KPMG | Google Cloud

AI-driven capabilities for security operations

Overview

The traditional security operations model, often strained by manual processes and fragmented tools, is proving unsustainable under the weight of modern demands. Forward-thinking organizations are turning to innovative solutions that integrate advanced technology with strategic expertise. Google's cutting-edge security tools, complemented by the industry knowledge of KPMG, offer a detailed approach that enhances visibility and delivers artificial intelligence (AI)-driven insights. This collaboration empowers security teams, enabling them to detect, investigate, and respond to threats with unparalleled speed and precision, redefining proactive security in an era of escalating cyber risks.

Introduction

Security operations are at a breaking point. The traditional model—built on siloed tools, manual processes, and reactive workflows—can no longer keep pace with the scale, speed, and sophistication of modern threats. This paper explores a new paradigm: an AI-driven, agentic security architecture that unifies detection, response, and intelligence into a cohesive, autonomous system. By focusing on the tools and technologies that enable this transformation, we believe organizations can shift from fragmented defenses to a proactive, intelligent security posture that's human-led and AI-powered.

Capabilities for modern defense

To escape the reactive loop that has defined security operations for decades, a fundamental shift in thinking is required. A modern security operations center (SOC) needs to be built around three pillars:

1. **Resilient:** Designed to dynamically adapt and scale to cyberattacks and data influxes
2. **Proactive:** Continuously monitoring and mapping potential attack paths to prevent threats before they can impact organizations
3. **Intelligent:** Leverages AI to analyze vast amounts of data for actionable insights and to automate workflows.

These capabilities form the foundation of an agentic SOC—one that is proactive, intelligent, and resilient by design. It's not about working faster; it's about building a system that works smarter, freeing human experts to focus on the challenges that truly require their intuition and skill.

Built to scale with context

To build a modern defense, first SOCs must address the foundational challenge of data. SOCs are drowning under the weight of their own telemetry, struggling to ingest at scale and, more importantly, make sense of it. Organizations need a SIEM built on hyperscale infrastructure with intelligence, adding valuable context, woven into the very fabric of their security operations.

This need for context is especially critical as the attack surface expands into new, complex domains.

A proactive security program requires a unified view that can identify asset relationships enriched with domain-specific intelligence. For example, take an AI chatbot.

If an end user is attempting to trick your chatbot into sharing sensitive data, then you want to detect the user's prompt injection attempts before something like data leakage occurs. Before you can block the user, you need to identify prompt injection attempts and the relationships between the model, application environment, and user endpoint.

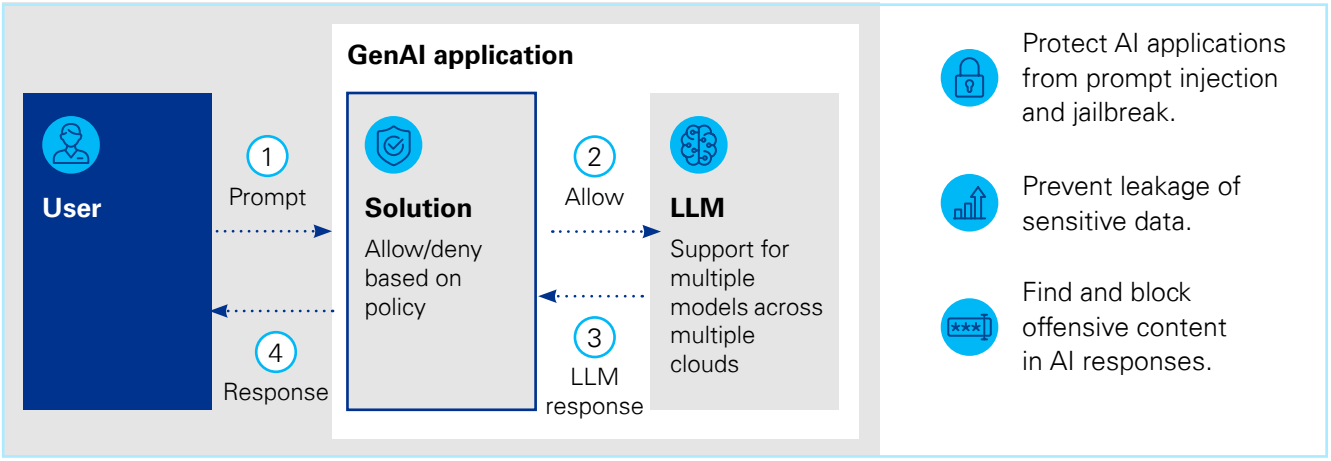

How to block screen malicious prompts and responses to and from a GenAI application

Browser is the new endpoint

This expanding perimeter extends all the way to the end user. In a world of hybrid work and Software-as-a-Service applications, most work is done in the browser. Much of an organization's sensitive data is now accessed and handled on the web, often from outside the traditional corporate network, creating significant risks of data exfiltration. This requires a new approach to endpoint security centered on an enterprise-ready browser. Organizations need the ability to enforce context-aware access and data protection policies to prevent both accidental and intentional data loss, no matter where their users are working.

The agentic SOC in action

An agentic SOC moves beyond the static, linear logic of yesterday's SOAR playbooks and instead pursues high-level goals by reasoning across a rich, precorrelated data set. For it to function, the AI agents require a scalable platform with a constant stream of shared, real-time context based on all security telemetry—from the cloud to the browser—residing in a single, unified security data fabric.

Google SecOps ingests and normalizes security data at hyperscale into a universal data fabric with threat intelligence built into the platform. Insights from Google Threat Intelligence, powered by Mandiant's frontline investigations, VirusTotal, and Chrome Safe Browsing, are not just bolted on; they are woven into the very fabric of the data. This level of event enrichment by default enables AI tools, like Gemini, to thrive. This means that when given the task "investigate this suspicious file download," Gemini can see the file's reputation, trace its origin to a phishing email, follow its subsequent execution on the host, and identify malicious command-and-control traffic—all within a single, unified platform.

This seamless context extends from the cloud all the way to the user. Telemetry from Chrome Enterprise Premium flows directly into the same security data fabric that feeds into Google SecOps. This allows an AI agent—or a human analyst—to follow the narrative of an attack from a malicious click on a web page to the resulting process activity on the host all through a single platform. That same principle applies to the cloud infrastructure itself. With Security Command Center, defenders gain a unified view of risk across their multicloud estate, from misconfigurations and vulnerabilities to threats against their AI workloads. These findings aren't isolated; they feed directly into the same data fabric, giving Google SecOps and Gemini the full picture of an organization's posture.

Google offers a resilient, proactive and intelligent defense powered by a system of AI agents working together to accomplish a shared goal. For example, a Gemini-powered AI agent could use Google Threat Intelligence to simulate an emerging attack, see if existing rules in Google SecOps fire, and if not, automatically generate a new, validated detection. At the same time, another AI agent could continuously scan the data fabric for the subtle, low-and-slow signals that precede a major breach.

This culminates into an agentic SOC drawing from an ecosystem of specialized agents, all acting on the same unified data fabric. When the agents reach the limit of their capabilities, they can escalate to human experts. AI agents and human expertise all working in concert can transform security from a reactive to preemptive defense.

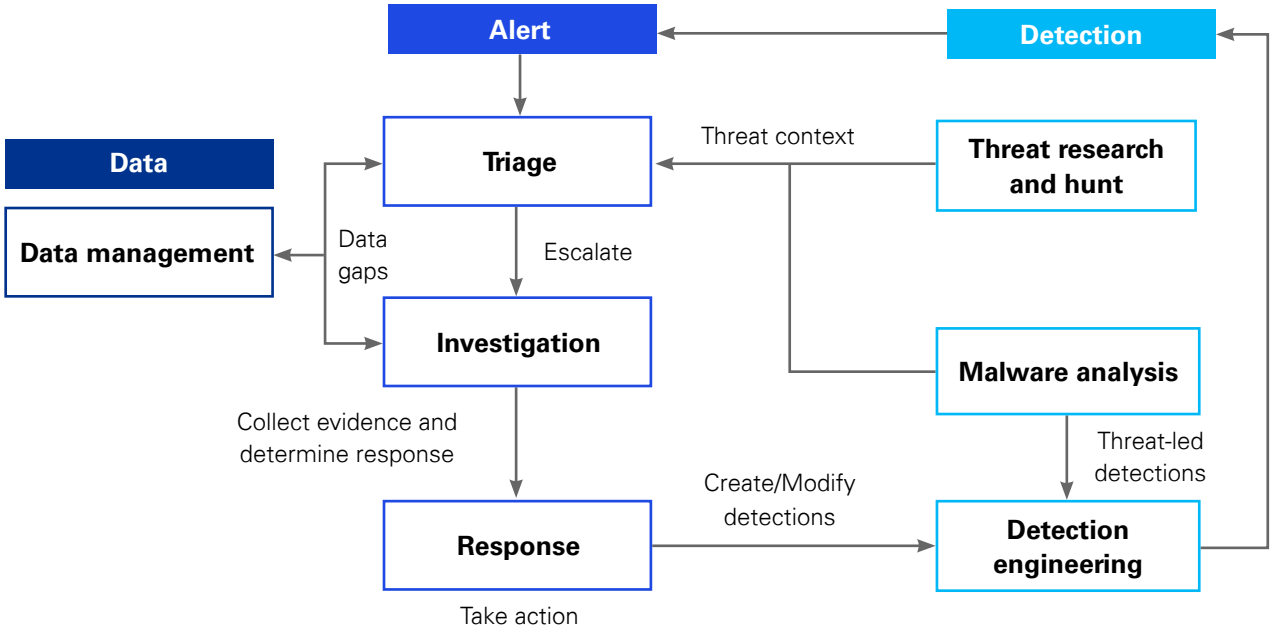

How agents can enhance SOCs with Google Security Operations.

Example scenarios and use cases

The true power of an agentic SOC is its ability to transform a single, low-context alert into a fully investigated, actionable incident in minutes. This is achieved not by a single technology, but through a collaborative system where AI agents and human analysts work together as a unified team.

Consider a common scenario: an employee downloads a seemingly benign file, which generates a security alert.

The AI teammate's initial triage

Leveraging Google SecOps' native agents as the first line of defense, a Malware Analysis Agent automatically examines the file's code and behavior, quickly identifying it as malicious. This verdict is immediately passed to an Alert Investigation Agent, which investigates the surrounding activity, correlates it with threat intelligence, and confirms the threat's urgency. Within minutes, the system has autonomously handled the initial, high-volume analysis and presented a validated, high-priority incident to its human teammate.

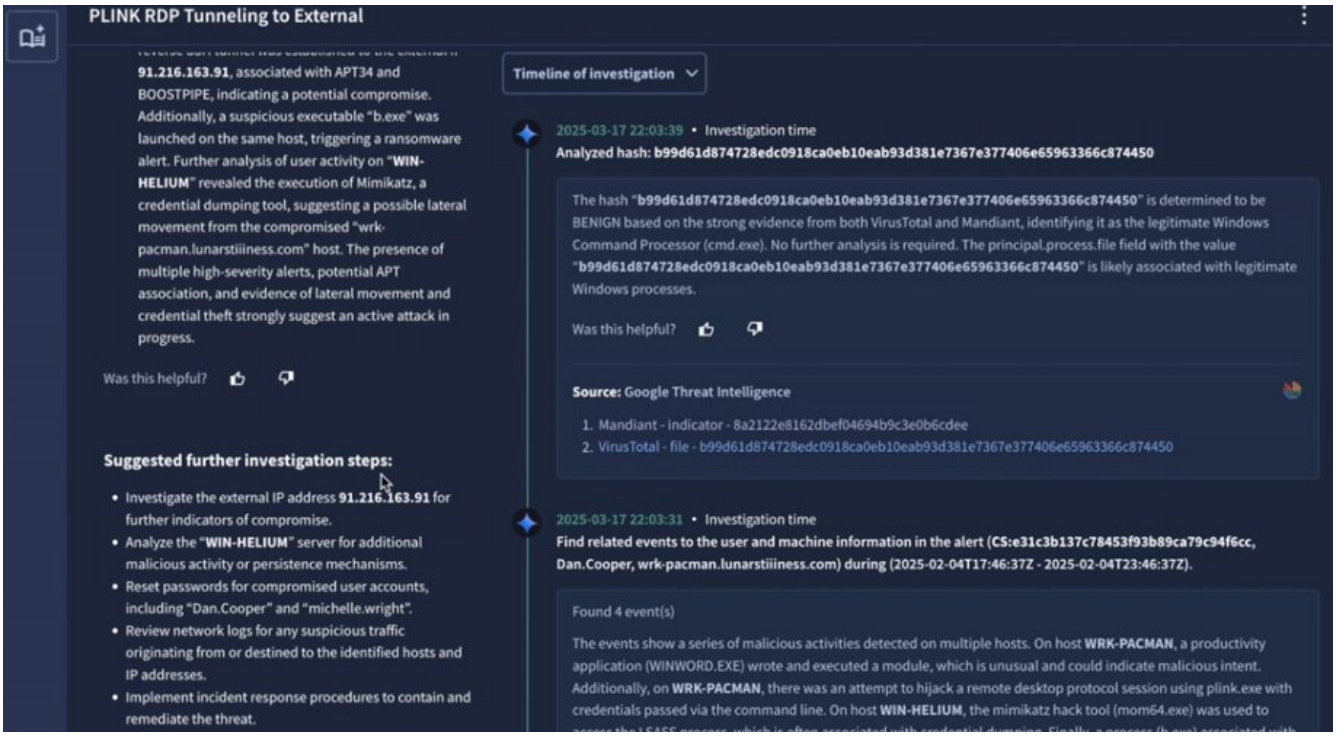

SecOps' Alert Investigation Agent suggesting next steps after completing an investigation.

The analyst

Directed investigation

With the initial triage complete, the human analyst steps in as the strategic commander. Freed from the time-consuming tasks of the initial investigation, the analyst leverages the built-in Gemini interface, using simple natural language prompts to direct the system to perform a broad search for any related activity across the network. This AI-assisted query allows the analyst to effectively determine the full scope of the incident and, based on the findings, make the crucial decision to initiate the appropriate SOAR playbooks to contain the threat.

The advanced partnership

Some may prefer custom development with unique risk profiles, governance models, and operational complexity. These organizations can leverage Google's broader technology suite to develop custom AI agents that act as powerful assistants for the human analysts. By using the open source SecOps MCP servers, with Google technologies like Agent2Agent, AgentSpace, and Vertex AI, organizations can build custom agents to handle industry-specific requirements, such as specialized threat detection for healthcare device protocols or automated analysis of financial transaction anomalies. These agents can be tailored to organization-specific needs, like automating proprietary incident response procedures or enforcing company-specific security policies.

The traditional process A linear, analyst-driven investigation	The agentic advantage A collaborative, AI-driven defense
Manual alert triage An analyst manually reviews, enriches, and validates an incoming security alert from a large queue.	**The AI teammate's triage** Native AI agents autonomously analyze, enrich, and validate the threat, presenting a high-priority incident.
Manual investigation The analyst performs broad searches and pivots between tools to determine the incident's scope.	**The analyst-directed investigation** The human analyst acts as a strategic commander, using AI-assisted queries to determine the incident's full scope.
Manual response Based on findings, the analyst manually initiates SOAR playbooks to contain the threat.	**The advanced partnership** The analyst directs a team of native and custom AI agents to perform complex tasks and automate the response.
Estimated time elapsed: **Hours...** **Outcome: Delayed, reactive response**	Estimated time elapsed: **Minutes...** **Outcome: Swift, collaborative response**

The path forward: Cliffhanger

The convergence of hyperscale data platforms, embedded threat intelligence, and agentic AI represents an inflection point for security operations. As we have explored, these capabilities provide the technological foundation to move beyond the limitations of traditional SOCs, enabling organizations to detect and respond to threats with unprecedented speed, precision, and autonomy. However, possessing these advanced tools is only the first step. The true challenge—and the greatest opportunity—lies not in technology itself, but in its successful integration into the fabric of an organization. Successfully deploying these solutions at scale requires more than just technical expertise; it demands a holistic, transformational approach that realigns people, processes, and governance to create a new operational reality.

Our final white paper in this series, Security transformation journey, provides the strategic blueprint for this evolution. We will move from the "what" to the "how," exploring the critical steps to implement an agentic security model at scale. We will detail how to design a new target operating model for a modern cyber defense operation, navigate a phased implementation roadmap from initial assessment to full optimization, and address the crucial human elements of change management and skills enablement.

The future of defense is not just smarter—it's unified, autonomous, and already within reach. Join us as we lay out the path to achieving it.

Contact Us

Steve Barlock
Principal, Advisory
E: sbarlock@kpmg.com

Anton Chuvakin
Security Advisor at Office of the CISO, Google Cloud
E: chuvakin@google.com

Niranjan Girme
Director, Advisory
E: ngirme@kpmg.com

Ash Elahi
Manager, Advisory
E: ashelahi@kpmg.com

Justin Horbacz
Sr Associate, Advisory
E: justinhorbacz@kpmg.com

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us: **kpmg.com**