Google Cloud
Security

Reinventing the SOC with agentic AI

Reducing risk, increasing speed, and upleveling your entire security team

Security teams are overwhelmed by rising threats, limited staff, and manual workflows that can't keep pace with threats. This guide shows how AI agents can help organizations reduce exposure, scale their defenses, and give every analyst the power to perform at their best.

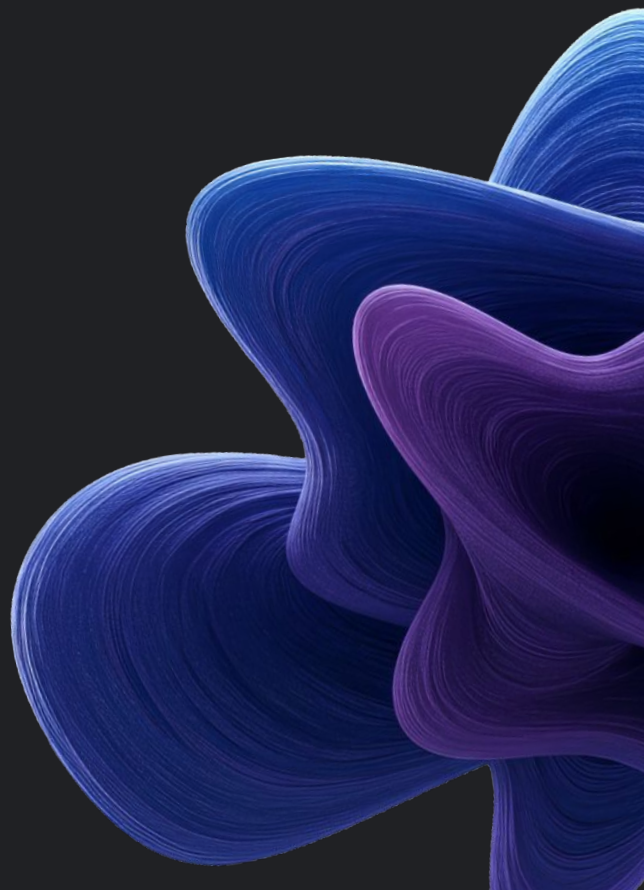

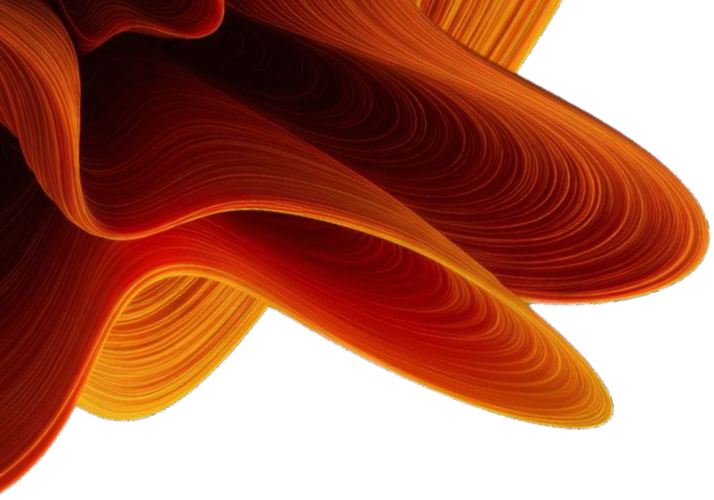

The SOC stretched thin

Security operations have reached a breaking point. Teams are overwhelmed by the scale of today's threats, buried under tool sprawl, and slowed by disconnected data. Attackers now use AI to generate endless variations of phishing, automate reconnaissance, and exploit vulnerabilities at machine speed — while security teams are still stuck reacting with manual workflows. Add a global cybersecurity talent shortage, and it's clear why defenders feel outpaced.

Google's vision: AI that empowers defenders — not replaces them

AI promises a path forward, but true SOC transformation will only happen where AI becomes a force multiplier for the analysts who are already stretched thin.

Google's vision embraces this shift: agentic AI that empowers defenders, not replaces them. AI agents act as trusted teammates that automate workflows and tackle routine tasks, ultimately accelerating investigations and multiplying human expertise.

91% of organizations have launched initiatives, yet

74% struggle to operationalize them.

Sources:
Globalization Partners / Boston Consulting Group

> *We are at an inflection point... choosing a partner who can secure this AI transformation is the most important decision you can make."*
>
> — John Ramsey, Google Cloud Security

Key takeaway: AI-enhanced defense represents the next stage of SOC evolution — where AI agents act as trusted teammates to scale expertise and speed.

Watch the clip:

How Google Cloud is redefining the SOC for the AI era

Google Cloud Security

Why the agentic SOC changes everything

The agentic SOC goes far beyond scripted automation. These systems can reason, plan, and take action — all with human oversight. Instead of following rigid playbooks, agents assess context, connect patterns, and recommend next steps, much like a seasoned analyst would. This allows complex investigative work to happen in seconds, not hours.

Google's Alert Investigation and Triage Agent

A powerful example is the Google Cloud Alert Triage and Investigation Agent, which helps security practitioners quickly and effectively identify threats by performing initial triage, saving valuable time. Designed to autonomously investigate alerts and provide comprehensive explanations, the agent autonomously gathers evidence, runs analyses, and delivers verdicts — which means numerous alerts that might otherwise go untriaged can now be processed by agents – enabling security analysts to prioritize alerts which require human attention.

> *We're ushering in the era of the agentic security operations center. Imagine a SOC where AI agents are specialized to handle specific tasks with unparalleled accuracy and quality."*
>
> —John Ramsey, Google Cloud Security

SOCs using the Google Cloud Alert Triage and Investigation agent report:

65% faster response times

70% lower breach risk

Source: Forrester TEI Study

Key takeaway: Analysts stay in control, but their workload shifts dramatically from manual tasks to high-value decision-making.

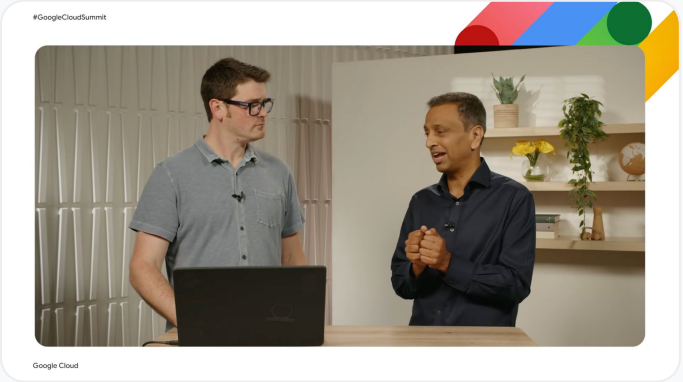

Watch the full session:

Empower your SOC with agentic AI for autonomous outcomes in Google SecOps

Inside the agentic SOC

The Google Cloud agentic SOC is built on specialized AI agents powered by Gemini and trained on Google-scale telemetry and frontline Mandiant intelligence. These agents take on the tasks that typically drain analyst capacity — and they do it with more context, more consistency, and far more speed.

Agents can automatically triage alerts, correlate signals across multiple tools, and summarize investigations into clear, auditable decisions. Others specialize in malware analysis, detection engineering, and continuous threat hunting — identifying patterns and generating response playbooks in real time. Every action is logged, explainable, and reviewable, ensuring transparency and analyst control.

Core capabilities enabled by agents

- Autonomous triage and enrichment
- Automatic malware classification and rule tuning
- Continuous hunting and playbook creation
- Fully explainable decision paths

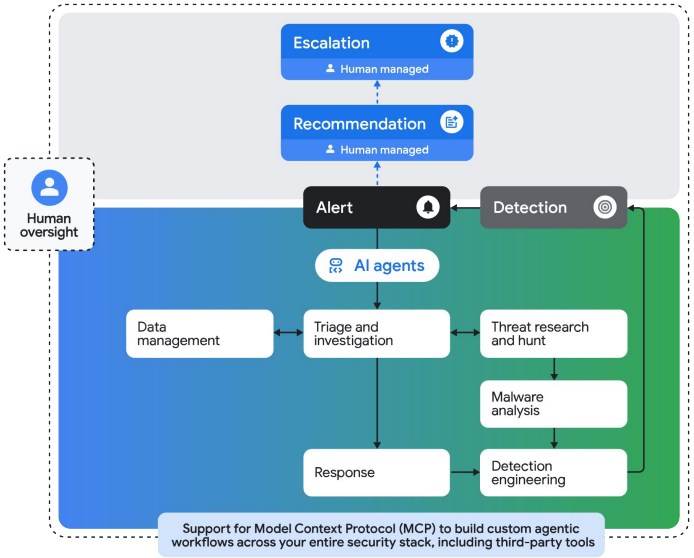

Focusing human analysts on critical decision-making

This agentic approach goes beyond accelerating work to changing the nature of how the SOC operates. Instead of chasing alerts, analysts review final outputs, validate reasoning, and guide the system.

Furthermore, support for Model Context Protocol (MCP) lets these agents connect with tools across an organization's environment, so SOC teams can orchestrate multi-vendor workflows using natural language.

> *The AI agent looks for who, why, what, how, and when for the alert — running searches, finding past and similar alerts, collecting and analyzing data, and dynamically reasoning through the problem. Today, this would take hours to resolve."*
>
> — Payal Chakravarti, Google SecOps

Key takeaway: With AI agents supporting continuous monitoring and surfacing context-rich, prioritized alerts, analysts make faster, more confident decisions to move from reactive to proactive defense.

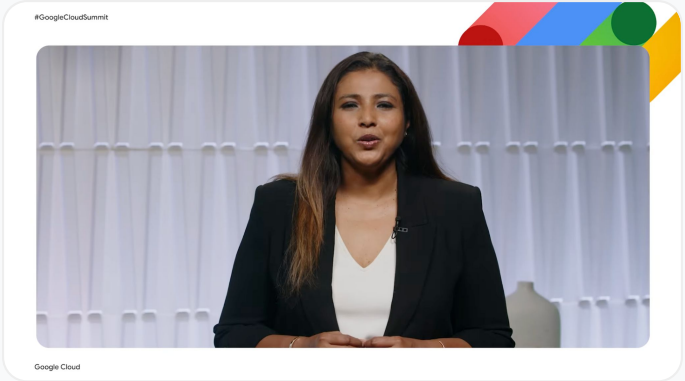

Watch the clip:

Inside the agentic SOC — autonomous workflows in action

The new threat frontier: When attackers use AI, too

AI is changing how defenders work — but also reshaping how attackers operate. Threat actors now use generative AI to craft highly convincing phishing messages, automate reconnaissance, and probe for model vulnerabilities. These techniques allow attackers to move faster, operate at higher volume, and evade traditional detection.

In an insidious twist, attackers are using AI-enhanced tactics to target the unique vulnerabilities of AI applications. Threats like prompt injection, data poisoning, and synthetic identity generation introduce new risks that didn't exist even a few years ago. Organizations need visibility into how their models behave, how they're accessed, and how attackers might try to manipulate them.

A holistic approach to defending AI from the inside out

Google's approach brings together real-time signals from Google, Mandiant, and VirusTotal to detect emerging AI-driven threats before they spread. Technologies like Model Armor apply guardrails at runtime, blocking malicious prompts and preventing data leakage, while Safe Browsing detects malicious content at the browser edge.

Common AI-driven threats

- Prompt injection
- Training data poisoning
- Synthetic identities
- LLM probing and multi-step manipulation

Key takeaway: When AI and agents act as trusted teammates, defenders gain the advantage — and the SOC becomes stronger, faster, and more resilient.

> *Generative AI allows threat actors to move faster and at higher volume."*
>
> —Aurora Bloom, Google Threat Intelligence Group

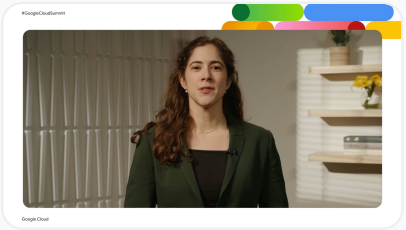

Watch the full session:

Hype vs. Reality — What to know about adversarial misuse of AI

Customer outcomes: What agentic defense delivers

Organizations adopting Google SecOps with AI agents are seeing meaningful, measurable improvements. With agents handling triage, correlation, and repetitive investigations, teams can focus on strategic work — and results come quickly. Investigations close faster. Fatigue drops. Accuracy improves because decisions are made with richer context.

Key customer outcomes

- 50% faster mean time to respond
  Source: Forrester The Total Economic Impact™ Of Google SecOps, July 2025
- Lower turnover and analyst fatigue
- Higher detection accuracy

Apex Fintech Solutions

Apex Fintech Solutions saw the impact immediately. Long investigation cycles were reduced to 15–30 minutes, thanks to autonomous triage and enriched context. Productivity increased, and analysts gained a clearer picture of risk across the environment.

Watch the customer story

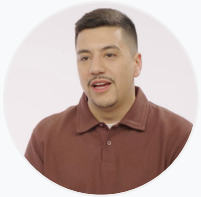

"Google Security Operations is the brain of our SOC."
— Hector Peña, Sr. Director, Information Security, Apex Fintech Solutions

Lloyds Banking Group

At Lloyds Banking Group, the UK's largest digital bank, security teams needed to move faster while managing risk at a national scale. By replatforming security operations on Google SecOps, Lloyds unified data from across a complex, multi-vendor environment, reducing low-value alert noise and allowing analysts to focus on high-fidelity threats.

Watch the customer story

"Google SecOps is right at the heart of our security transformation, and we chose it because of the engineering-led approach. It's also a really open ecosystem."
— Matt Rowe, Chief Security Officer, Lloyds Banking Group

Key takeaway: Agentic defense allows organizations to achieve the dual goal of operational efficiency and stronger protection.

Resources for CISOs

- Cyber Crisis Planning & Response Services
- Mandiant AI Security Consulting Services

Watch the clip:

Incident response planning in the age of AI
Tabletop exercises: Simulating end-to-end incident scenarios

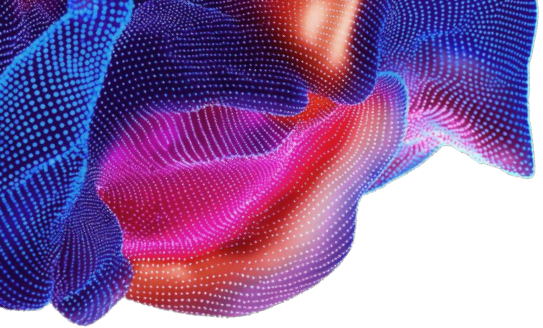

Moving your SOC forward

The challenges in front of the SOC can't be solved by adding more point solutions — the bloated security stack is already a major problem. Modernizing the SOC requires leveraging new tools and technologies to fundamentally redefine how security teams work. AI and agents help organizations scale expertise, improve outcomes, and protect data and systems with greater confidence. The path starts by understanding where your SOC is today and what autonomous defense can unlock.

> *The possibilities, as you can see, are limitless... Agentic AI is here to drive that step function improvement and efficiency that the world of security has been waiting for."*
>
> — Payal Chakravarti, Google SecOps

Ready to get started?

Google provides a set of resources to help teams assess maturity, measure impact, and identify high-value opportunities for AI-driven operations.

See where your SOC stands	Estimate your ROI	Enhance your threat hunting
Take the Google SecOps Maturity Assessment.	Read the Forrester TEI report: The Value of SecOps.	Learn more about Mandiant services for AI and threat detection.

Read the companion guide
Secure AI innovation without interruption

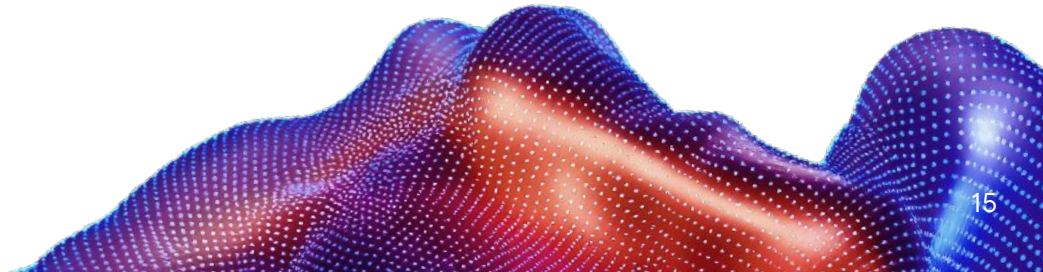

Google Cloud

Security