Google Cloud
Security

Secure AI innovation without interruption

A guide to rapid AI adoption

Abstract: AI is transforming how organizations operate, but innovation can't scale without security. This guide shows how Google Cloud helps teams build, deploy, and govern AI responsibly — with visibility, control, and trust built into every layer, so innovation accelerates instead of slowing down.

The promise and paradox of AI innovation

AI is reshaping industries at remarkable speed, unlocking new efficiencies and new revenue opportunities. But rapid adoption also brings new risks. Unsupervised model development, agentic workflows, and shadow AI can create blind spots that expand faster than traditional security can keep up with. Leaders want the benefits of AI, but they need confidence that their data, models, and users remain protected as they scale.

That's why trust has become the foundation of every AI strategy. Security must move from a limiting factor to an accelerant — and a major competitive advantage — that makes responsible innovation possible.

54% of organizations use public AI like GPT-4, Gemini, or Claude

52% say sensitive data exposure is their top security risk

27% are confident they can secure AI used in core business operations

Source: CSA Report: The State of AI Security and Governance 2025

> ***As organizations move from experimentation to operational deployment, strong security and mature governance are the key differentiators for AI adoption."***
>
> —Dr. Anton Chuvakin, Security Advisor at Office of the CISO, Google Cloud

Google's approach: Embed protection across the AI life cycle

Google's approach embeds protection into every layer of the AI life cycle through the Secure AI Framework (SAIF) and AI Protection capabilities in Security Command Center (SCC), combining visibility into your AI assets, AI risk assessment, model and agent guardrails, and real-time threat detection so organizations can move faster with confidence.

> *The entire AI supply chain introduces new risks... We see the entire life cycle as a new attack surface. Across every stage, inadequate security controls, especially robust identity management, leave organisations vulnerable to data poisoning, model theft, and tampering attacks."*
>
> —Sai Sirish Reddy Kommareddy, Security Engineer, Yahoo

How Google secures every layer

Google Cloud Security

Infrastructure

Secure-by-design cloud foundation with hardened compute, network isolation, confidential computing, and built-in governance from the ground up.

Data

Sovereign controls, sensitive data protection, and cryptographic transparency to protect data wherever it lives — and control who can access it.

Models

Runtime guardrails and protections to prevent prompt injection, data leakage, and unsafe outputs across first- and third-party models.

Platform

Secure orchestration of AI services, tools, and workflows with policy enforcement, monitoring, and posture management built in.

Agents

Visibility and governance over agent behavior, permissions, and interactions — ensuring autonomous systems act safely and as intended.

Key takeaway: Innovation can't move faster than trust — and trust begins with security.

Watch the full session:

Protect your AI innovation with Google Cloud IAM

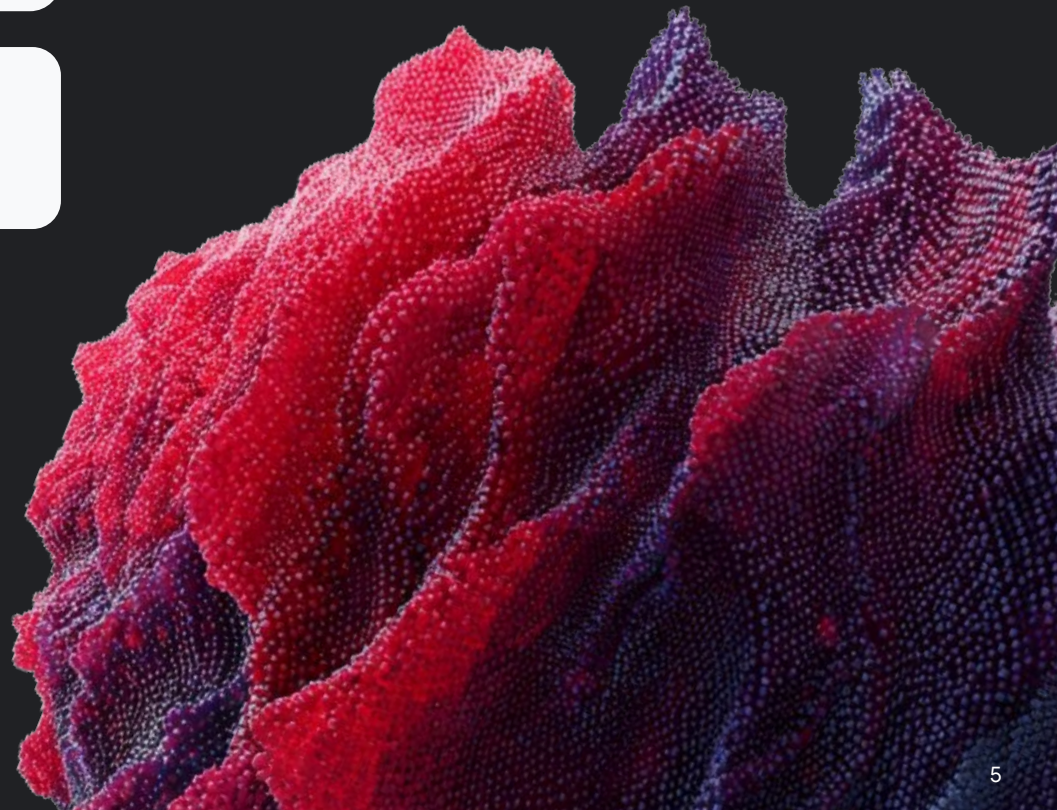

Securing the foundations: Identity and access governance

Every AI system, no matter how advanced, depends on strong identity and access control. As organizations adopt more AI workloads and autonomous agents, the number of non-human identities grows fast. Without clear governance over who and what can access sensitive data, innovation quickly becomes exposure.

Google Cloud IAM: Built for the age of AI

Google Cloud IAM provides that foundation. With fine-grained control for both human and machine identities. For AI agent access management, Google Cloud IAM provides an agent identity principal that better addresses agentic risks compared to traditional service accounts. Automated policy intelligence highlights overly broad permissions and recommends least-privilege fixes, ensuring the right users (and the right agents) have access to only what they need, and nothing more.

> *AI agents are here. They are already on the job and increasingly getting entrenched into enterprise workforce... A large part of [agentic AI value] stems from agents' ability to take actions autonomously... And that can also result in unintended and harmful outcomes which current systems are not yet equipped to mitigate..."*
>
> —Aniket Patankar, Lead Product Manager, Cloud AI Security

Defense-in-depth for access management

Google Cloud Security

IAM Allow and PAM Grant

Allow (or grant JIT) access permission on specific resources.

IAM Deny

Prevent members from using permissions, regardless of the roles they're granted.

Principal Access Boundary

Restrict resources on which members can act on.

What strong IAM enables

- Clear visibility into who and what is accessing data
- Automated enforcement of least-privilege access
- Continuous governance as AI agent identity and access

Key takeaway: Confident IAM fuels faster AI innovation, because you can't secure AI without knowing exactly who and what has access.

Watch the full session:

[Protect your AI innovation with Google Cloud IAM]

> *We provide you with a full spectrum of human identity capabilities... and for non-human identities, we enable you to securely have your workloads talk to Google Cloud APIs... What this means is that your nonhuman identities can be secured, and you don't have to worry about impersonation or dormant identities lying around that can be misused by an attacker."*
>
> —Ravi Shah, Group Product Manager, Google Cloud

Ready to build this foundation? Read our technical guide:
[Building a Secure Foundation to Protect AI Workloads in Google Cloud]

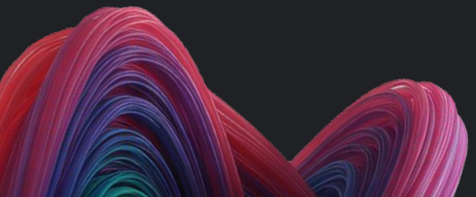

Securing the AI life cycle: From discovery to defense

AI innovation is dynamic — and oftentimes messy and non-linear. Models evolve, datasets grow, and agents interact in new ways. To keep pace, security must evolve continuously across the entire AI life cycle — from discovery to deployment to defense.

> *Organizations lack an ability to discover and visualise agentic AI interaction chains... and proactively monitor for emerging urgent systemic risks, for example, cascading failures from multi-agent systems."*
>
> — Aniket Patankar, Lead Product Manager, Cloud AI Security

Lack of visibility and governance of agent ecosystem

- Agent, MCP sprawl
- Shadow AI
- Cascading failures in multi-agent systems

Compromised agent identities and excessive permissions

- Excessive agency
- Agent impersonation
- Misuse of authority

Unsafe agent execution and tool misuse

- Unintended actions
- Resource overload
- Agent misalignment

Data and model integrity violations

- Data poisoning
- Prompt injection
- Model compromise

AI Protection in Security Command Center: End-to-end AI security

AI Protection in Security Command Center (SCC) provides visibility across every model, dataset, and AI agent running in your Google Cloud environment. This gives teams a real-time inventory of AI assets, including the shadow AI systems that often slip through the cracks.

The pillars of Google Cloud's AI Protection (AIP)

Discover

Automatically inventory AI models, datasets, and agents — so you see what's actually running, not just what's documented.

Secure

Apply runtime guardrails with Model Armor to block prompt injection, prevent data exfiltration, and filter unsafe outputs.

Manage

Detect, investigate, and respond to AI threats using integrated Mandiant and Google Threat Intelligence.

Key takeaway: AI moves fast, and visibility and protection must move with it.

Watch the full session:

Unsupervised autonomy: How to secure AI agents and limit risk

Explore the future of AI-driven threats in the Cybersecurity Forecast 2026.

Governing agents and autonomous systems

As organizations adopt agentic AI — systems that reason, plan, and act with increasing autonomy — governance becomes essential. These agents are powerful, but without clear controls, they can introduce new forms of operational and security risk.

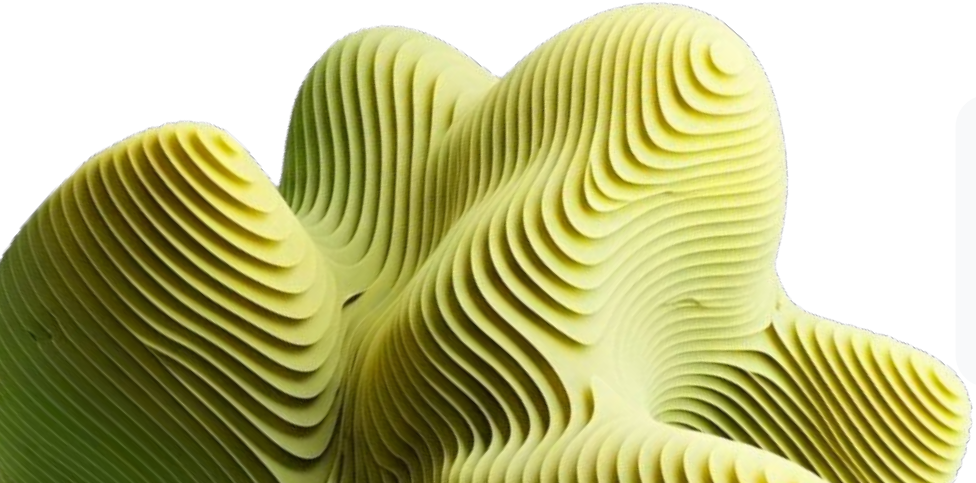

Visibility into the entire agentic ecosystem

AI Protection in Security Command Center extends beyond model protection to secure the entire agentic ecosystem — mapping each agent, the tools they can access, and interaction paths between systems. Intuitive visualizations show how agents behave and potential risks they might run into.

The outcomes of agent governance:

- Full visibility into the agent ecosystem
- Safe delegation of tasks to autonomous systems
- Compliance with emerging AI regulations

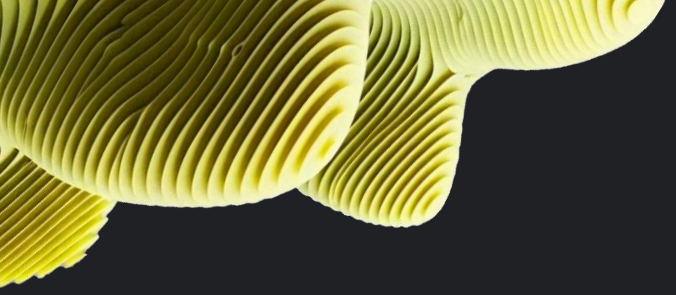

> *AI protection empowers customers to build and deploy AI in a secure, compliant, and private manner... and helps customers manage risks holistically across the entire AI lifecycle."*
>
> —Aniket Patankar, Lead Product Manager, Cloud AI Security

Key takeaway: Governance builds confidence, because autonomy without oversight is risk.

Watch the full session:

Unsupervised autonomy: How to secure AI agents and limit risk

Learn how governance accelerates innovation in new research covered in "The State of AI Security and Governance."

Building trust through transparency and sovereignty

In the AI era, where and how data is processed matters as much as what that data is powering. Organizations — particularly those in regulated industries — need assurance that sensitive data stays within defined boundaries and is protected by clear, enforceable controls.

Sovereign Cloud from Google: Control, transparency, and compliance at scale

Google's Sovereign Cloud helps organizations define and enforce their own sovereignty model. With Data Boundary, they get full visibility and control into where data is stored and who can access it. Capabilities like Cloud External Key Manager (Cloud EKM) and Key Access Justifications ensure customers maintain cryptographic control — even from Google — while meeting evolving regional regulations. The result: Organizations can adopt AI confidently without slowing down innovation.

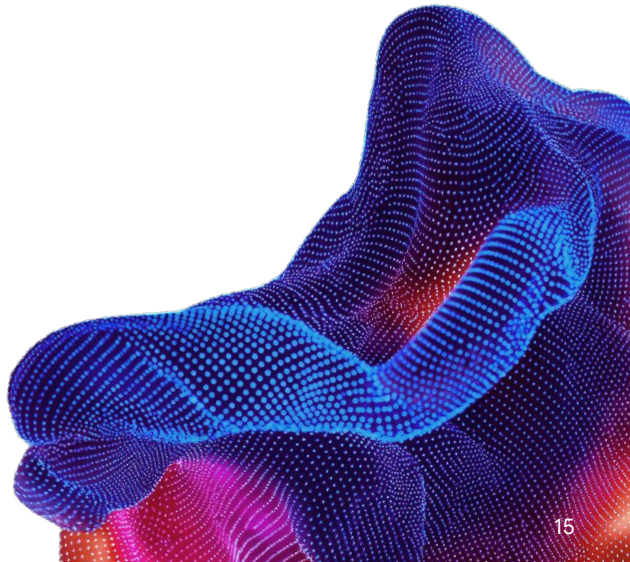

Global demand for sovereign solutions is growing

1

Regulations

The need for more digital and physical control over organizational data and operations.

2

Geopolitical

Changing dynamics and planning for new scenarios.

3

Faster Innovation

Developing capabilities in-house can be time consuming and expensive.

The outcomes of sovereignty:

- Regionally compliant AI operations
- Clear, provable data residency and access controls
- Customer-controlled encryption and key management
- Faster innovation without regulatory risk

Key takeaway:

You can't restrict your way to trust. Trust is earned through transparency and control.

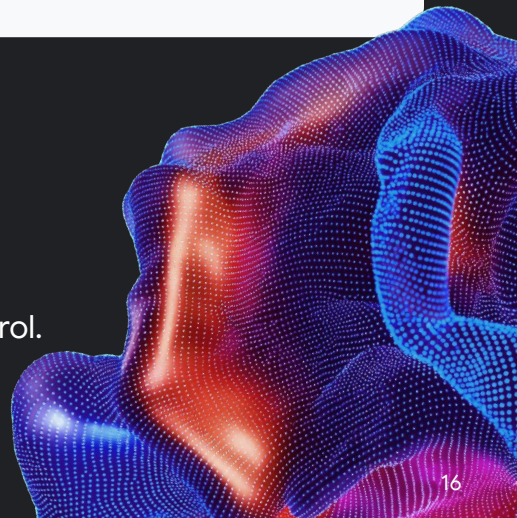

Watch the full session:

Sovereign Cloud from Google: The power of choice

> *Nations around the world are implementing new rules that demand more digital and physical control over how organizational data is stored, processed, and managed... The world is changing quickly, and organizations need to plan for new scenarios."*
>
> — Matt Garr, Group Product Manager, Google Cloud

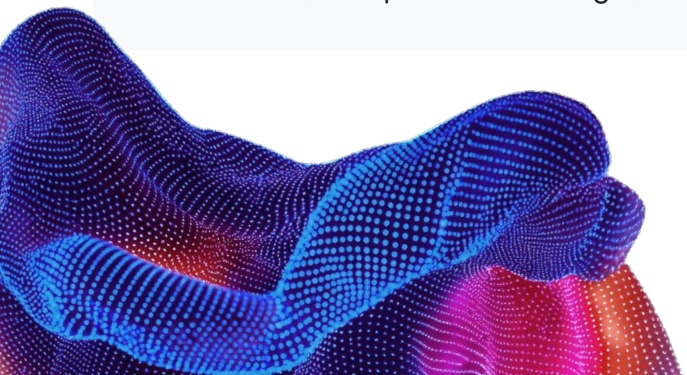

> *Our customers are telling us that they need to innovate faster, not get bogged down in building and managing complex infrastructure. They need to leverage the power of the hyperscale cloud, but in a new way that meets these new stringent requirements. At Google Cloud, we believe the answer is not a one size fits all approach. That's why we've built the industry's most comprehensive portfolio of sovereign solutions designed to give you the power of choice.*
>
> —Matt Garr, Group Product Manager, Google Cloud

Real-world outcomes: Securing AI at scale

Organizations across industries are already using AI Protection capabilities in Security Command Center and Sovereign Cloud to scale AI safely and confidently. The results are clear: more visibility into AI assets, faster threat detection, and easier compliance — all while accelerating innovation.

Case study: **Deutsche Telekom**	
Problem	Needed to run AI at scale while meeting strict EU data-residency and sovereignty requirements.
Solution	Built compliant AI workloads on Sovereign Cloud from Google with EKM, Access Justifications, and regional oversight.
Results	• Accelerated AI adoption without regulatory blockers • Stronger customer trust through transparent sovereignty • Simplified audit readiness across markets
Key quote	*“For my fellow data architects in regulated industries, you don’t have to choose between innovation and compliance. With the right technical approach, you can achieve both and build platforms that position your organization for the AI-driven future that’s rapidly approaching.”* — Ashutosh Mishra, VP, Data Architecture & Governance, Deutsche Telekom Read the full case study

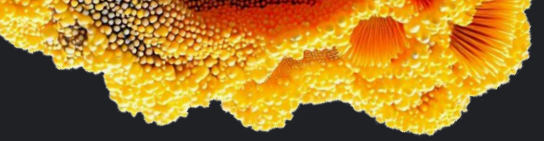

Case study: **Snap**	
Problem	Rapid AI adoption created “shadow AI” risks and limited visibility.
Solution	Implemented AI Protection capabilities in Security Command Center and Model Armor to discover models, enforce guardrails, and prevent unsafe prompts.
Results	• Full visibility into AI model use and risk • Reduced exposure from unmanaged AI tools • Faster, safer innovation with “guardrails, not gates”
Key quote	*“We are early into journey to adopt AI securely. My team’s core ethos is to be proactive enabler for business. We believe in providing guardrails instead of gates and participating in security solutions rather than operating through decree.”* — Shrikant Pandhare, Head of Infrastructure Security, Snap

Key takeaway:

When teams gain clear visibility and control over their AI workloads, innovation accelerates while risk moves out of the way.

Watch the full session:

Sovereign Cloud from Google: The power of choice

Moving your AI innovation forward

AI has moved from experimental curiosity to essential strategic priority. Executive and board commitment to deploying AI is high and only rising, and virtually every organization is moving quickly to embed AI into their core business and security workflows. But that momentum presents a new challenge: According to the Cloud Security Alliance, 3 in 4 organizations (72%) aren't confident in their ability to secure AI — even as they aggressively accelerate AI initiatives.

This gap between enthusiasm and execution is now the biggest barrier to scaling AI successfully and responsibly. Organizations that close this gap (by gaining visibility and strengthening governance and security across their AI stack) will move faster, reduce risk, and build a dramatic competitive advantage around their ability to turn AI ambition into measurable and sustainable business value.

Security can be a business enabler. Read how early adopters of AI-driven security are achieving measurable ROI and accelerating their time to market in "The ROI of AI in security."

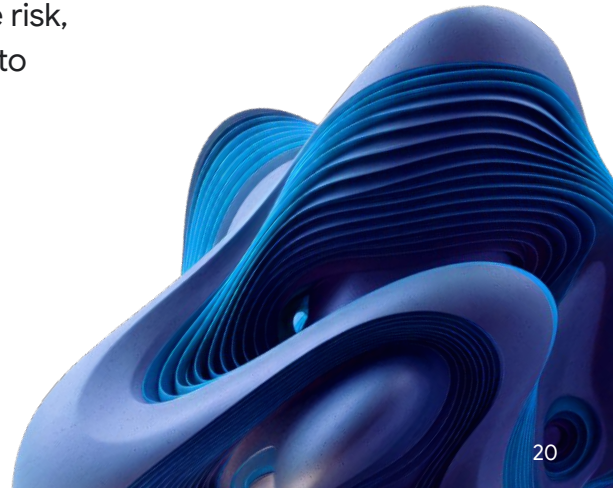

Ready to get started?

Google Cloud provides the frameworks, controls, and expertise needed to innovate responsibly from model to market. Whether you're evaluating risk, building governance, or deploying agents, Google offers a clear path forward.

Secure AI Framework (SAIF)	AI Protection in Security Command Center	Mandiant Consulting
Best practices for responsible AI adoption.	Comprehensive governance and threat defense.	Expert-led threat readiness and AI-specific assessments.

Read the companion guide
Reinventing the SOC with agentic AI

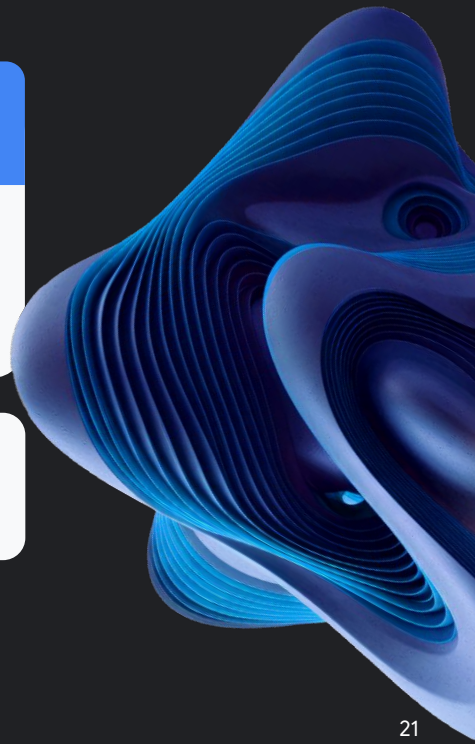

Google Cloud

Security