Google Cloud Deloitte.

# Scaling security defenses:

## The future of SecOps with agentic AI

“Right now it’s an overlapping two-dimensional problem: AI is needed to accelerate SecOps, and SecOps needs to adapt to secure AI. Security leaders must safeguard operations and their own transformation—embedding governance, detecting model exploitation, and enhancing security processes simultaneously.”

**Naresh Persaud,** Principal, Cyber Risk Services, Deloitte & Touche LLP

# Introduction

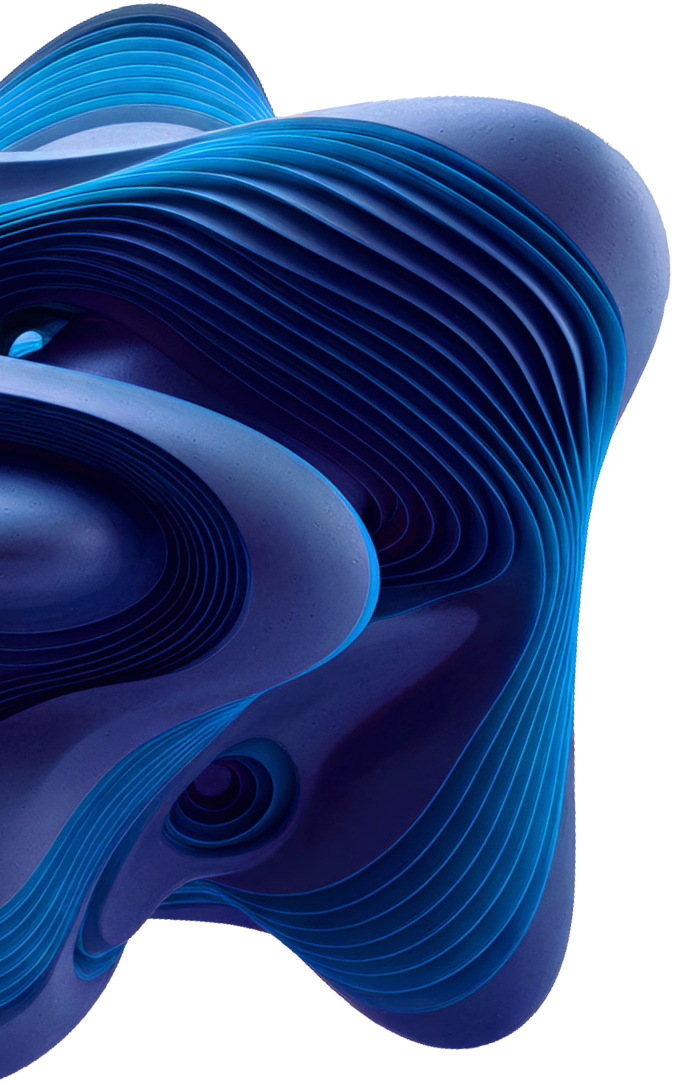

The modern security operations center (SOC) was built for calmer times. The conventional SOC's structure, once a reliable keep, now creaks under the weight of its own inefficiency with larger backlogs. As CPO Magazine contributor Dave Brown writes in *Beyond Burnout: Three Ways to Reduce Frustration in the SOC*[1], 71% of data analysts report burnout, and 64% consider abandoning their posts within just a year. Still, the Trojan horses keep pounding at the gates. This untenable condition calls not for just patchwork repairs, but a rebuilding of the fortress.

First, a shift in mindset and method is needed. Modern SecOps teams should move from disconnected and reactive to a proactive engineering-focused command center. Second, chief information security officers (CISOs) must take a technological leap. When deployed wisely, agentic AI becomes the vigilant steward of the digital realm, sensing, reasoning, and responding to threats with unmatched speed and precision. By combining human defenders with agentic AI analysts, organizations can move beyond automation to an autonomous force at the heart of security operations—lifting the weight of analyst overload.

Read on to learn how leaders can strengthen their security operations with agentic artificial intelligence (AI).

1. Dave Brown, Beyond Burnout: Three Ways to Reduce Frustration in the SOC, CPO Magazine, August 13, 2025

# Strained defenses

## Combatting analyst burnout and lost talent

The modern SOC should function like a well-guarded fortress. Each alert is a signal of incoming bombardment, demanding immediate attention and decisive action to preserve the stronghold's integrity. Yet, defending against threats isn't nearly that simple, or romantic. Organizations face a constant flood of data, alert fatigue, and analyst burnout that threatens to sideline even the most vigilant guardian.

To complicate matters further, at least 50% of these alerts — perhaps even higher — are just "noise," Jeffrey Giannetti, brand contributor for Forbes, writes in their *Make False Positive Reduction A New Year's Resolution* article.[2] Giannetti writes that these false positives waste at least a quarter of analysts' time, sapping SOC teams of valuable time, energy, and efficiency. The sheer volume of alerts makes an effective

2. Jeffrey Giannetti, Make False Positive Reduction A New Year's Resolution, Forbes

response difficult. Given this pressure, it's not surprising that 71% of analysts burn out and 64% consider leaving their jobs within just a year.[3]

The root of the problem lies in an operational model built for a simpler time, relying heavily on manual monitoring and reactive responses. This often leads to inefficiencies and wasted effort, and siloed systems are also a problem. Much like a fortress designed without modern surveillance, the traditional SOC struggles to manage today's complex data environment. This creates a cycle of inefficiency that fuels alert overload, analyst exhaustion, and talent loss. Beyond analyst burnout, threat response itself can be hampered when teams lack the ability to take immediate action, making rapid, automated responses a critical factor. Addressing this challenge requires more than additional staff or tools—it demands a fundamental transformation of the SOC model, one that employs agentic AI for proactive and scalable defense.

3. Dave Brown, Beyond Burnout: Three Ways to Reduce Frustration in the SOC, CPO Magazine, August 13, 2025

## The agentic SOC shift from reactive to autonomous

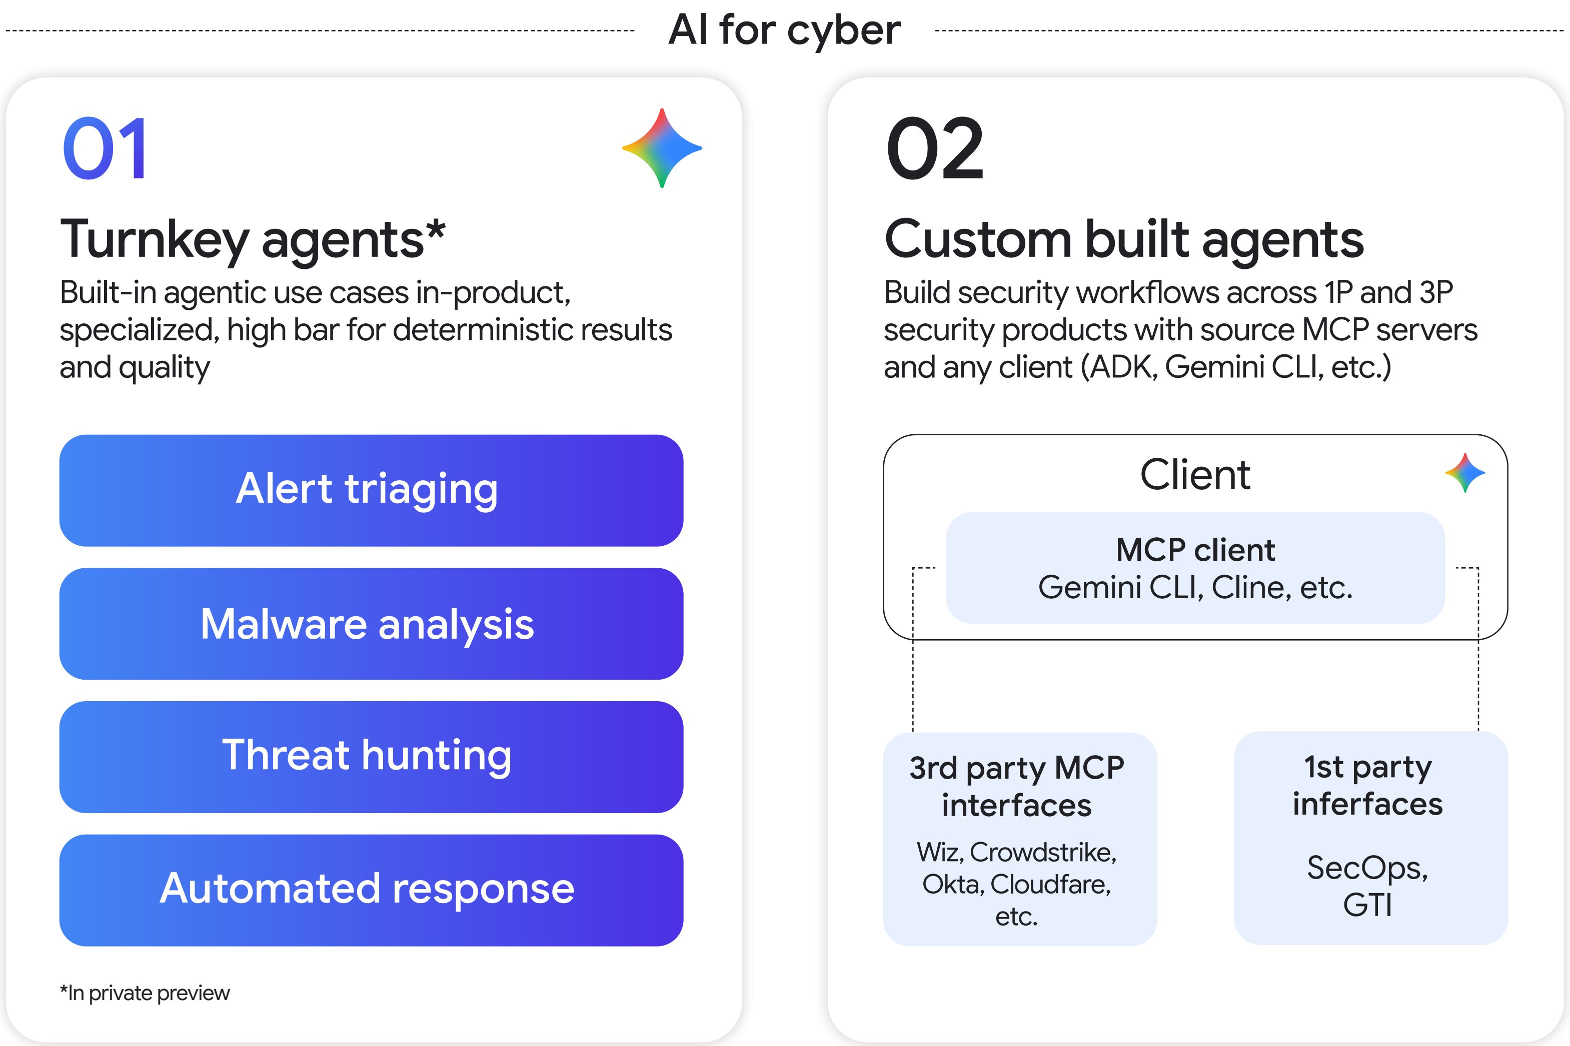


Source: Google Cloud

# Raising the drawbridge

## Scaling security operations with agentic AI

CISOs should fortify their security posture, and technology is the mortar that strengthens the walls. With the addition of agentic AI, SecOps teams can strengthen their defenses by detecting and responding to threats quickly and efficiently. By prioritizing alerts and automating routine responses, agentic AI helps free analysts from the endless siege of alerts, so they can focus on the most critical threats. When every guard knows which threat to focus on, scattered watchposts become a united SecOps front.

Additionally, agentic AI scales the expertise of seasoned SOC analysts across multiple agents. This allows generalist security professionals to handle more complex problems and build their skills faster. Beyond replicating knowledge, AI uncovers edge cases and detects patterns and anomalies that might otherwise be missed.

For one large Fortune 500 client, adopting Deloitte’s Digital Analyst for SecOps had an immediate impact, with additional soft benefits as the AI solution evolved. Initially, the client started with fragmented data, multiple ticketing systems, and a largely manual process. The deployment of the Digital Analyst agent consolidated the data and created a predictable and repeatable model to ingest and manage events. The solution demonstrated geometric scale—the client was able to resolve 50% of the incident volume with the agent, which allowed the human analysts more time to focus on the incidents that needed their attention. The average time to handle an incident was reduced by 10-20 minutes, and the mean time to detect (MTTD) was reduced by 60%. As the solution matured, supervised learning feedback is enabling incremental accuracy and efficiency gains.

"With AI-driven systems, security teams can handle more alerts and scale operations geometrically with only fractional human involvement needed as the volume of incidents increases. This can help teams reduce incident backlogs and security risk."

**Naresh Persaud,** Principal, Cyber Risk Services, Deloitte

# From watchtower to stronghold: Agentic AI in SecOps

| Use case name | Scenario | AI solution description | Indicative impact |
|---|---|---|---|
| Advanced phishing detection | Current phishing solutions generate high number of false alarms and often struggle in adapting against the advanced phishing attacks that evade traditional defenses. | Leveraging intelligent data ingestion, content analysis, and threat detection, offering automated responses to attacks while also assisting in training users to recognize and avoid phishing scams. | Improved effectiveness by 60-70% through smart detection and remediation capabilities |
| AI-empowered ransomware detection and response | Rapidly evolving ransomware are able to bypass old defenses and traditional methods are lagging timely detection and often produce slow or inaccurate results. | Enabling AI agents to monitor activities and early alert triage by detecting threats at their onset, preventing escalation into ransomware attacks and allowing for automated responses. | Improved detection and response capability leading effectiveness gain by 40-60% |
| Digital SOC analyst | Security operation centers rely on analysts to identify and handle security events, but the manual work and large number of alerts make this process time-intensive and challenging as systems grow more complex. | Using large language models to use past incidents and threat data to summarize alerts, identify causes, and using AI agent model to triage, respond to alerts, helping SOC analysts work more efficiently. | First-level response achieves efficiency, reduces blind spots, and lowers overall risks. Cost efficiency gain by 30-35% |
| AI penetration testing | Conventional penetration testing depends heavily on a small number of specialists and manual processes, which can overlook sophisticated threats. | A modular, agent-based AI approach, complex cyber tasks are divided among specialized agents, with flexible workflows that improve efficiency and significantly augment, accelerate, and strengthen offensive cyber capabilities. | 30-40% efficiency gain of pen testers effort though methodical discovery of controls, fast analysis, and reporting |
| Enhanced anomaly detection for OT | Organizations face challenges securing operation technology (OT) because of limited threat visibility and complex attacks. | An agentic solution that fine-tunes anomaly and novelty detectors for OT datasets, expands attack type classification and contextualizes host metadata can transform detection and response. | Anomaly scores and severity with up to 90% precision are improving effectiveness and consistency across investigations |
| Automated vulnerability testing and remediation | Typical vulnerability management involves overload of findings, inefficient quality and tedious path verification process, leading to inefficient remediation and limited visibility of the risk. | An agentic AI based solution can help drive continuous compliance, risk exposure, resource reprioritization, while driving efforts for cost optimization as well. | 40-50% efficiency gain on current efforts and driving towards building a transformed exposure solution |

# The autonomous future

## SecOps with agentic AI

Agentic AI is transforming what SecOps teams can achieve. But every advantage comes with risks. A key challenge in adopting agentic AI is establishing trust within the organization. In SecOps, agentic AI often acts as a force multiplier, helping teams scale their efforts. However, early planning and preparation are essential to ensure adoption is successful.

Having the right partners and strategy is critical. Making strategic decisions early helps guide how agentic AI is adopted and optimized, helping to ensure that organizations get its full benefits. But not everyone is on the same page. Executives often have a mixed mindset toward AI. There is fear of unknown risks, particularly among more risk-averse CISOs. Some leaders worry about missing out and are eager to charge ahead at full speed. Balancing risk-taking with risk management is vital.

It's also important to implement a structured adoption blueprint that covers talent, tools, methodologies, and secure risk frameworks. Aligning internal and external audits, maintaining visibility and controls, and embedding safety mechanisms—much like brakes in a car—can build confidence. Additionally, ongoing oversight of these controls ensures they remain effective as AI and autonomous agents take on more of the cybersecurity workload.

Organizations should adopt flexible tools, but tools alone are not enough. Technology must be applied in the right ways through adaptable models and tailored implementation. Moreover, organizations need a technical architecture that's both scalable and resilient. They also need the right talent model and a rationalized operating approach that answers critical questions like:

- How do we scale our workforce?
- How many level-one analysts can be augmented or replaced?
- What is the evolving role of level-two analysts?
- What does the new operating procedure look like?
- What does the operating model look like across organizations?

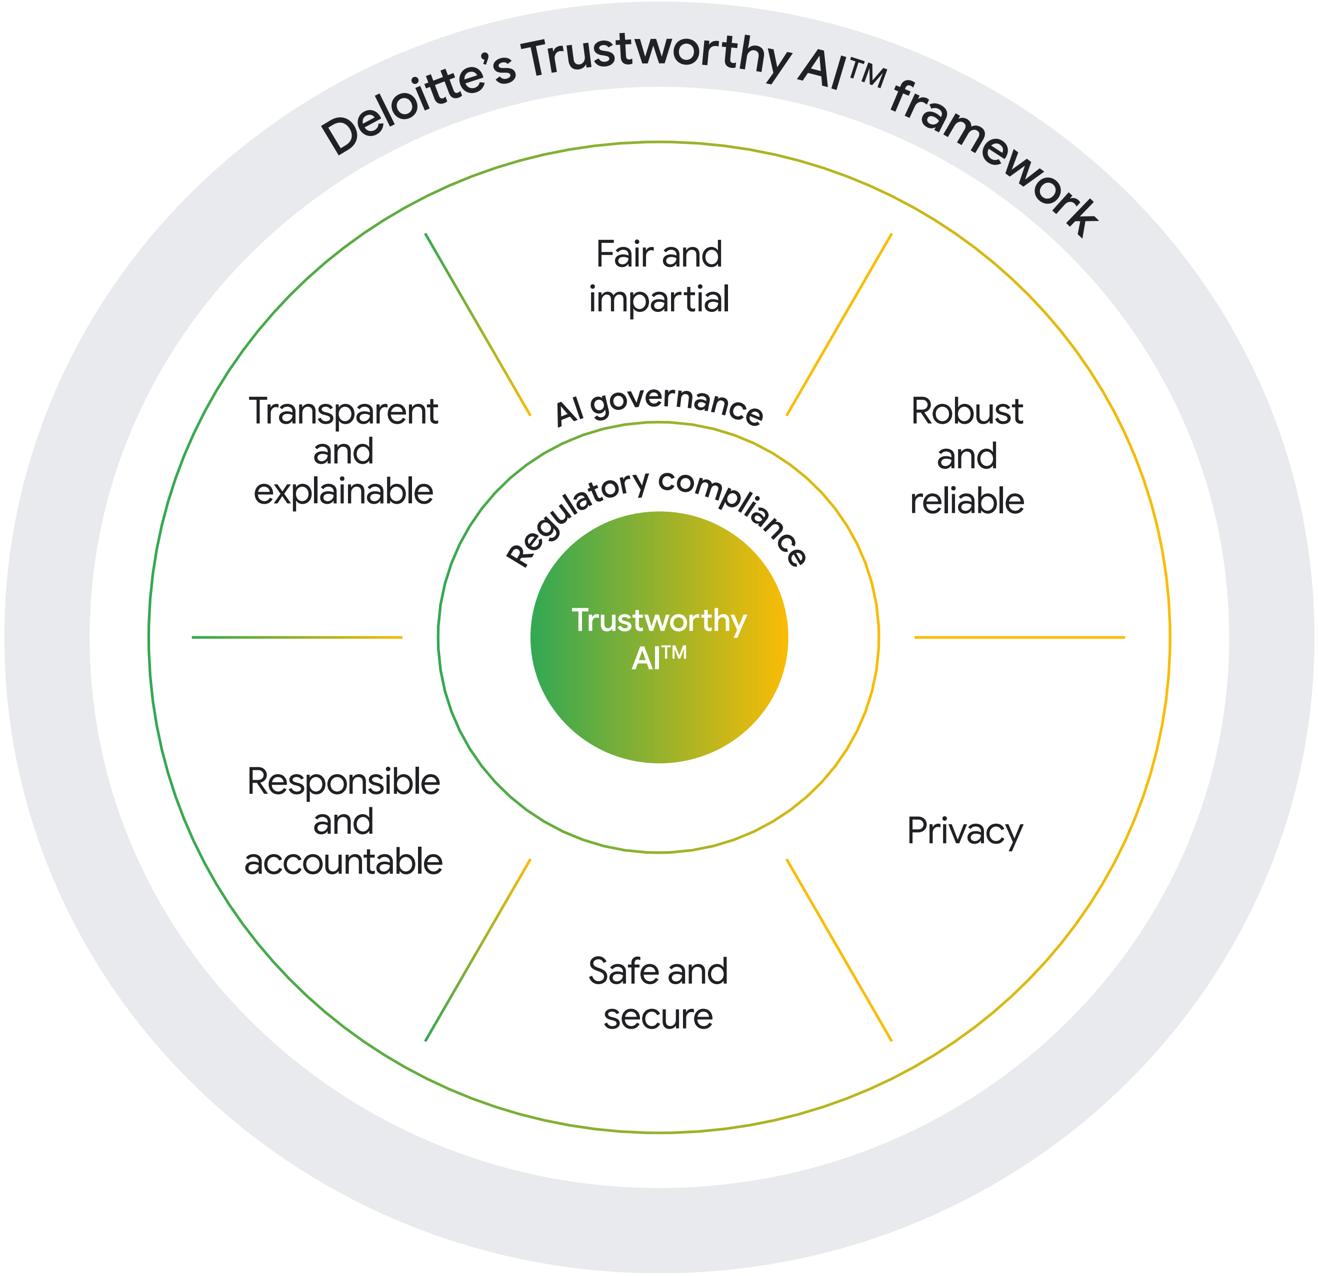

To sustain and mature AI-powered security operations, organizations should transform their talent models—relying on both AI-skilled resources and professionals with deep security expertise.

Strong governance and transparency are vital. It's not enough to simply deploy a tool—leaders must demonstrate to auditors and stakeholders that systems are properly governed. If something goes wrong, how do you troubleshoot? How do you validate performance and tune up the system over time? A Trustworthy AI™ Framework is essential to ensuring visibility and control. That framework should be fair and impartial, transparent, responsible, robust, and prioritize privacy and safety.

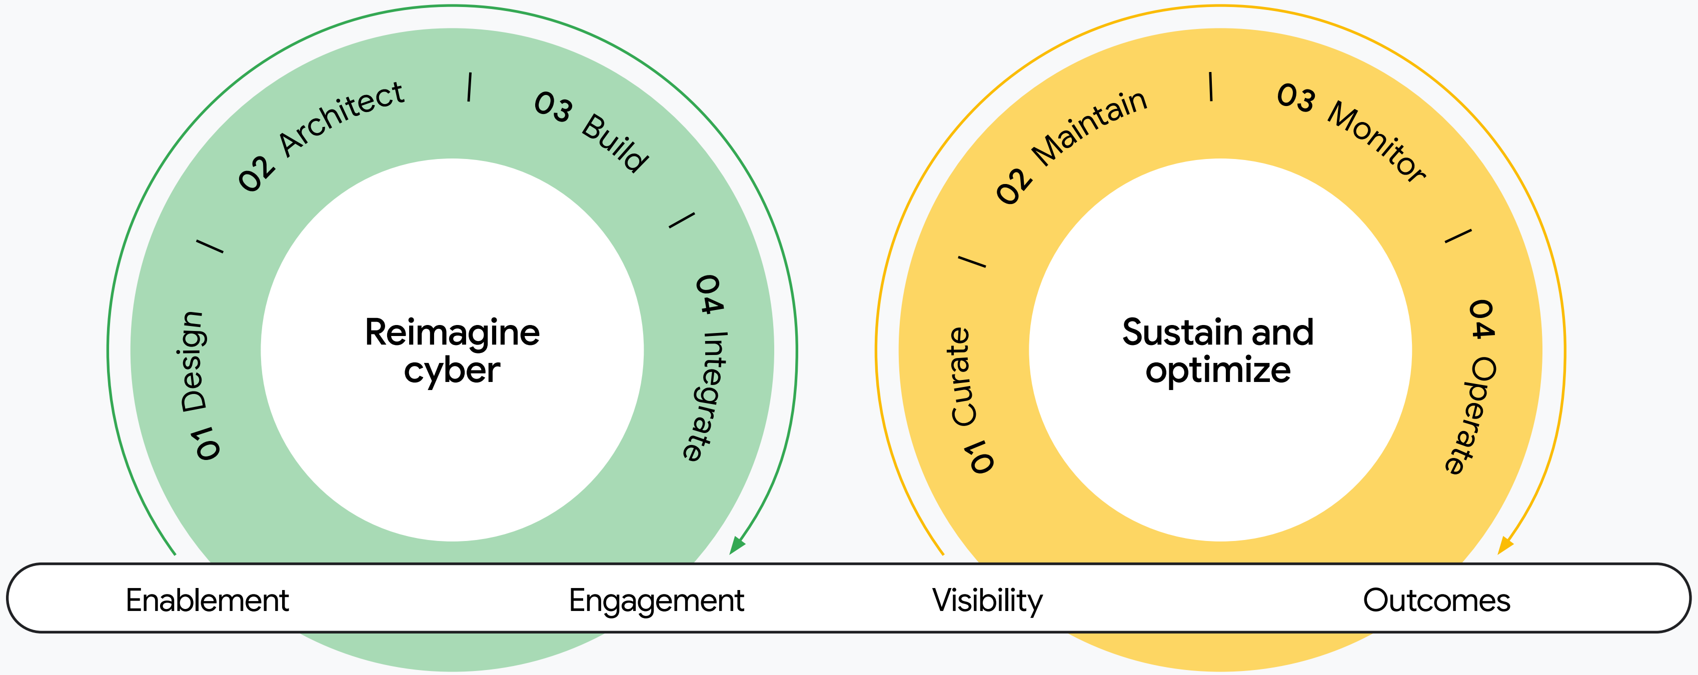


**01 Design:** AI Policy Advisor | AI Auditor | AI Solution Analyst
**02 Architect:** AI Engineers | AI Learning Specialists | AI Developers
**03 Build:** AI Solution Engineers
**04 Integrate:** AI Model Builders | AI Data QA

**01 Curate:** Selection | Organization | Preperation
**02 Maintain:** Manage performance | Fine tuning | Enhancing
**03 Monitor:** Quality/accuracy | Availability | Security/Integrity
**04 Operate:** Service level agreements | Service management | Operating procedures

**Define the architecture** across cyber services
**Design the AI components**/experience
**Build the models** (integrating with LLMs)
**Integrate the data systems** to operationalize

**Curate the workforce of bots/agents** and models
**Maintain/support the software** and systems
**Monitor the workload** and results
**Operate/support the sytems** and processes

# The agentic SOC shift from reactive to autonomous

**Reimagining the cyber org built on and run by AI**

**Securing the AI-enabled systems and data that support the business**

Acquiring the right talent and agentic capability to run the cyber organization and achieve the automation and outcome benefits

Defining the new operating processes/ roles and responsibilities which incorporate AI-enabled automation

Executing on a technology roadmap that brings AI capabilities to refresh the platforms and systems underpinning cyber

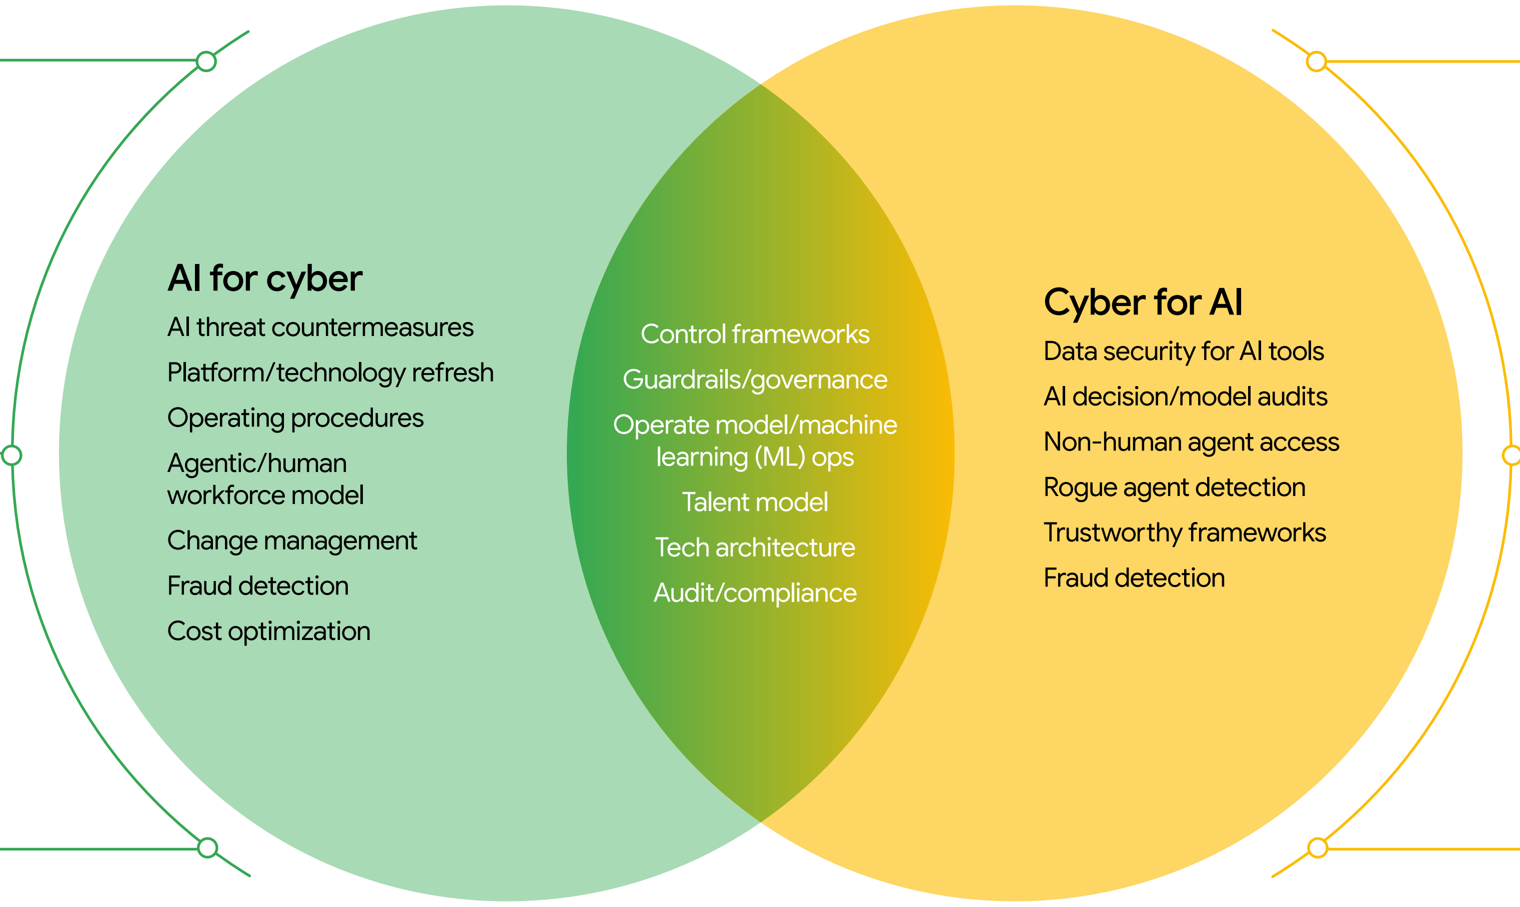


Defining new security services to enable business leaders to securely deploy and adopt AI-enabled services

Assessing the risk and implementing controls to address compliance/security requirements of adoption

Establishing the governance framework and change management to build trust in AI adoption

**Reinventing how cybersecurity works in an AI-powered organization**

Source: Deloitte Development LLC: Reimagining the modern cyber organization

> "Adopting agentic AI requires building trust within the organization. Having the right partners and upfront strategy is critical for effective adoption and optimization."

**George Stamos,** Global Partner Development Manager at Google Cloud

# Guarding the gates with an agentic-enabled SOC

Agentic AI is more than a tool—it's the elite guard of modern SecOps. By combining human specialist knowledge with autonomous intelligence, organizations can reduce alert overload, detect threats at lightning speed, and take decisive action without delay. But a resilient security operations center is more than walls and towers — it requires a foundation of strategy, governance, and trust. Once these bricks are laid, agentic AI can transform a flood of alerts into clear signals and connect scattered watchtowers into a unified command. For CISOs, the path forward is clear: fortify the walls, empower the defenders, and let agentic AI extend its reach—turning today's SOC into an ironclad fortress ready for any attack.

Don't let analyst overload weaken your defenses. Book a complimentary AI readiness assessment to learn how agentic AI can help prioritize threats, reduce analyst burnout, and strengthen your security posture.

Contact us

Learn more about Deloitte and Google Cloud. For more information, contact your Deloitte or Google Cloud representative.

# About

## Google Cloud

Google Cloud is the new way to the cloud, providing AI, infrastructure, developer, data, security, and collaboration tools built for today and tomorrow. Google Cloud offers a powerful, fully integrated and optimized AI stack with its own planet-scale infrastructure, custom-built chips, generative AI models and development platform, as well as AI-powered applications, to help organizations transform. Customers in more than 200 countries and territories turn to Google Cloud as their trusted technology partner.

## Deloitte

Forge the future. Unite Deloitte's deep industry insights with the transformative power of Google Cloud. Together, we'll craft tailored strategies to address your unique challenges and aspirations to spark progress for your enterprise. Unlock the Google Cloud ecosystem and create bold new possibilities, transform operations, and capitalize on measurable results. With Deloitte and Google Cloud, you won't adapt to change. You'll shape it. Forging a future that's sustainable, impactful, and uniquely yours. www.deloitte.com/googlecloud

# Disclaimer

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited (DTTL), its global network of member firms, and their related entities (collectively, the “Deloitte organization”). DTTL (also referred to as “Deloitte Global”) and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/about to learn more.

This communication contains general information only, and none of Deloitte Touche Tohmatsu Limited (“DTTL”), its global network of member firms or their related entities (collectively, the “Deloitte organization”) is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser. No representations, warranties or undertakings (express or implied) are given as to the accuracy or completeness of the information in this communication, and none of DTTL, its member firms, related entities, employees or agents shall be liable or responsible for any loss or damage whatsoever arising directly or indirectly in connection with any person relying on this communication. DTTL and each of its member firms, and their related entities, are legally separate and independent entities.