

White Paper

Agentic AI Security Innovation with AhnLab AI PLUS

What if AI could handle part of your job on the way to work this morning?


Some people would find that convenient, while others might feel anxious that AI could eventually replace their jobs.

What matters is that this is no longer a hypothetical scenario. AI is rapidly evolving beyond simply providing answers and is increasingly becoming capable of performing tasks alongside humans. At the center of this transformation is agentic AI.

This white paper explores the concept of agentic AI, examines response strategies against AI-driven attacks, and highlights benefits of adopting AhnLab AI PLUS, our agentic AI security platform.

What Is Agentic AI?

The evolution of AI can broadly be divided into three stages:



In simple terms, AI is evolving from “informing” to “performing.” The core characteristics of agentic AI include:

Autonomy

Makes decisions without constantly waiting for human instructions.

Goal-driven

Performs multi-step tasks to achieve objectives.

Tool use

Actively interacts with internal/external systems.

Memory

Retains previous results and uses them for future actions.

Agentic AI is therefore not simple automation; it is a goal-oriented autonomous execution.

How Is Agentic AI Different?

The biggest difference between generative AI and agentic AI is that the former is reactive, while the latter is action-oriented. Generative AI responds to questions, and humans remain at the center of the workflow. Agentic AI, on the other hand, receives a request, determines the necessary steps, uses appropriate tools, and produces results autonomously.

From a cybersecurity perspective, the transition is from “Summarize this log” to “Identify suspicious activity, gather relevant information, isolate affected systems, create a ticket, and generate a report.”

Security operations are shifting from being human-centered to AI-assisted and increasingly AI-driven.




Figure 1. Generative AI vs. Agentic AI

Recently, interest in agentic AI has surged because three key areas have matured:

LLM reasoning capabilities
Now supports complex multi-step problem solving.

Agent frameworks and protocols
Improves tool integration and execution.

API-centric enterprise
Allows security tools, ticketing systems, cloud services, and authentication platforms to be connected.

The technology and infrastructure are now ready, and market demand is growing rapidly.

According to Gartner, 60% of organizations are expected to provide personalized one-to-one customer interactions through AI agents. AI is no longer merely an assistant - it is becoming an operational component of organizations.

How Cybersecurity Roles Will Change

For cybersecurity teams, there are four major trends.

Agentic AI oversight
Invisible AI agents introduce new attack surfaces.

IAM for AI agents
AI agents must be managed as operational identities.

AI-driven SOC
AI will increasingly handle repetitive security analysis tasks.

Preemptive security
Security shifts from post-incident response to early prevention.

Historically, security admins primarily operated tools and manually managed them, responding to alerts. In the future, their role will increasingly become that of an orchestrator - designing AI agents, defining policies, approving exceptions, validating outcomes, and overseeing autonomous systems. In simple terms, AI is not replacing humans; it is changing their responsibilities.




Figure 2. Agentic AI security model

Let's explore expected changes in each cybersecurity domain.

#1. Endpoint Security

The potential use cases of AI agents in endpoint security are quite obvious.

Autonomous threat detection and isolation	When suspicious ransomware activity is detected, the AI agent can immediately isolate the affected endpoint, extract related Indicators of Compromise (IoCs), and deploy blocking policies across other endpoints.
Automated vulnerability prioritization	Instead of relying solely on CVE scores, the AI agent can assess actual risk based on factors such as exposure level, asset criticality, and the likelihood of exploitation, then dynamically adjust patching priorities.
Automated forensic data collection and reporting	The AI agent can automatically gather forensic artifacts, including memory dumps, process trees, and network connection histories, and then draft investigation reports.

The primary benefits are reduced Mean Time to Respond (MTTR) and more efficient use of threat analysis resources.

#2. Network Security

Key agentic AI use cases in network security include:

Autonomous segmentation and policy optimization	Continuously analyzes traffic patterns and optimizes security policies accordingly.
Dynamic zero trust policy enforcement	Enforces zero trust policies dynamically based on context, rather than relying on static rules.
Autonomous DDoS and C2 blocking	Automatically blocks threats that require rapid response, such as DDoS attacks or C2 (Command and Control) communications.
Autonomous remediation of misconfigurations	Automatically detects and corrects configuration errors, such as misconfigured security groups or publicly exposed storage in cloud environments.

After all, AI becomes a critical driver of policy automation and response speed.

#3. SOC

Modern security operations are increasingly evolving into multi-agent systems:

Detection agent	Identifies events and anomalies
Threat intelligence agent	Provides reputation and attacker context
Forensics agent	Collects evidence automatically
Response agent	Performs containment and blocking actions
Reporting agent	Creates tickets and reports

Recently, this style of AI-agent-based security operations has been formalized under the concept of an Agentic SOC. The key advantages of this model are speed and collaboration. It helps eliminate the bottlenecks that occur when a human analyst must perform every step of the process sequentially.




Figure 3. Agentic SOC Workflow

For many security professionals, AI-agent-driven security operations will likely be the first area where the impact of agentic AI becomes noticeable in day-to-day work. Traditionally, when an alert was generated, analysts had to manually perform reputation lookups, review historical records, examine similar incidents, assess the impact, and determine appropriate response actions.

However, when Agentic AI is integrated into security operations, it can analyze and consolidate the information required for decision-making and present it in an actionable format. As a result, security professionals can spend less time gathering and organizing information and focus more on making high-value decisions.

Risks Introduced by Agentic AI

Agentic AI provides tremendous opportunities for innovation to enterprises and security professionals. The challenge is that attackers have access to the same opportunities.

AI agents are highly attractive tools for attackers as well. They can dramatically increase the speed, scale, and sophistication of attacks. The ways in which attackers can leverage AI agents include the following:

AI-powered autonomous attack campaigns	Using AI agents to autonomously execute the entire attack chain, from vulnerability scanning and initial compromise to lateral movement and data exfiltration.
Advanced spear-phishing	Learning publicly available information about targets to generate highly customized messages at scale, while operating as a “conversational phishing agent” that responds to victims in real time.
Vulnerability weaponization	AI agents can immediately analyze publicly disclosed CVE information and generate exploit code, reducing the “golden time” between patch release and attack execution.

In addition to using AI to enhance attacks, organizations must also consider attacks directed at the AI agents they use. Related threats can be categorized into five areas:

Prompt injection	Injecting malicious prompts to manipulate the LLM’s original instructions and induce unintended behavior.
Agent hijacking	Stealing tokens or credentials to impersonate AI agents and gain control of systems.
Tool poisoning	Manipulating execution by compromising external tools or plugins used by AI agents.
Memory poisoning	Corrupting long-term memory sources such as databases to undermine decision-making logic.
Vulnerability weaponization	Granting agents more privileges than necessary, thereby amplifying the impact of a single compromise.

In summary, while it is important to successfully deploy and utilize AI agents, the primary challenge is to first establish a structure that controls the permissions and execution paths associated with AI.

Therefore, as organizations enter the Agentic AI era, they need to prepare for AI security and governance. This concept has recently gained significant attention in the cybersecurity industry.

Agent inventory	Ensure visibility into all AI agents operating within the organization (including SaaS) and clearly define their owners and managers.
Least privilege	Configure read-only access by default, restrict network segments accessible to agents, and separate permissions on a per-task basis.
Human-in-the-Loop policies	Define which actions can be automated and which require approval, and establish a “kill switch” capable of immediately stopping agents in the event of malfunction.
Machine identity	Treat agents as machine actors and operate dedicated credentials and audit logging systems for them.

Agentic AI security innovation with AhnLab AI PLUS

As attackers increasingly leverage AI, defenders' use of AI has become essential. Centered on our proprietary agentic AI security platform, AhnLab AI PLUS, we are driving customers to accelerate their transformation toward AI-driven security.

[AhnLab AI PLUS](#) is our AI-native security platform, strategy, and initiative that seamlessly unifies our technologies, products, and services around AI. With AhnLab AI PLUS, we are evolving from “AI-powered security,” which merely applies AI to existing products and services, to “AI-native security,” in which products, services, and security operations are designed around AI.

AhnLab AI PLUS implements an agentic AI architecture that integrates data distributed across security domains into a unified security context, enabling multiple AI agents to collaborate and enhance security operations. This establishes an AI-native security operations framework in which AI understands critical risks, prioritizes them, and helps people make faster and more accurate decisions.

The architecture of AhnLab AI PLUS and the role of each layer are as follows.




Figure 4. AhnLab AI PLUS

Layer	Role
AI Decision Intelligence	· Feed insights for risk prioritization and decision-making across security operations. · Understand which events matter most among the vast volume of security alerts.
AI-Native Operation	· Automate and enhance the entire security operations lifecycle—from detection, triage, and analysis to response—through AI-native operations. · Improve the speed and scalability of security operations.
Access/Exposure Management	· Minimize threat exposure through continuous exposure management (CTEM). · Continuously validate access and reduce potential risks.
Unified Security	· Seamlessly integrate AI with security platforms that bring together existing protection capabilities.
Security Data Fabric	· Collect and normalize data from native and third-party sources.

Our objective is to achieve a true “AI-native security,” in which products, data, and security operations are redesigned around AI from the ground up. AhnLab AI PLUS connects data and security capabilities across diverse security domains—including endpoint, network, cloud, email, and OT—into a unified security context.

Alerts and telemetry generated across individual products and services are utilized as normalized data. AI analyzes the relationships among individual alerts, the progression of an attack, and its potential impact, reconstructing signals from multiple products into a single incident and presenting response priorities. This allows security operators to move beyond reviewing alerts from each product individually and instead identify critical risks and response actions based on a unified incident and its broader context.




Figure 5. AI-native sensors triaging, investigating, and prioritizing threats

Figure 5 provides a clearer view of how this works. Over the past month, 20 AI-native sensors generated 1,969,908 alerts. Assuming two minutes of manual review per alert, analyzing them all would require 65,664 hours—an amount of time that is physically impossible to process manually. Our AI automatically analyzes approximately 2 million alerts, triages 16 as requiring actual investigation, and identifies 13 as real threats. The entire process of identifying these threats takes only 4 minutes and 23 seconds. When a user selects an AI-investigate threat, they can review the detailed analysis and the response recommendations.

Also, we monitor and validate AI decisions, actions, and data access, while requiring user approval for key actions before execution. This reflects the principle of “controlled autonomy”, in which AI is granted autonomy while remaining governed through permissions, policies, and auditing processes. In this way, we have designed an AI-native security operations framework that improves the speed and efficiency of security operations while maintaining safety and trust.

In addition to agentic AI security, we also deliver AI security and governance capabilities. SecureBridge, provided by our subsidiary AhnLab CloudMate, detects and blocks confidential and critical information contained in prompts to prevent such data from being exploited, while allowing safe prompts to be transmitted as they are. This enables organizations to use AI without concerns about data leakage.

Benefits of AhnLab AI PLUS

From the user's perspective, the tangible benefits of AhnLab AI PLUS are:

	See What Matters	Identify and prioritize the most critical threats and risks among countless alerts.
	Understand Faster	Quickly understand why a threat is significant and which assets are affected.
	Act with Confidence	Receive guidance on the investigation and move quickly from analysis to response.

By leveraging AhnLab AI PLUS, we are achieving measurable improvements across security operations, including event processing, threat analysis, and monitoring.



AhnLab AI PLUS is not about presenting customers with more alerts or information. Instead, it is to reduce the time required to identify what matters, understand why it matters, and determine what actions need to be taken. Through this process, organizations can transform their security operations with intelligent AI.

FAQ

#1. What is AhnLab AI PLUS?

AhnLab AI PLUS is AhnLab's AI-native security platform and strategy that integrates our technologies, data, products, and services around AI to make security operations more intelligent and autonomous.

#2. Why AhnLab AI PLUS?

Because emerging AI-driven threats must be addressed through unified security operations rather than individual products. AhnLab AI PLUS delivers distinctive AI security capabilities by learning from the security data and intelligence that we have accumulated for more than 30 years.

#3. What changes for users?

AI enables users to operate security faster and more efficiently by quickly understanding critical risks and translating those insights directly to investigation and response. This helps minimize human error caused by manual tasks, improve productivity, and allow users to focus on critical decisions and higher-value work.

#4. How does it transform security platforms?

We already provide AI-driven event analysis capabilities, such as "AI Insight," in solutions including [EDR](#) and [XDR](#), and are about to deploy AI-native triage, investigation, and response capabilities. It is expected to automatically reduce the volume of tens of thousands of alerts generated each day and use AI-driven investigation and risk scoring to distinguish true positives from false positives, so analysts can focus only on actual threats. In doing so, it can significantly reduce manual workload and analyst fatigue.

Conclusion

In a single sentence, agentic AI can be defined as: “an autonomous AI system that plans, takes actions using tools, and achieves goals on its own.” For security professionals, it represents both a powerful automation opportunity and a new attack surface. While it is important to use AI agents effectively, organizations must also establish mechanisms to control and validate their actions.

Agentic AI is no longer a distant future technology. It is more of a present reality that is already spreading. The first question readers should consider after reading this article is:

“How will our organization use AI agents, and are we prepared to use them effectively and securely?”

If security professionals truly understand the context of this question, they will keep staying successful in the rapidly changing AI era.



AhnLab

AhnLab, Inc.

220, Pangyoyeok-ro, Bundang-gu, Seongnam-si, Gyeonggi-do, 13493, South Korea

www.ahnlab.com / global.sales@ahnlab.com

© 2026 AhnLab, Inc. All rights reserved.