Security Accelerator for Gemini Enterprise

Turn security from a roadblock into an AI accelerator

Simplify AI governance, harden your security, and proactively address AI-specific security risks with Mandiant experts

Benefits

- **Risk reduction:** Evaluate governance and align architecture to industry frameworks
- **Validated guardrails:** Test against OWASP Top 10 for LLMs to mitigate jailbreaking and prompt injection.
- **SOC readiness:** Prepare defenders with AI-specific playbooks and MITRE ATLAS-aligned offensive exercises.
- **Improved defenses:** Enforce secure configurations via threat-led STRIDE assessments of VPC service controls and encryption.

The challenge

Enterprise AI deployments can span a full spectrum: from the Gemini Enterprise Agent Platform for developers, to the Gemini Enterprise app for everyday workflows, and industry-tailored solutions like Gemini Enterprise for Customer Experience. Beyond the architecture, these global AI initiatives often require securing alignment across a web of diverse stakeholders—from legal and compliance to product, engineering, and security. Compounding this complexity is the urgent need to secure this new, expanding attack surface against sophisticated, AI-specific threats like prompt injection and data poisoning.

A comprehensive approach to securing and scaling your AI initiatives

Mandiant provides comprehensive security accelerators aligned to the current stage of your Gemini Enterprise rollout. We work with you to establish a secure operations foundation, validate your controls and scale your workflows, while empowering your engineering and security teams to deploy AI at scale with total confidence. The Mandiant Retainer provides flexible access to these services via pre-paid funds.

Services Overview

Aligned with industry frameworks

The Security Accelerator for Gemini Enterprise is a set of services built on Mandiant's frontline intelligence and industry-standard frameworks like NIST AI RMF, ISO 42001, and Google's Secure AI Framework (SAIF).

Mandiant experts work to strengthen your AI governance, threat model your Gemini Enterprise deployment, and prepare your security operations center (SOC) for AI-driven data exposure. Explore our key services areas aligned with the different stages of your Gemini deployment:

Area 1: Gemini Enterprise Governance Service - Pre-Gemini deployment

Mandiant Retainer - starting at 42 units

Align your key stakeholders to eliminate policy blind spots and establish clear data protection guardrails before your Gemini Enterprise deployment goes live. Mandiant maps your planned deployment against industry frameworks, to document actionable steps that accelerate AI adoption while keeping your sensitive corporate data secure.

Service Methodology

Mandiant experts facilitate a structured interactive workshop with your Risk, Compliance, and Security stakeholders to map your intended Gemini deployment against industry standards (NIST AI RMF, ISO 42001, Google SAIF). This collaboration establishes critical data handling guardrails, utilizing Mandiant's data tiering methodology to define exactly what corporate data is safe for Gemini to ingest or index. Next, Mandiant conducts a targeted gap analysis to identify where current corporate policies fail to address the unique nuances of generative AI.

Deliverables

- Customized "traffic light" acceptable use policies (AUPs): Clear guidelines defining what AI usage is green-lit, requires approval, or is strictly prohibited.
- AI governance gap analysis: A documented review of your planned Gemini Enterprise deployment against Google SAIF and other industry frameworks to identify critical policy blind spots with recommendations on how to address them.

Area 2: Security Assessment and Threat Modeling - Moving from Gemini pilot to production

Mandiant Retainer - starting at 49 units

This targeted technical architecture review provides you with the confidence that your Gemini Enterprise security controls are correctly configured to support your AI initiatives before you go live. By proactively identifying and addressing complex risks like privilege escalation through agents and data leakage, you empower your engineering teams to deploy Gemini Enterprise securely and at scale.

Service Methodology

Mandiant will perform a targeted technical validation of your key Gemini Enterprise components and a facilitated threat modeling session limited to one (1) Gemini use case and up to three (3) data stores.

Mandiant experts collaborate with your engineers to analyze the architecture and data flows of your specific Gemini Enterprise use case. By mapping trust boundaries, we conduct a targeted threat modeling session using industry-leading frameworks and frontline intelligence to identify unique risks within your AI workflows. Following this, Mandiant performs a focused review of your underlying Google Cloud infrastructure to ensure essential controls, such as the following are implemented:

- Data leakage prevention through security perimeters (VPC-SC)
- IAM access controls
- Logging
- Prompt screening protection through Model Armor
- Organizational policies relevant to your Gemini Enterprise use case

Mandiant reviews data stores to ensure access scopes in Gemini Enterprise (utilized by agents) align with least-privilege access and enterprise requirements. Finally, Mandiant conducts a threat modeling exercise to evaluate your Gemini Enterprise architecture against spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege.

Deliverables

- Threat model matrix and configuration analysis: A detailed report of threat tiering paired with precise, step-by-step configuration guidance to improve your Gemini Enterprise security posture.

Area 3: Incident Response Readiness- Moving from Gemini pilot to production

Mandiant Retainer - starting at 37 units

Transform your SOC from reactive to ready. Ensure your defenders can confidently detect and respond to the speed and scale of AI-driven threats. We stress-test your incident response procedures against real-world scenarios, giving you the actionable playbooks needed to isolate compromised AI agents without disrupting your broader Google Workspace environment.

Service Methodology

To ensure your defenders are prepared for AI-specific threats, Mandiant experts facilitate a targeted two-hour Tabletop Exercise or Playbook Workshop. Using MITRE ATLAS-aligned scenarios, we simulate threat scenarios targeting Gemini such as prompt injection, jailbreaking and attempts to bypass Model Armor, to evaluate your current detection and response capabilities and identify critical logging blind spots. Mandiant stress-tests your workflows to confirm analysts can isolate threats swiftly without impacting the wider Google Workspace tenant.

Deliverables

- One (1) AI incident response playbook tailored for your organization
- OR one (1) Table top after-action-report (AAR) detailing SOC visibility gaps to help you immediately mature your response posture.

Area 4: Technical Validation and Offensive Security - Scaling Gemini enterprise-wide

Mandiant Retainer - starting at 31 units

Validate your defenses work in the real world. This rigorous, time-boxed offensive engagement provides the ultimate validation that your Gemini guardrails can withstand sophisticated, real-world attacks before full production rollout. By proactively assessing vulnerabilities like complex prompt injection or unauthorized data access via connectors, you receive the precise technical remediation steps needed to confidently secure your AI environment.

Service Methodology

To validate your security guardrails against advanced adversaries, Mandiant experts execute a targeted, two-week Red Team engagement against your fully provisioned Gemini environment.

Using the OWASP top 10 for LLM applications and MITRE ATLAS frameworks, Mandiant maps the exposed attack surface (e.g. chat interfaces, embedded Workspace extensions, API endpoints for custom agents) and systematically attempts to bypass system instructions and Model Armor via complex prompt injection techniques.

This assessment tests your data connectors for unauthorized pivot paths and attempts to bypass data loss prevention (DLP) controls by forcing the model to exfiltrate sensitive data.

Deliverables

- **Red Team technical report:** A detailed technical breakdown outlining findings, prompts used, OWASP mapping, telemetry triggered and the specific configuration remediations required to lock down your environment.

Measurable AI Security Outcomes

The Security Accelerator for Gemini Enterprise goes beyond theoretical advice. Mandiant's comprehensive approach hardens your environment against the most critical AI failure points, delivering tangible protection against:

- **Unauthorized Data Access:** We validate Workforce Identity Federation (WIF) and ACL mapping to ensure Gemini respects source permissions (like SharePoint or Salesforce), preventing internal information disclosure.
- **Agent and Credential Hijacking:** We audit Agent Designer configurations and Model Armor policies to block prompt injection attacks that attempt to extract hardcoded API keys or manipulate agents into "Confused Deputy" actions.
- **Sensitive Data Sprawl:** We implement and test Chrome Enterprise Premium DLP and Cloud Sensitive Data Protection (SDP) to automatically mask PII/PHI and block unauthorized data exfiltration to personal accounts.

Ready to secure your AI strategy with expert guidance? Contact us today to discuss how we can help you accelerate your Gemini Enterprise journey and build your AI future on a secure foundation.