AI Really Means and the Future of Security Operations

Ontinue

Ontinue

Agentic Artificial Intelligence (AI) has taken the cybersecurity spotlight, quickly becoming the buzzword of the year, replacing "generative AI" in conference keynotes and vendor roadmaps alike. While many organizations are eager to label/relabel their solutions "agentic," there is a lot of confusion about AI technologies, the difference between their capabilities, and how they augment human intelligence in a Security Operations Center (SOC).

At Ontinue, we use agentic AI (as well as other types of AI and automation) to deliver managed Security Operations (SecOps) at scale with the speed and quality required to alleviate burden on in-house security teams defending their organizations from today's threat landscape. In this whitepaper, we break down the spectrum of AI and its evolution, what makes agentic AI agentic, and how Ontinue is using the technology to improve security outcomes for customers.

The whitepaper also explains the future of agentic AI and the inevitable paradigm shift that will occur between managed SOCs, customers, and third-party tools and technologies.

The Evolution of Automation in Security Operations

For decades, the cybersecurity industry has embraced automation as the answer to alert overload and limited staffing. Deterministic workflows, scripted playbooks, and the use of machine learning have helped security experts respond faster to cyberattacks. These tools filter noise, flag anomalies, detect known threats, and can trigger prescriptive automated response actions in specific situations. Upon the success of this wave, the industry next started using generative AI.

In cybersecurity, generative AI has primarily been used to a) translate plain-language questions into structured queries, and b) generate recommendations for next steps during incident investigations. For example, a non-expert might ask, "What browsers were used last week?" Generative AI can help write the complex KQL or SQL needed to interrogate logs, filter for specific header fields like User-Agent, and return a usable result. It synthesizes data based on statistical patterns. This reactive assistance has lowered the technical bar for many SecOps tasks, including incident investigation and summarization, and made tooling more accessible.

The advent of generative AI technology and its evolution also gave rise to agentic AI. Built on the foundation of generative AI, agentic AI introduced cognitive autonomy: the ability to interpret intent, plan multi-step actions, adapt dynamically to changing information, and enable autonomous actions toward defined goals.

It's important to note up front that each of these technologies have important applications in a modern SOC. Like any tool, the key to success is applying the right tool to the right challenge.

To further understand this progression and the use cases for each of these three technologies, consider the following summary and diagram, in basic terms:

Deterministic Automation handles repetitive, rules-based actions. It's fast and precise, but rigid—unable to handle ambiguity or novel scenarios. Its application is ideal for resolving recurring events.

Generative AI Assistants bridge the gap between human language and machine logic. The technology assists and accelerates queries and reporting based on user prompts. Generative AI assistants are LLM-based task-oriented tools that can generate content, such as SQL queries, answers to questions, and recommendations in specific situations.

Agentic AI, *enabled* by generative AI, provides cognitive autonomy. It doesn't wait to be told what to ask—it figures out the right questions, answers them, adapts, and solves problems in real time all *proactively*.

Demystifying the Hype

With the introduction of agentic AI, many cybersecurity vendors have hastened to relabel Generative AI assistants as "AI agents." This has created a confusing situation in which AI agents now fall on a spectrum of autonomy. It has also sowed doubt about the credibility of truly agentic systems, capable of autonomously carrying out complex tasks that once fully required humans to handle.

It's important to remember that generative AI assistants are designed to carry out specific tasks, i.e. they are prompted with a specific objective and produce outputs, but their internal logic is coded and therefore they operate within a narrow scope. This means they are best suited to simpler use cases, such as answering questions within a pre-defined dataset, generating code, or producing recommendations based on a specific knowledge base.

4 Simple Questions to Tell the Difference Between an AI Assistant and Agentic AI

1. Can your AI proactively initiate actions without a human prompt, based on context or events?
2. Can you describe a situation where your AI made a decision or took action without human input?
3. Does your AI wait for instructions, or can it monitor data streams and trigger workflows autonomously?
4. What safeguards or approval processes exist when your AI acts autonomously?

Agentic AI takes autonomy further, allowing for applications that can resolve more complicated use cases.

For example, agentic AI can be used to investigate Tier 2 and Tier 3 security incidents, which has previously been the domain of "humans only." To do this, agentic AI must be able to:

- **Classify threats** using input from multiple sources (e.g., EDR, logs, telemetry).
- **Investigate autonomously**, determining what evidence is needed, and retrieving it without human input.
- **Refine its hypothesis** as it learns, removing false positives, adding new findings, and adapting its conclusions.
- **Prioritize results** into a structured plan that's ready for human approval or execution.
- **Continuously improve** by incorporating human feedback—both explicit (e.g., "approve," "dismiss") and implicit (e.g., what actions analysts take).

This kind of autonomy mirrors how skilled human analysts and threat hunters operate. Agentic AI doesn't just read alerts but thinks through them. It doesn't just summarize data; it makes sense of it, continuously learning, making decisions, and acting independently. We will take a look at how this works more closely later in this paper, but first, lets explore what makes agentic AI agentic.

What Makes Agentic AI Agentic?

An agentic AI system is based on four basic principles: reflection, planning, tool usage, and multi component collaboration. It uses sophisticated reasoning and iterative planning to autonomously solve difficult multi-step problems and execute tasks on its own–without human intervention or prompting, including:

- Using advanced Large Language Model (LLM) and machine learning (ML) algorithms.
- Executing actions in real-time.
- Achieving specific objectives without explicit coded business logic to do so.
- Building memories.
- Responding to dynamic scenarios.
- Leveraging sensory inputs.
- Learning and optimizing through continuous feedback to think, reason, and plan like a human.

A Practical Agentic AI Example: Ontinue's Autonomous Investigator Agent

To help illustrate the autonomy of an agentic AI system, let's take a detailed look at Ontinue's Autonomous Investigator agent. This agent, built from scratch as part of Ontinue's ION SecOps Platform, conducts Tier 2 and Tier 3 investigations of incidents escalated to Ontinue's Cyber Defense Center. It also produces detailed reports complete with recommendations and root cause analysis for Ontinue's security analysts to review. To execute such complicated tasks without being prompted, the Autonomous Investigator breaks the uber tasks into component parts, as noted below, just like a human analyst would. The agent replicates the work of a team of SOC analysts and generates its findings in minutes. The components and capabilities are shown in the diagram below.

In Ontinue's Cyber Defense Center, the autonomous investigator works in tandem with deterministic automation. In fact, 97% of incidents that Ontinue handles on behalf of customers were automatically resolved using deterministic automation in 2025. The Autonomous Investigator reviews these incidents retroactively to validate that they were closed appropriately, providing an additional check on the verdict and preemptively compiling a thorough report if a human wants to review the incident in detail.

For the 3% of incidents that cannot be handled by deterministic automation and are escalated for human review, the Autonomous Investigator proactively investigates and produces its report, giving Ontinue's security analysts a big head start. A typical investigation takes less than 10 minutes with the Autonomous Investigator, allowing analysts to spend more time per incident on deeper investigation and response. The result is that 99.5% of all incidents handled by Ontinue's Cyber Defense Center in 2025 were resolved without escalation to customer security teams. When an incident requires escalation to a customer, it is escalated fast, and it is accompanied with a thorough investigation report, complete with recommended actions that Ontinue security analysts have reviewed and vetted. This saves customer security teams time and effort they can repurpose for their own higher value work.

Ontinue

The Scalability Paradox in MDR: Why Ontinue Built the Autonomous Investigator

The industry has struggled with how to automate Tier 2 and Tier 3 investigations; most MDR providers still rely on human-only models for resolving these types of escalated and difficult cases. In practice, this means they pass complex incidents right back to internal security teams, shifting the burden back to the customer rather than solving the problem at the managed SOC level. This doesn't scale.

As threats grow faster, more evasive, and more frequent, human-only approaches can't keep up. The future of MDR demands more autonomy, more precision, and more context—without sacrificing trust.

Paradoxically, as threats and AI technology constantly evolve, human defenders must maintain full control of cases; human attacks require human defenses. At Ontinue, the humans-in-the-loop model also helps tailor final decisions to an organization's operational model and risk tolerance, monitor and adjust any natural hallucinations inherent in AI, and provide the final say on closing or directing cases to customers for further investigation or action.

Ontinue predicts that even with the most sophisticated, futuristic AI technology, human experts will always be in the SOC.

IDC recently recognized Ontinue's use of Agentic AI as a "leap forward in MDR capabilities," highlighting how it enables automation in incident investigation that was previously unattainable.

Agentic AI Innovation: Multi-Agentic AI is Next

With the creation of AI agents, the obvious next step is to expand the technology to multi-agentic AI. In the near future, Ontinue expects multi-agentic AI, a collection of AI agents that communicate and collaborate with each other, will be the next wave of innovation.

It is important to note that even with multi-agentic AI, there will still be humans in the SOC. Ontinue believes that multi-agentic AI and future technologies still will not replace humans, but rather continue to augment their intelligence to further prevent and defend ever-changing cyberattacks.

What is True Multi-Agentic AI

Multi-agentic AI enables communication and collaboration between multiple AI agents that have deep knowledge across very specific domains, such as an organization's attack surface, the cybercriminals and their motives (financial gain, espionage, geopolitical turmoil), attack techniques and exploits leveraged. These agents will work together to protect organizations against cyber threats at scale. For example, detection, response, and prevention will all be "single agents," and they will autonomously interoperate with each other, third party agents and customer agents across multiple environments and attack surfaces. Inter-agent communication and collaboration are difficult because they require a new infrastructure for the agents and new permissions between a SOC and the customer. Another area that is tricky for collaboration is getting the multiple agents to communicate and act with nuanced, contextual information, such as the industry of each customer, where certain employees are located and travel to, and job responsibilities.

Ontinue is continuously and quickly innovating within the field of agentic AI to "agentify" and seamlessly automate all aspects of MXDR for faster, more sophisticated and deeper threat prevention, detection and response. Every customer's environment is different; defenders need to be able to apply their tools and knowledge to specific conditions in unique environments to provide optimal security.

Innovating the Future of Security Operations with Multi-Agentic AI

Examples of future agents for specific types of security and defensive stages include:

Defender Agent for Detection and Response

This agent would enable broader coverage, speed up analysis and containment, and provide even more time for human SOC analysts to investigate cases. With a higher level of autonomy dedicated to detection and response, this agent also lowers the burden for customers because it contains threats faster and sends fewer alerts.

Cyber Advisory Agent for Prevention

This agent integrates customer-specific recommendations for security posture improvements and "contributed memory," such as meeting notes and other critical information, to drive automated risk mitigation tasks. Example tasks will include providing specific follow-up posture hardening actions after an incident, prioritizing existing security posture improvements tasks, creating new potential threat hunt scenarios, and developing dynamic risk labels for assets/environments. The agent will also flag and prioritize detected exposures at organizations and suggest how to optimize current detections.

Third-Party Customer Agents

This includes the ability to: streamline access to policies, settings, and Microsoft's new agentic platform; tap into customer memory, including past investigations, so that every incident becomes smarter; and tap into Sentinel Graph to get contents.

In the future, it's possible that third-party agents could implement auto-patching (through Microsoft Immune or Intune Agent) by triggering the patching agent within Intune. Not patching software is a perennial problem; in fact, 80% of attacks that come into Ontinue's Cyber Defense Center are from unpatched systems. Integrating this type of third-party customer agent could provide a massive shift in securing the world.

Imagine the possibilities. Ontinue already has.

Ontinue

Ontinue offers nonstop SecOps through an AI-powered managed extended detection and response (MXDR) service. Ontinue ION MXDR combines powerful proprietary AI with the industry's first collaboration with Microsoft Teams to continuously build a deep understanding of our customers' environments, informing how we prevent, detect, and respond to threats.

Continuous protection. AI-powered Nonstop SecOps. That's Ontinue.

CONTACT US

LEARN MORE