Google Cloud

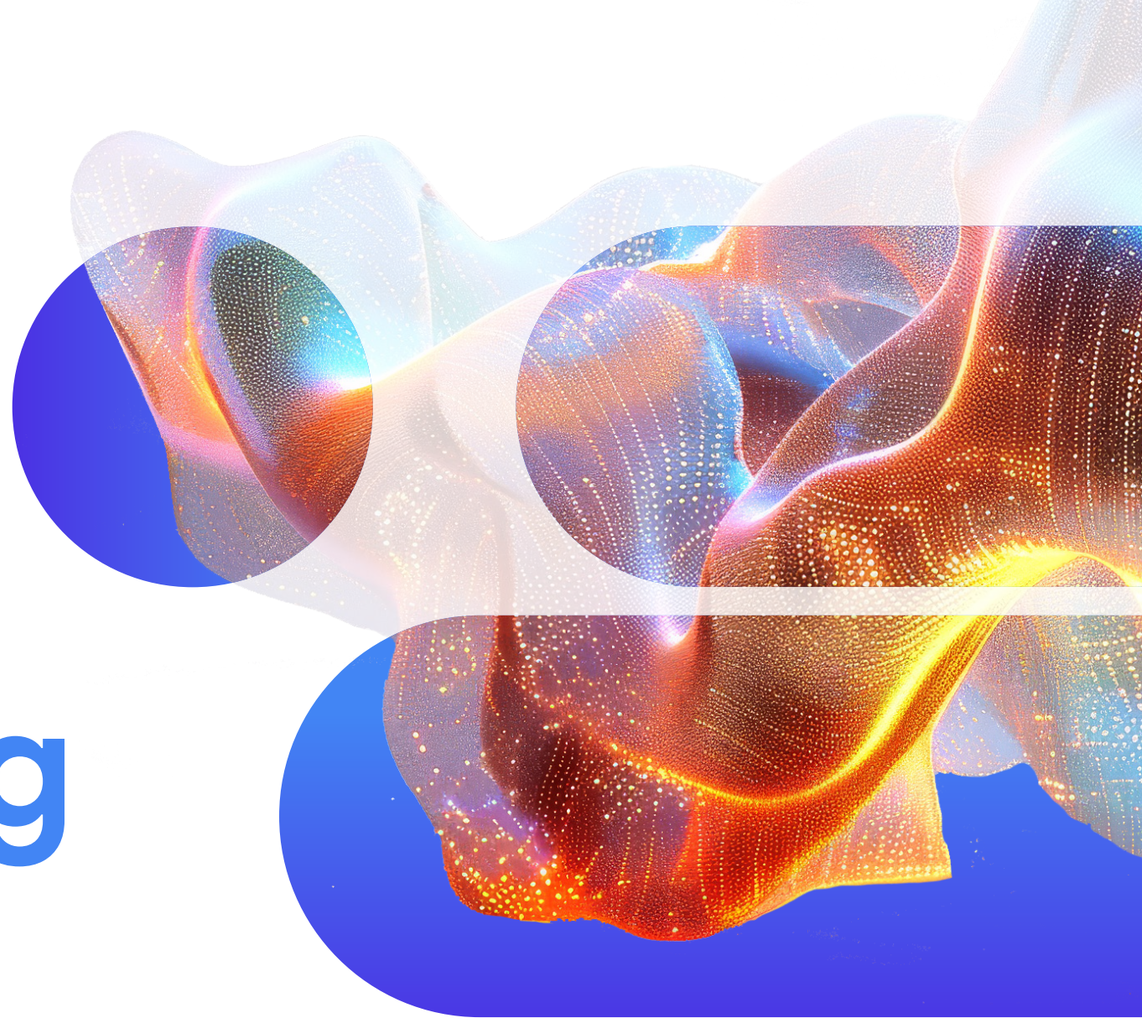

How AI Agents are Transforming Cybersecurity

Product Vision for Agentic Defense

Executive summary

Agentic defense establishes AI agents as a force multiplier for the modern security team. By deploying a connected, multi-agent system that autonomously orchestrates and automates complex security tasks, organizations can counter advanced threats at machine speed, dramatically improving security productivity and significantly reducing mean time to detect and respond.

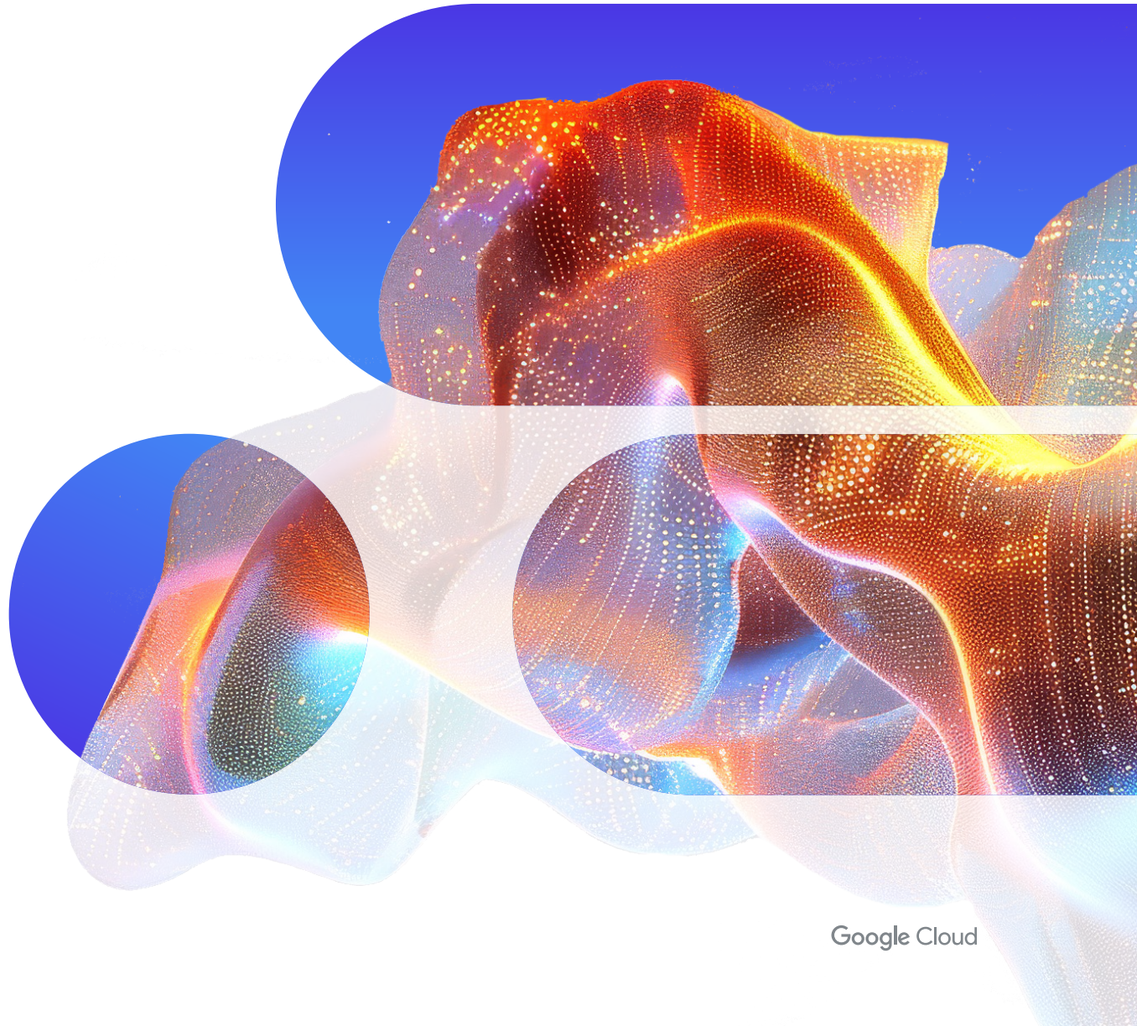

The Security Imperative: Challenges in the Modern SOC

The AI era has fundamentally shifted the security paradigm. As adversaries leverage AI to operationalize sophisticated campaigns at machine speed, security teams—already facing resource constraints—are challenged at every organizational level.

CISOs: Navigating a Dynamic Threat Landscape

The attack surface has moved from static to hyper-dynamic, driven by the rapid adoption of AI. Relying on human reaction times is no longer viable, and can expose organizations to significant financial and reputational risk from agentic attacks.

SOC Managers: Overcoming Operational Paralysis

Security Operations Centers (SOCs) are overwhelmed by data volume and keeping pace with threats. The incessant flood of alerts and false positives obscure genuine threats, demanding a new model for prioritization.

Security Practitioners: Mitigating Toil

Analysts are burdened by manual, repetitive and time-intensive processes, such as threat hunts and rule creation, which prevent them from focusing on strategic investigations. This operational toil is a primary driver of burnout, dissatisfaction, and high turnover among security talent.

Agentic defense

The pace and sophistication of cyberattacks are accelerating and reliance on manual human intervention is no longer viable. To survive this shift, organizations must pivot to an agentic SOC, composed of goal-driven AI agents that function autonomously rather than waiting for human prompting, to deliver positive outcomes on behalf of users.

Agentic attacks require agentic defense.

Fast, accurate detection and response

AI agents observe, reason, and act to accelerate detection and response, continuously fill gaps in coverage, and autonomously investigate and respond in seconds.

Highly automated processes

Systems of collaborating AI agents fully automate routine tasks while continuously streamlining and optimizing processes.

Empowered analysts

AI-driven, human-led processes enable SOC analysts to focus on genuine threats, address strategic questions, and ultimately perform faster and better.

AI agents are blazing a path to transforming the SOC.

The Journey

Google Cloud's vision for agentic defense provides organizations with a path from a manual SOC to an agentic SOC, a multi-agent system that collaborates with human analysts to orchestrate security workflows at machine speed. The path covers these stages:

Manual

Humans cyclically and routinely execute tasks

AI Assistance

Assistants improve analyst productivity

AI Task Automation

AI agents automate specific tasks

Agentic SOC

Collaborative agents run autonomous processes with human supervision

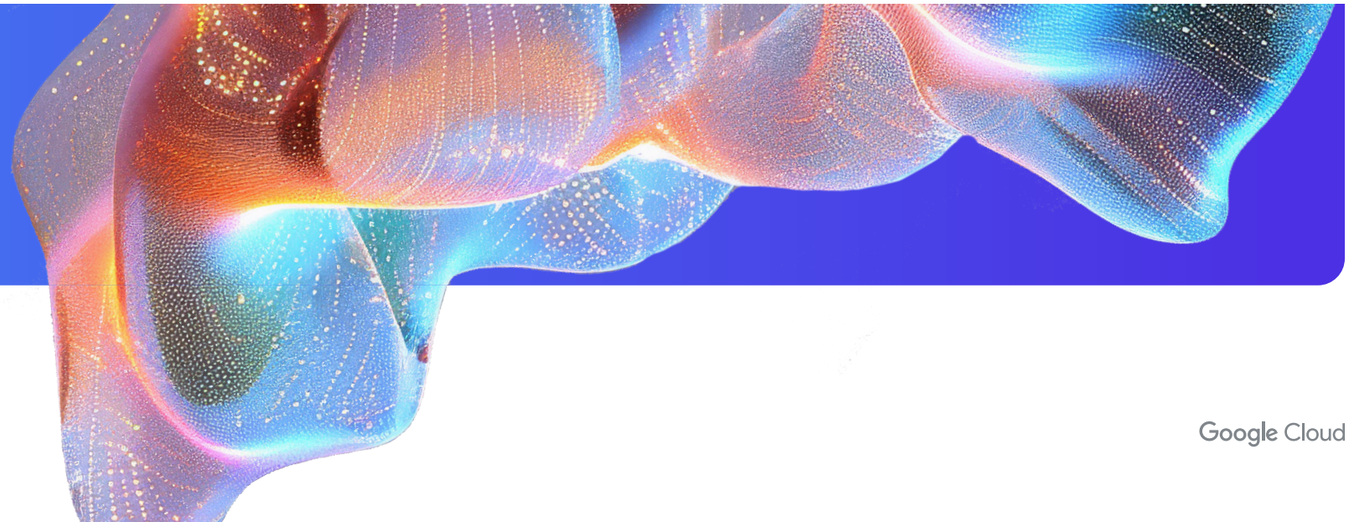

AI Assistance

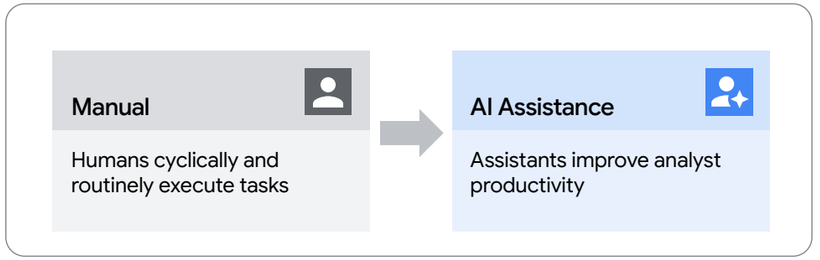

Many SOCs are already using AI assistance to boost productivity and improve analyst decisions today.

Performing manual tasks for analysts

AI assistants can enable analysts to perform many repetitive, tedious tasks faster and more reliably. This includes tasks such as:

- Generating and executing structured and natural language data queries
- Creating event timelines
- Enriching data with threat intelligence
- Creating and refining attack detection rules
- Building incident response playbooks

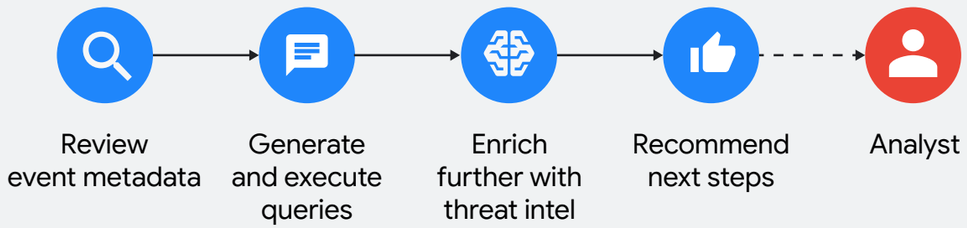

Improving decisions

AI assistants also help analysts make better decisions, for example by providing more complete summaries of cases and alerts and by recommending actions based on past SOC responses to similar challenges.

But assistance is only the first step

As valuable as AI assistants can be, it is important to recognize their limitations. In particular:

- Analysts still have to drive processes by accessing the AI tool, creating prompts, and copying data between security tools
- Analysts must make final decisions and initiate all activities to contain attacks, remediate vulnerabilities, notify affected parties, etc.

As a result, analysts are still burdened with repetitive tasks that often become bottlenecks slowing down timely responses to fast-moving events. In addition, AI assistants are typically focused on limited tasks and have no visibility into, much less ability to optimize, end-to-end processes.

Task Automation by AI Agents

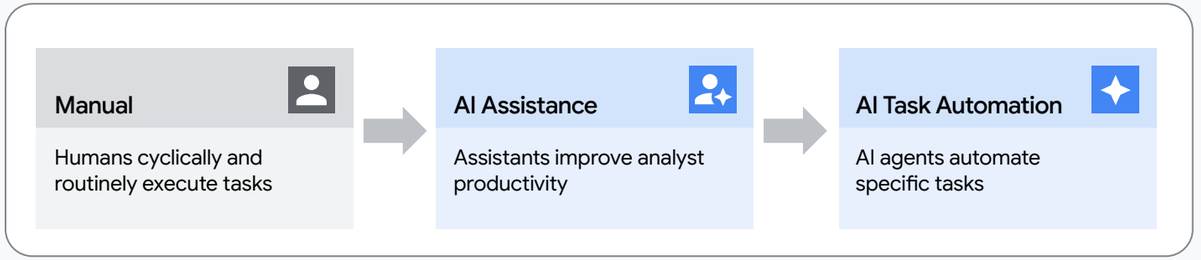

AI agents make a fundamental shift beyond AI assistance because they are fully capable of making accurate decisions and can also initiate actions on behalf of users.

Agents perform multiple tasks at machine speed

AI agents are goal-driven; they know their goal, can autonomously execute actions, and can learn and adapt to improve accuracy and speed. They don't need to wait for a prompt to start working, or wait for an analyst to make final decisions and initiate activities. That means:

- Analysts are freed from many repetitive, tedious tasks
- Decisions can be made and many containment, remediation, and other actions can be performed at machine speed
- Tasks are performed consistently, 24x7, with continuous improvements when conditions change and new information is available

For example, a triage and investigation agent can automatically be triggered by an alert or a similar event; it does not require an analyst to initiate the process or input prompts.

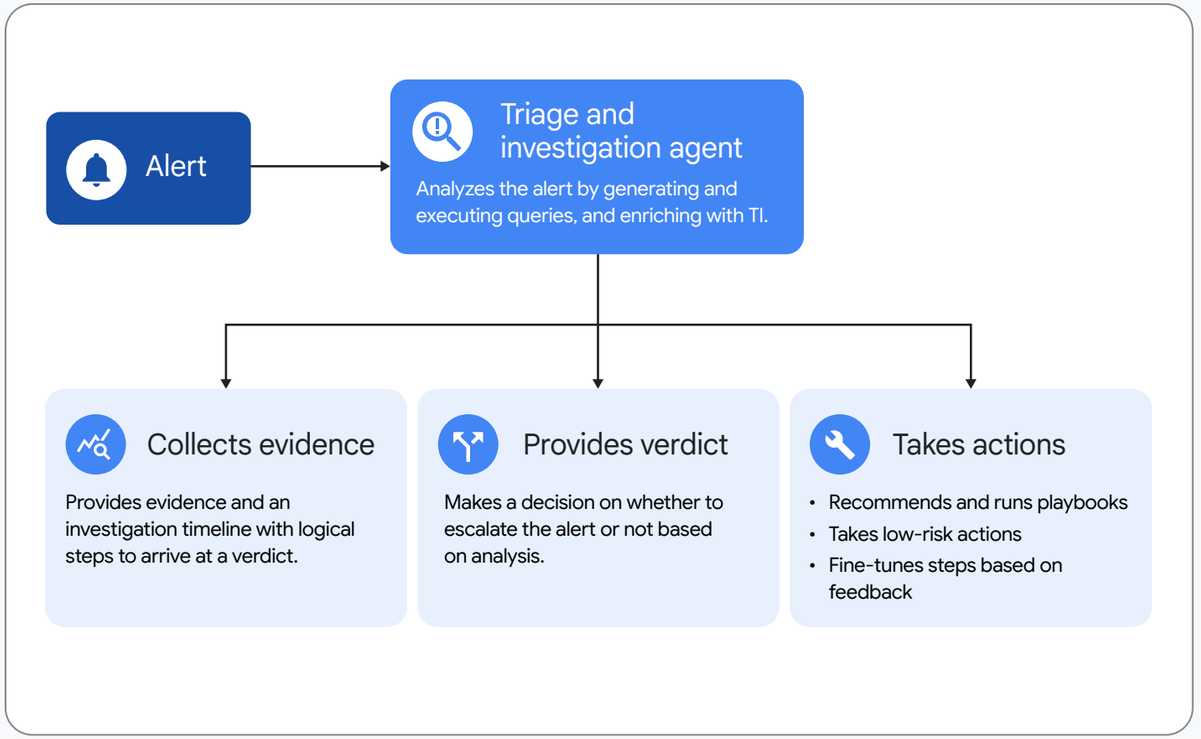

It steps through a series of actions, which might include generating and running queries, checking user permissions, analyzing timelines and process trees, enhancing the data with threat intelligence, weighing alternative courses of action, deciding whether it is necessary to act on the alert, and if so whether to escalate it to an analyst or take immediate action such as running a playbook.

This agent can also fine-tune the steps it takes based on results and feedback, such as modifying queries to gather and correlate new kinds of data, changing the conditions under which it escalates to analysts, or modifying the playbook.

Elevating the role of analysts

While AI assistants make SOC analysts more productive, they still require human participation. AI agents transform the relationship. They can perform a wide range of tasks, on their own initiative, but under the direction or supervision of humans. They elevate the role of analysts from alert triager to agent manager and strategic thinker.

With semi-autonomous agentic processes, decisions and actions can be taken either with or without the participation of humans. That is, AI agents can act autonomously where appropriate, but can transfer decisions to analysts that require human judgment because the information is ambiguous, the situation is novel, or actions involve high risks to the organization. SOCs often face novel threats that AI agents have not been trained to address and these instances still require human intervention.

The Agentic SOC

Today, organizations can go beyond deploying goal-driven AI agents to perform specific tasks.

In an Agentic SOC, multiple AI agents collaborate with each other and with external resources to manage and improve *end-to-end processes*.

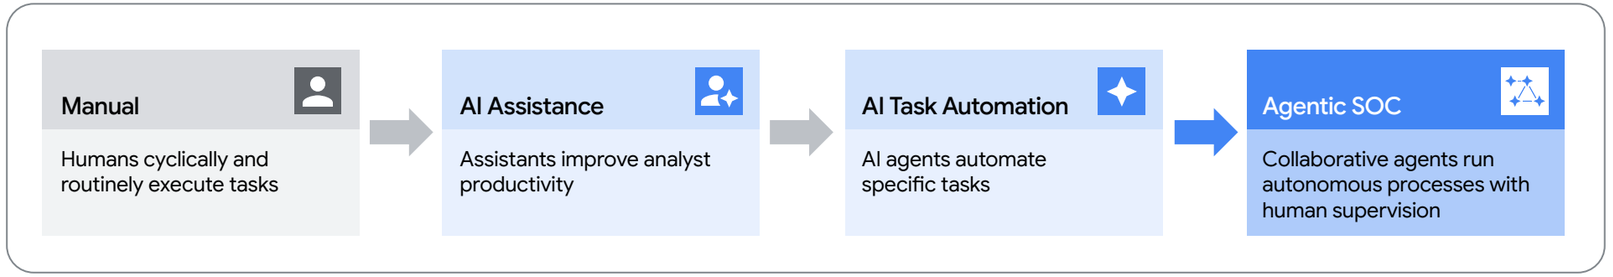

The Agentic SOC

A fleet of autonomous SOC agents performing core tasks at machine scale and speed

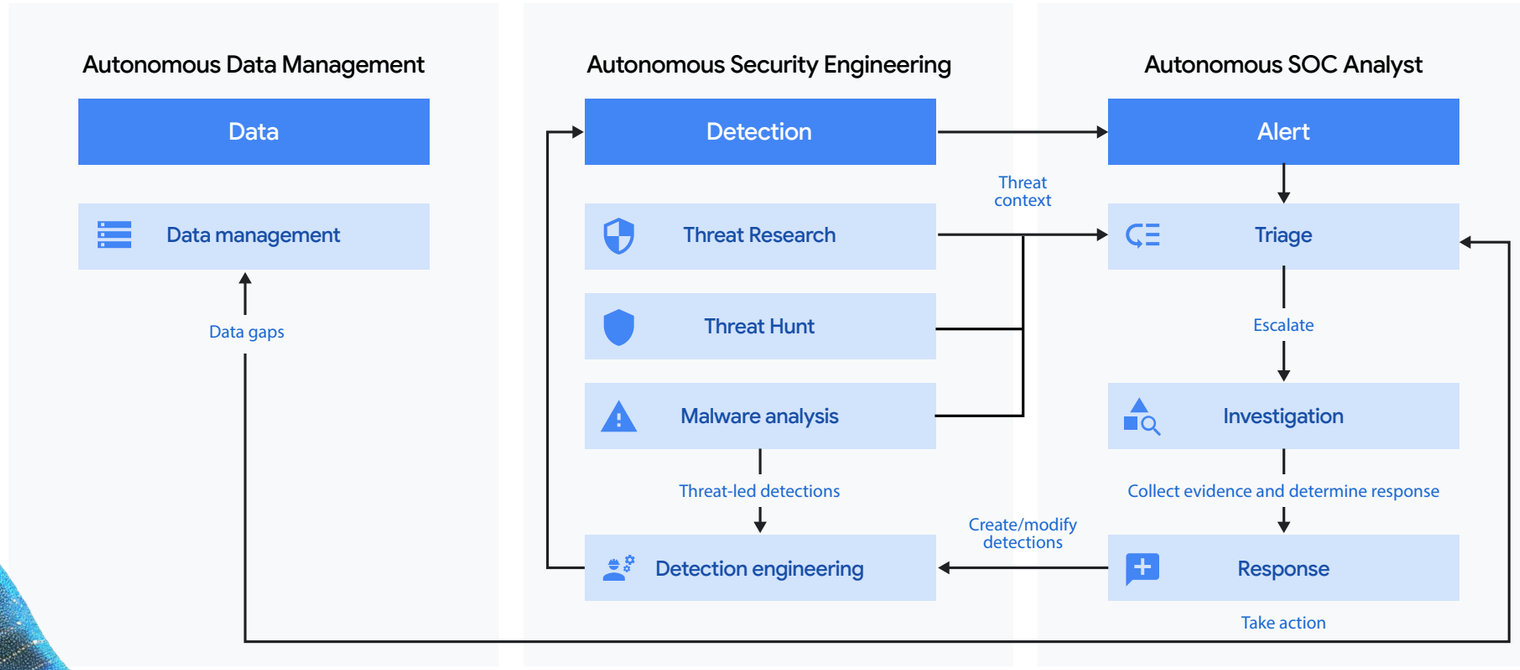

Orchestrating a system of AI agents

For example, although a triage and investigation agent does a lot of work by itself, its value is multiplied when it is integrated into a system of collaborating task-specific AI agents across security operations workflows.

For instance, it might:

- Draw on the services of a specialized **data management agent** to collect, normalize, and organize data from a wide range of data sources, network and system logs, and file stores.
- Pass data to specialized **response agents** that can bring in additional context, create playbooks, or take actions to contain attacks.
- Interact with a **threat hunting agent** to immediately search for indicators of compromise to help determine the scope and severity of security incidents
- Leverage a **detection engineering agent** to generate persistent detection rules

In an Agentic SOC, the agents work together to combat attacks at machine scale and speed. To observe and act like an expert human analyst, agents are trained on real-world insights and are continuously learning and applying improvements throughout the cycle. Agents also leverage active threat intelligence to proactively uncover and defend against the latest threats.

For example, Google security agents can generate new detection rules in a fraction of the time based on new insights from Google Threat Intelligence.

Through the work of detection engineering agents, security teams see fewer false positives and can more rapidly uncover indicators of compromise.

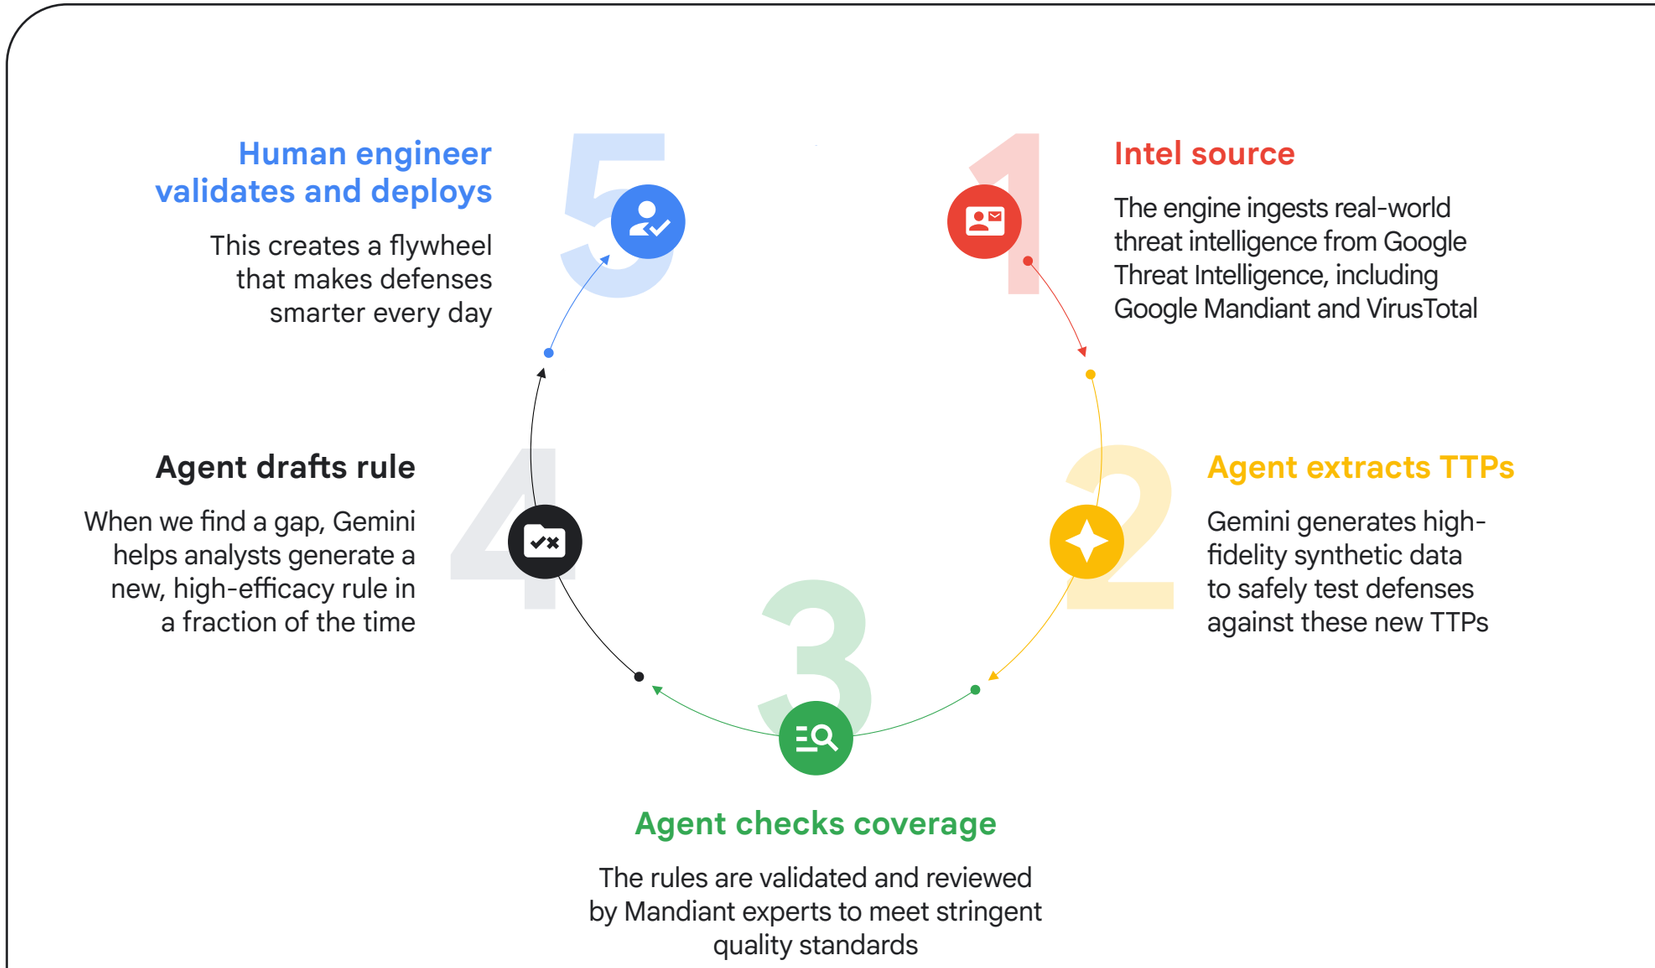

An open ecosystem for models and agents: Flexible, standards-based collaboration

Part of the power of Agentic SOCs derives from the ability of multiple AI agents to not only collaborate with each other, but also with external resources such as data sources, security and IT operations tools, business applications, and development environments.

Our vision for agentic defense ensures engineers and agents can easily interact with data and take actions. We support Model Content Protocol (MCP) and the Agent2Agent (A2A) protocol, which provide AI agents with a common communication channel to share insights and coordinate actions.

The use of flexible, standards-based interfaces also make it easier for users to build their own enterprise-ready security agents and workflows. Organizations can plug in agents and tools they have developed internally, as well as third-party data sources.

Service providers and standards groups will continue to refine and enhance the A2A protocol, MCP, and other standards that enable collaboration of AI agents with each other with external resources. The scope of the standards will expand to facilitate more sharing of data and insights, and also to address issues related to authentication, privacy, resilience, compliance and auditability, as well as end-to-end observability and measurement across multi-task processes.

Human in the loop, Human on the loop

AI agents act as force multipliers to augment security operations teams, not replace them. Google Cloud's vision embraces AI agents that act as trusted teammates to automate workflows and tackle routine tasks, ultimately accelerating investigations and multiplying human expertise.

In some agentic workflows, agents make decisions to escalate based on guidelines and policies. When that happens, the agents provide clear summaries and references of the sources they used. This information enables analysts to understand and validate the reasoning underlying the agents' actions and recommendations. Ultimately, the analysts guide the system.

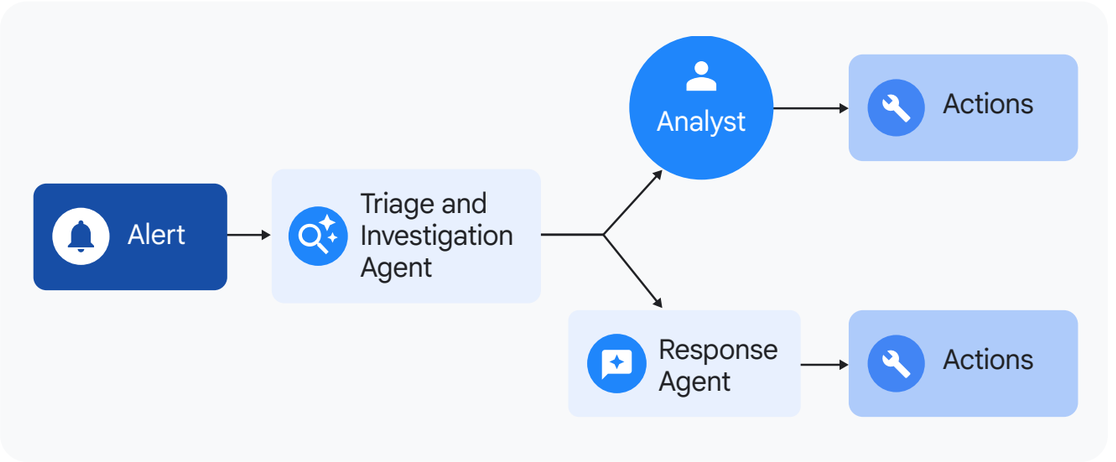

Impacts Across the Security Organization

The agentic SOC not only benefits analysts, it can have a major impact across almost every corner of the security organization.

With agents, security professionals spend much less time on tedious, repetitive tasks. They turn their attention to more productive and rewarding activities such as advanced threat hunting, business risk analysis, security architecture, AI coaching, and strategic planning.

According to recent research from Omdia, the agentic SOC is already delivering a measurable impact: of early adopters already report a positive contribution from agentic AI solutions in their security operations solutions.

This transformation benefits the entire security organization, not just analysts, with over 87% of adopters reporting a meaningful positive impact across all key SOC functions/KPIs surveyed, including threat detection times, analyst throughput, and alert fatigue.

The chart below highlights just a few of the security roles that are empowered by agentic AI processes.

Furthermore, 50% of security practitioners believe agentic AI will improve their personal career prospects over the next few years (source: Omdia). This optimism is even higher among more experienced professionals (56%) and early adopters of agentic solutions (69%), indicating that familiarity with the technology increases confidence in career advancement.

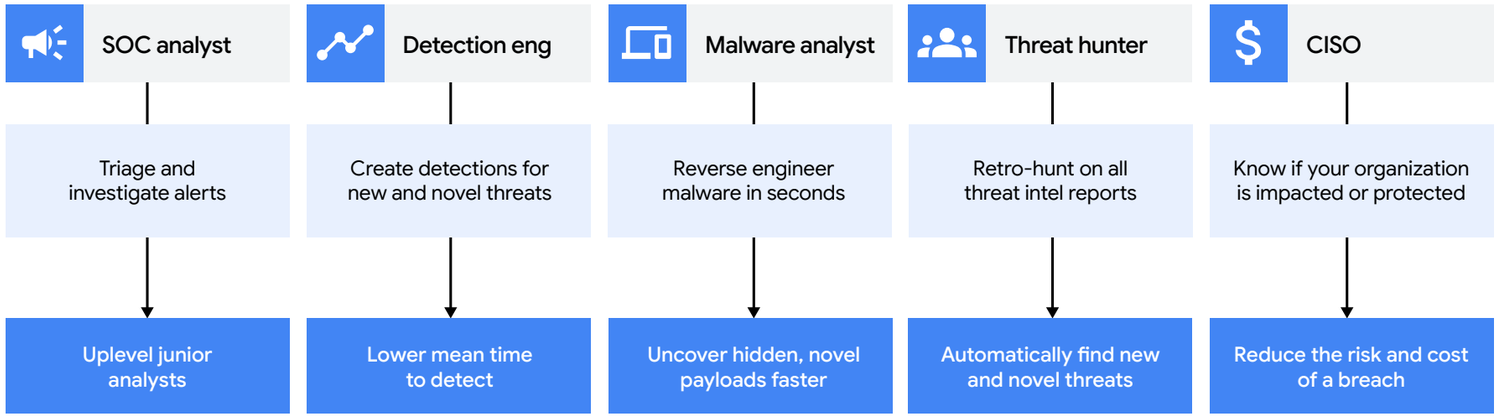

Google's approach to agentic defense is different

Our mission is to be the world's most trusted security partner in the AI era, helping every organization accelerate their security transformation with the same AI-powered cloud technology that runs Google. We bring AI innovation to security at scale, delivering:

- **A full AI stack built on the most secure cloud.** From chips to models, and research from DeepMind we co-design every layer, so security is not bolted on, it's engineered into the foundation.
- **Unmatched visibility and applied intelligence.** Google has the world's largest threat observatory with telemetry from Chrome, Android, Gmail, YouTube, VirusTotal and Mandiant. We apply the most relevant and prioritized intelligence to our tools.
- **Mandiant frontline expertise.** Because our agents are trained on the specific cognitive workflows of elite Mandiant analysts and everything Google sees and knows, they can observe and act like an expert, providing high quality results.

We take a disciplined approach to hill climbing to achieve greater levels of accuracy. Without AI-enabled tools, an average enterprise analyst spends 30 minutes triaging a single alert. AI agents cut that time down by orders of magnitude and reduce the volume of alerts that need intervention, generating dramatic cost savings. Google customers like Lloyds Banking Group and Sunrun are already seeing results. A Total Economic Impact study from Forrester found that our approach results in a reduction in mean time to respond by 50% and mean time to investigate by 65%, as well as a 70% reduction in the risk and cost of a breach. This is compelling evidence that faster detection directly correlates to bottom-line protection.

Build what's next

Conclusion: The Future of Agentic Defense

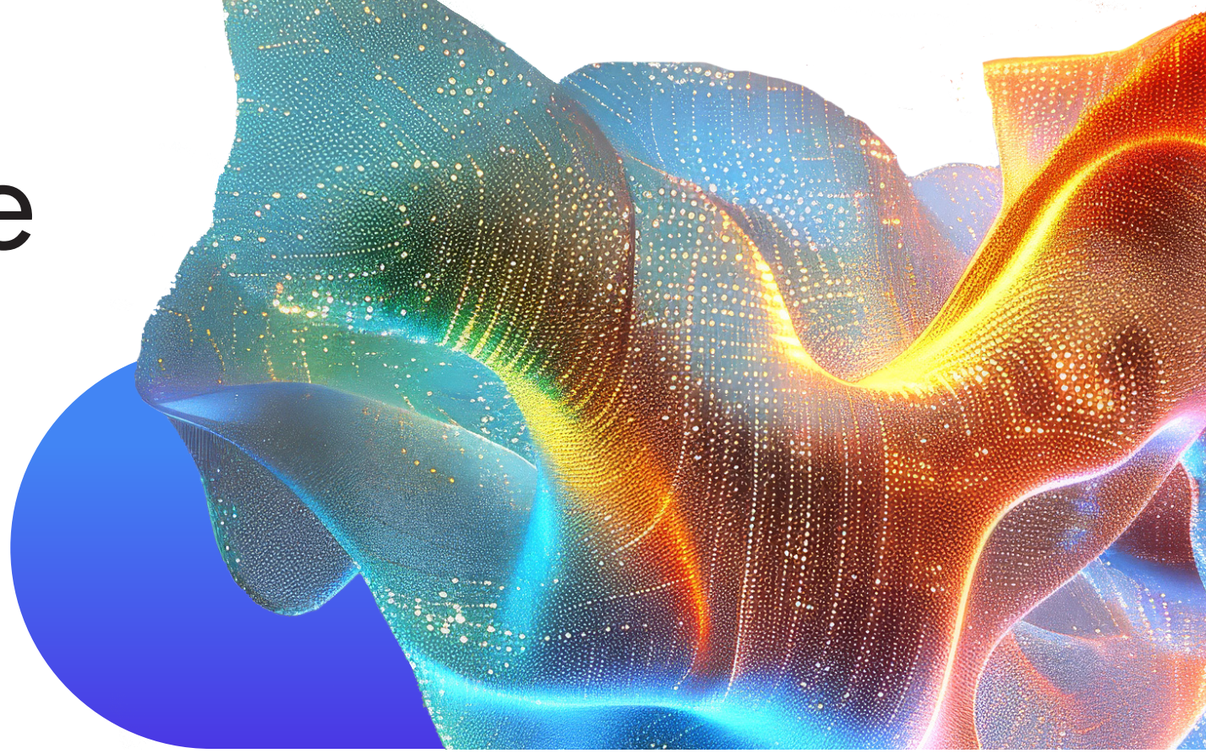

The transition to an agentic SOC is no longer a luxury but a matter of urgent necessity to defend against adversaries operating at machine speed. As AI continues to permeate security operations, we see organizations moving beyond assistive tools toward agentic defense.

In the near future, we'll see the development of more task-specific AI agents as organizations become more familiar with their use. SOC processes will have increasing levels of automation and autonomy. The organizations that successfully integrate these autonomous capabilities will be best positioned to outpace emerging threats and maximize business resilience.

As Agentic AI processes mature, organizations will be able to observe and measure results in more detail and tie those results to business outcomes such as fewer breaches, higher availability, lower operations costs, and reduced risk. AI will continue to be a powerful tool for optimizing the security lifecycle to maximize business outcomes.

Learn more

Visit the Google Cloud Agentic SOC webpage to learn more, get in touch with our team, or register for an upcoming event.