KPMG | Google Cloud

# Security transformation journey

## Overview

Building on the strategic foundation laid in earlier papers, this white paper presents a practical implementation guide ready to embrace an artificial intelligence (AI)-powered agentic security operations center (SOC). It presents a three-phased approach to transformation, supported by the experience of KPMG as a Google Cloud Security Premier Partner. The focus is on delivering immediate value while building toward autonomous, intelligence-led security operations.

## Introduction

In our previous white papers, we established the need for a new security paradigm and introduced the AI-driven capabilities that make it possible. This paper shifts focus—from the "why" and "what" to the "how."

Transitioning from a reactive, tool-centric SOC to an intelligence-led Cyber Defense Operation is not just a technological upgrade—it's a full-scale transformation. It demands more than new tools; it requires experienced leadership, strong governance, and a clear direction.

This white paper presents a structured approach for that journey. It guides organizations from traditional, human-limited operations to an agentic SOC—where automation, context, and decision-making are deeply integrated. A Google Cloud Security Premier Partner like KPMG brings the strategic insight and change management experience needed to lead this shift. Patchwork fixes are no longer sufficient. The path forward requires intelligent, adaptive, and agentic operations.

## Agentic AI SOC implementation roadmap

Given every organization has unique processes, risk tolerances, and maturity levels, this transformation strategy is not intended to be a one-size-fits-all solution. It's a three-phased flexible approach to guide teams gradually toward agentic AI capabilities.

### Phase 1: Assessment and agent readiness

**Goal:** The objective of this foundational phase is to move beyond generic automation strategies and develop a tailored approach for AI adoption. It begins with a thorough analysis of the current SOC environment to identify high-impact opportunities and prepare systems, data, and processes to support intelligent automation.

**Actions**

- **Thorough SOC assessment**
  The journey starts with a current-state assessment of the SOC environment. This includes mapping existing workflows, identifying bottlenecks, and isolating repetitive tasks that consume analyst time. Key metrics—such as alert volume, false positive rates, and mean time to resolution—are captured to prioritize automation opportunities and establish a baseline for measuring impact.
- **Identify and prioritize the automation backlog**
  Effective automation starts with knowing what to automate and how. This step involves reviewing and categorizing the existing backlog of SOC workflows to identify where automation can deliver the most value. Some tasks are ready for immediate implementation using tools like Google SecOps SOAR playbooks. Others may need to be re-engineered to support AI-driven reasoning—for example, a phishing investigation workflow that currently relies on manual correlation of email headers, user behavior, and threat intelligence might need to be redesigned so an AI agent can autonomously analyze indicators, trace activity across systems, and determine containment actions. The goal is to distinguish between rule-based processes that benefit from straightforward automation and complex investigations that can benefit from agentic intelligence.
- **Establish technical readiness for AI**
  Before AI agents can deliver value, the environment must be ready to support them. The goal here is to establish the technical foundation required for agentic operations.
  - **Unified data and application programming interfaces (APIs):** AI agents depend on rich, contextual data to make informed decisions. That requires breaking down silos while helping ensure critical systems—like CMDBs, HR platforms, and identity providers—are accessible via APIs. Without this integration, agents can't distinguish between routine activity and potential threats—such as differentiating a scheduled data migration from a data exfiltration attempt.
  - **Standardized operating procedures (SOPs):** AI agents benefit from clear, structured guidance to operate effectively, especially in environments where consistency and accountability are critical. Documenting and standardizing investigation and response procedures provides a framework that helps agents act reliably and transparently. While some believe AI thrives in unstructured scenarios, SOPs serve as foundational guardrails that support intelligent decision-making rather than constrain it. These SOPs define how agents respond to specific incident types and ensure automation aligns with organizational expectations and risk tolerance.
- **Establish robust governance and guardrails:** AI adoption demands trust. This step introduces a governance framework to help ensure AI agents operate safely, predictably, and within defined boundaries:
  - **Define trust boundaries:** Set clear policies that determine when agents can act independently and when they must escalate to humans.
  - **Implement a secure AI framework:** Adopt a structured approach like Google's SAIF to embed security, accountability, and responsible AI practices into every stage of development and deployment.

**Outcome**

By the end of this phase, the organization can have a data-driven strategy and a prioritized roadmap for agentic transformation. The technical and procedural groundwork can be in place, creating an environment where AI can be integrated effectively and responsibly.

## Phase 2: Augmentation with built-in AI agents

**Goal:** This phase delivers immediate value by embedding Google's built-in AI agents into investigation workflows. The objective is to automate high-effort tasks, reduce manual load, and enable analysts to respond faster and more accurately.

**Actions:**

- **Deploy built-in agentic capabilities:** Google SecOps provides powerful, prebuilt AI agents designed to tackle the most common and time-consuming SOC tasks. This phase involves operationalizing these agents:
  - **The Alert Triage Agent:** This agent handles the initial investigation for incoming alerts—gathering context, analyzing evidence, and delivering a verdict with a transparent audit trail. Tasks are now completed in a fraction of the time it would have taken an analyst to complete, freeing up Tier 1 analyst capacity and increasing investigative rigor through consistent, repeatable analysis.
  - **The Malware Analysis Agent:** Integrated with Google Threat Intelligence, this agent autonomously reverse-engineers suspicious files—analyzing obfuscated code to deliver a clear verdict on whether a file is malicious. It replaces a time-consuming task that previously required highly specialized human expertise.
- **Empower analysts with Gemini:** Gemini in Google SecOps brings AI directly into the analyst's workflow. Analysts can use natural language to craft complex UDM queries, accelerate investigations, generate AI-powered case summaries, and build detection rules from plain text descriptions.
- **Operationalize a human-in-the-loop model:** AI agents handle the investigation, but human analysts validate findings and approve response actions. This model builds trust in AI decisions while serving as real-time training—analysts learn from AI reasoning while AI improves through human feedback.

**Outcome:** The SOC becomes an augmented operation where human-machine teaming is the standard. Alert fatigue drops, investigation quality stays consistently high, and analysts are freed from repetitive tasks to focus on validating AI findings and resolving complex escalations.

## Phase 3: Selective autonomy with custom AI agents

**Goal:** With augmentation in place, this phase shifts focus toward expanding autonomous capabilities through a strategic approach. Starting with high-impact, low-effort use cases, organizations can progressively scale their agentic capabilities as performance and confidence improves.

**Actions:**

- **Develop and deploy custom AI agents:** Organizations that have already adopted native agents like the Alert Triage Agent and Malware Analysis Agent may find value in extending these capabilities—especially when their security environment spans multiple tools, platforms, and workflows beyond Google SecOps. In such cases, purpose-built custom agents can be designed to operate across broader contexts, enabling intelligent, coordinated, and scalable response.

For example, a phishing response agent can be built to autonomously analyze suspicious emails, extract indicators of compromise, correlate them with threat intelligence, and initiate containment actions such as disabling links, quarantining attachments, and notifying affected users—even across third-party email gateways or endpoint protection platforms. While some of these actions may be achievable through SOAR playbooks, a purpose-built agent adds contextual reasoning, cross-system orchestration, and adaptive decision-making that goes beyond rule-based automation.

Organizations comfortable with vendor-built agents may choose to accelerate adoption using prepackaged capabilities. Others may prefer custom development to align with unique risk profiles, governance models, and operational complexity.

This follows a structured development lifecycle:

1. **Design:** Using Vertex AI Agent Builder, teams define the agent's purpose—such as phishing detection and response—along with its functions and operational boundaries.
2. **Develop and ground:** Using the Python-based Agent Development Kit, engineers codify the agent's logic. The agent is grounded by connecting to enterprise APIs and data sources prepared in Phase 1, such as email gateways, threat intelligence feeds, and user identity platforms. This context allows the agent to distinguish between benign and malicious activity and take appropriate action.
3. **Orchestrate and deploy:** For more complex scenarios, multiple agents can collaborate using the open Agent2Agent (A2A) protocol. For instance, a phishing response agent may work alongside a user behavior analysis agent to validate anomalies before executing containment. Once tested, agents are deployed on the managed Agent Engine, which handles infrastructure and scaling.

**Performance monitoring:** Implement thorough key performance indicator (KPI) tracking to measure agent effectiveness and guide expansion:

- **Metrics assessment:** Monitor KPIs such as response time, false positive rates, and successful resolution rates to validate agent performance.
- **Continuous improvement:** Use performance data to refine agent logic, decision thresholds, and containment strategies.
- **Strategic scaling:** Based on demonstrated success, progressively expand agentic capabilities into additional security domains—like lateral movement detection or insider threat response—and more complex use cases.

- **Transition to a human-on-the-loop model:** In this mature operating model, the AI agents act autonomously on high-confidence incidents within defined parameters. Human experts—now serving as AI supervisors—intervene only for novel, complex, or high-impact cases escalated by the agents.

**Outcome:** The result is a true agentic SOC. Autonomous agents handle the majority of incident response, enabling human experts to focus exclusively on strategic tasks like threat hunting, defense design, and crisis management.

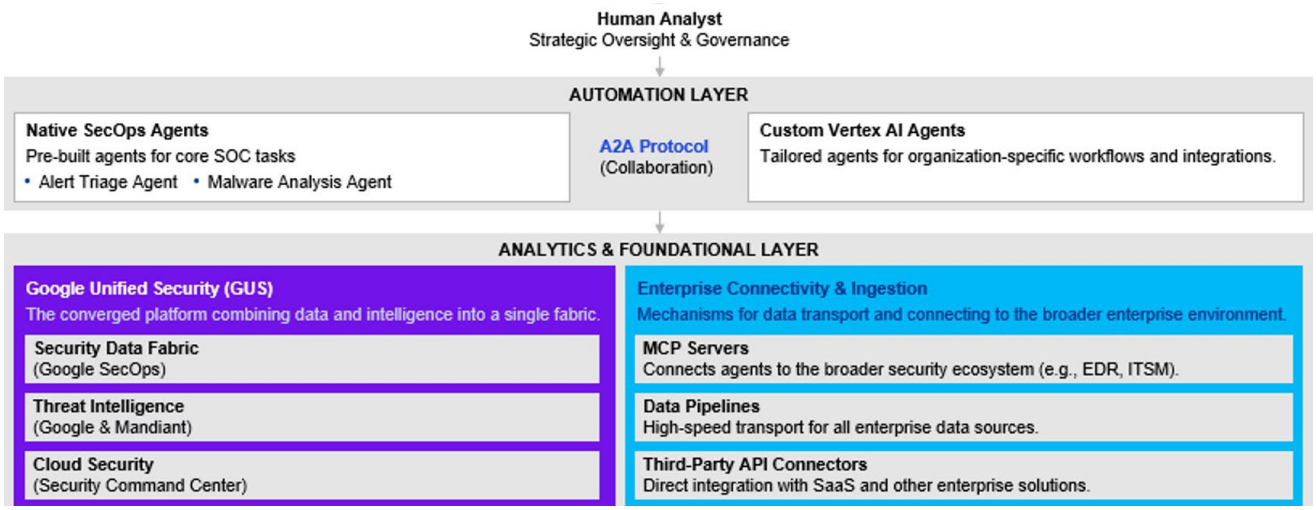


*The visual illustrates how built-in SecOps agents and custom vertex AI agents interact with enterprise data sources through a unified security fabric, enabling automated threat detection and response. Human analysts provide strategic oversight and governance, ensuring safe and intelligent operations across the SOC.*

# Accelerating transformation with a partner

Adopting agentic capabilities in a SOC is a complex transformational journey that affects technology, processes, and people. Attempting it without specialized guidance risks delays, missteps, and missed value. Partnering with a Google Cloud Security Partner like KPMG is a strategic decision—one that enables speed, precision, and alignment with business goals.

## Specialized guidance is essential

KPMG brings deep experience in Google Cloud Security, recognized as a Premier Partner across SecOps Managed Security Service Provider and Reseller Programs. Backed by a significant investment in generative AI, data analytics, and cybersecurity, KPMG guides organizations toward advanced agentic capabilities. This experience means organizations are not only adopting new technology but also working with a team that applies it effectively in real-world, high-stakes environments.

## The KPMG implementation model

KPMG offers a structured, full-spectrum approach to modernizing security operations:

- **Strategy and assessment:** Transformation begins with a clear baseline. KPMG conducts a SecOps capability assessment to benchmark tools, processes, and maturity against industry standards. Based on these findings, they develop a SecOps strategy and roadmap that aligns security goals with business priorities, timelines, and risk profiles.
- **Architecture and implementation:** KPMG designs the target architecture and leads the migration to Google SecOps through the strategic design, configuration, and deployment of the platform—including SecOps SIEM, SecOps SOAR, and Google Threat Intelligence—tailored to the organization's threat landscape and operational needs.
- **Agentic operations services:** Beyond the core Google SecOps platform deployment, KPMG enables true agentic operations through its Cyber Intelligence Automation Orchestration (CIAO) services. These services are designed to extend Google SecOps capabilities by orchestrating AI agents across SIEM, SOAR, and threat intelligence workflows. By automating investigation flows; enhancing cross-platform collaboration; and delivering real-time, context-rich insights, CIAO provides an established path to proactive, AI-driven defense.

## Beyond implementation: Managed services and continuous improvement

Platform deployment is only the beginning. KPMG provides ongoing services to help organizations operate, scale, and mature their agentic SOC:

- **Managed detection and response (MDR):** For organizations without 24/7 in-house coverage, KPMG offers MDR services. Our security specialists operate the client's Google SecOps platform directly, delivering continuous monitoring, threat hunting, and incident response as an extension of the internal team.
- **Holistic defense posture:** KPMG integrates its services with the broader Google security ecosystem to build a thorough defense strategy. This includes Mandiant offerings such as Cyber Defense Center Development, Incident Response Retainers, and Red Team Assessments to validate and strengthen defenses.
- **KPMG Trusted AI:** Governance is central to successful agentic operations. KPMG helps organizations define escalation policies, simulate failure scenarios, establish audit-ready controls for AI-driven decisions, and communicate governance strategy to executive stakeholders. These capabilities are essential for building trust, maintaining compliance, and helping ensure safe operations as autonomous agents take on greater responsibility.

The role that KPMG plays extends beyond technical implementation. It provides the structure, safeguards, and leadership alignment needed to help the agentic SOC operate safely and deliver sustained value.

## Conclusion

This white-paper series has mapped a complete journey—from the persistent challenges of traditional SOCs to the strategic advantages of an agentic, AI-driven defense. We began by identifying the operational gaps that make the status quo unsustainable. We then explored how Google's Unified Security platform enables intelligence and automation at scale. Finally, we provided a practical implementation guide for building a mature agentic SOC.

The agentic SOC is more than a technological upgrade—it's a strategic evolution that redefines how organizations defend, respond, and adapt. It replaces reactive alert triage with proactive, resilient, and business-aligned security operations.

To move forward, security leaders should:

1. **Assess current maturity** to define a realistic transformation roadmap.
2. **Adopt a new operating model** that prioritizes intelligence, automation, and integration.
3. **Invest in people**, upskilling analysts to thrive in a human-machine teaming environment.

The future of security isn't about choosing between humans and machines; it's about building a partnership between them. Those who embrace this new paradigm will be better equipped not only for today's threats, but also for the challenges of tomorrow. To help accelerate progress and ensure success, the next step is to engage with a trusted partner who can help guide this transformation.

## 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**