The future of security operations is powered by AI agents.

AI agents are orchestrating the way for a new, streamlined approach to security operations.

Security Operations Centers (SOCs) are overrun.

Genuine threats are obscured by the constant stream of data, alerts, and a high false positives rate.

83%
of alerts are false positives, which can needlessly drain resources.[1]

Analyst burnout and turnover results from overwhelming staff with repetitive, manual tasks.

67%
of alerts go unaddressed due to the sheer volume of results.[1]

Wide attack windows arise due to labor-intensive tasks that stall mean time to resolution.

71%
of analysts report not being able to investigate all the daily alerts they receive.[1]

Security Orchestration Automation and Remediation (SOAR) enabled some automation, but provided only incremental benefits.

Can AI agents help?

With their ability to reason, act, observe, and adjust actions based on new information, AI agents have the potential to help security teams **identify** and **respond** to threats **faster**.

In the near term, AI agents will take over manual tasks like alert triage and investigation.

For example, the **Google SecOps Alert Triage and Investigation Agent** can enrich alerts, analyze command line interfaces, build process trees, and determine an alert verdict with next-step recommendations for humans.

But they don't stop there...

The agentic SOC of the future:

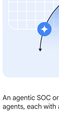

An agentic SOC orchestrates a system of task-based AI agents, each with a specific role, to achieve a common outcome.

Upon receiving a security alert, the agentic SOC cycles through a process of:

This dynamic process of evaluating, acting, and re-evaluating enables the system to **autonomously adapt to a changing security environment in real time**.

Google already has the tools

Technologies like Agent2Agent (A2A) and Model Context Protocol (MCP) provide a preview of how **multiple SOC agents can share the same security data sources, regularly communicate, and adapt their actions.**

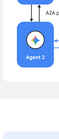

A2A
Provides AI agents with a common communication channel to share insights and coordinate actions.

MCP
Enables the AI agent to access data, tools, and other resources needed to inform the defined action.

Delivering better security outcomes through automation.

The automation of repetitive and manual tasks not only enhances analyst job satisfaction, but improves threat prioritization and resolution time.

Human-led and AI-driven security operations.

While some SOC workflows will transfer to AI agents, human analysts will continue to run:

Advanced threat hunting: For the truly novel, sneaky attacks that only a human can intuitively grasp.

Strategic risk and business context: Translating tech risks into real-world business impact.

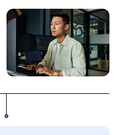

Crisis communication and management: Handling the messy, human side of a major incident.

Security architecture and AI coaching: Designing the overall defense and fine-tuning the AI itself.

The future isn't a human-less SOC, but a human-led, AI-powered one.

The trust-but-verify model will be paramount because of the unwritten knowledge and insight into real world impact humans bring.

Learn more about how Google is helping organizations supercharge security with AI.

Visit our webpage

[1] 2025 Work+ AI: State of Threat Detection, July 2025

Google Cloud