AVIATRIX

SOLUTION BRIEF

Contain Gemini Enterprise Agent Platform

Powered by Aviatrix + Google Cloud

Aviatrix Validated Containment Architectures are lab-tested containment deployment blueprints for the AI platforms enterprises are actually running. Ship-ready, policy-included, validated before they arrive. This Validated Containment Architecture covers Gemini Enterprise Agent Platform.

The Threat

Gemini Enterprise Agent Platform runs production AI agents in a Google-managed tenant project. By default, those agents have unrestricted outbound internet access. Their tool calls, MCP connections, and any calls to non-Gemini model providers are indistinguishable on the wire from legitimate work. There is no inline enforcement point the customer's security stack can observe.

When prevention fails and detection is too slow, containment decides whether the incident becomes a catastrophic breach. Two threat vectors follow directly from the default-allow posture:

- A prompt injection or compromised dependency redirects an agent tool call to attacker-controlled infrastructure. No Google-native control stops the socket from opening to a non-Google host.
- An agent reaches an unsanctioned model provider, such as OpenAI or Anthropic, carrying sensitive data retrieved from sanctioned RAG sources. VPC Service Controls is blind to non-Google destinations. The call looks identical to legitimate work on the wire.

AI workloads are the bullseye of the modern threat model. They are ephemeral, highly privileged, and rapidly shipped, and the default cloud posture is open. Closing that default-allow posture is the operator's responsibility, not the platform's.

THE CASCADE, MARCH 2026

A coordinated supply chain operation compromised LiteLLM and four other AI infrastructure packages in twelve days, using trusted update channels to harvest credentials from AI development environments. The attack executed as legitimate code through trusted pipelines. Detection saw nothing anomalous. Aviatrix Distributed Cloud Firewall customers in default-deny mode blocked the command-and-control egress at the network layer before credentials left the environment. This Validated Containment Architecture pre-assembles that posture for Gemini Enterprise Agent Platform.

The Architecture

The Aviatrix Validated Containment Architecture for Gemini Enterprise Agent Platform inserts transparently at Layer 3, between the agent and the internet, with no changes to the agent code or runtime. All non-Google agent egress leaves the Google-managed tenant project through a Private Service Connect interface (PSC-I) network attachment into a customer VPC, where the Aviatrix gateway enforces default-deny containment policy.

The architecture ships in two deployment shapes, both governed by one Distributed Cloud Firewall policy model:

- **Managed runtime shape:** Agent Platform's managed runtime fronted by PSC-I, with Aviatrix inserted transparently at Layer 3. No HTTPS_PROXY, no SDK change, no agent redeploy.
- **GKE shape:** GKE-hosted custom ADK runtime, with Aviatrix Kubernetes SmartGroups for pod-level identity and selective transparent TLS decryption for URL-path enforcement. Requires Controller 9.0+.

How it works in practice

Component	What It Does
PSC-I network attachment	Routes all non-Google agent egress out of the Google-managed tenant project into the customer VPC. The RFC 1918 egress next-hop Google requires under VPC Service Controls, satisfied at the route layer.
Aviatrix gateway (transparent L3 insertion)	Performs NAT and acts as a transparent forward proxy. The agent makes the same outbound calls it always made. There is no HTTPS_PROXY and nothing on the agent changes.
Default-deny DCF policy	Every agent starts denied. Only explicitly approved tool destinations, MCP servers, and RAG endpoints are reachable via WebGroup allow-lists. Everything else is blocked and logged in CoPilot FlowIQ.
vca-gemini-enterprise-agents-shadow-model-deny rule	Placed ahead of the tool allow-list. Denies and logs every attempt to reach api.openai.com, *.anthropic.com, api.mistral.ai, *.perplexity.ai, or any other unsanctioned model provider. VPC Service Controls does not see these destinations. This rule does.
SmartGroup East-West isolation	Deny rules between the agent spoke and every other spoke in the fabric. A compromised agent cannot reach adjacent workloads. Blast Radius stops at the agent VPC.
UDP/53 deny	Blocks DNS-tunneled exfiltration to external resolvers.
Selective TLS decryption (GKE shape)	Scoped tightly to a destination FQDN SmartGroup. Enables URL-path filtering to block compromised supply-chain paths on sanctioned hosts while leaving all other traffic encrypted end to end.

Three Things Your Current Stack Can't Do

Google's native controls each address a different axis. None of them is the network egress firewall for non-Google destinations across your whole estate.

01 Block a shadow-model call at the network

VPC Service Controls stops Google-to-Google exfiltration. It is blind to an agent calling api.openai.com. This Validated Containment Architecture ships a named shadow-model deny rule that sees those calls and blocks them before the socket opens.

02 Enforce egress policy across every cloud your agents run on

Secure Web Proxy does FQDN allow-listing, but it is GCP-only and per-region, with its own console and policy language. This Validated Containment Architecture uses the same Distributed Cloud Firewall policy model that governs your Bedrock AgentCore, Azure AI Foundry, and self-hosted Kubernetes agents. One control plane. One policy model. One audit log.

03 Contain lateral movement from a compromised agent

No Google-native control isolates the agent spoke from adjacent workloads in the customer fabric. SmartGroup East-West deny rules in this Validated Containment Architecture ensure that a compromised agent cannot reach anything beyond its own VPC. The Blast Radius is bounded by architecture.

Compliance Evidence

For SOC 2, HIPAA, PCI-DSS, and FedRAMP environments, auditors require architectural proof of enforcement, not a policy document. CoPilot FlowIQ per-connection logs are the continuous audit evidence that the control is operating.

The same policy model and log format apply whether the agent runs on Gemini Enterprise Agent Platform, Bedrock AgentCore, Azure AI Foundry, or your own cluster. Auditors see one control, not one per platform.

Evidence Artifact	What It Proves
CoPilot FlowIQ per-connection logs	Continuous audit trail of every egress decision, attributed to SmartGroup, WebGroup, and DCF rule with timestamp. Maps to SOC 2 CC6.6/CC6.7, HIPAA §164.312(e)(1), PCI-DSS v4 Req 1.3, FedRAMP SC-7.
vca-gemini-enterprise-agents-shadow-model-deny block logs	Named, timestamped log of every attempted reach to unsanctioned model providers. Proves the control is operating even when no breach occurs.
East-West SmartGroup deny logs	Proves lateral movement from the agent spoke to adjacent workloads was architecturally prevented.
DCF policy-as-code in Terraform	Version-controlled, peer-reviewed policy with an approval trail. Adding a tool endpoint is a pull request, not a change ticket to another team.

Cross-platform log consistency	One policy model and log format across every agent platform. Auditors see one control, not one per platform.

The proof of enforcement is the architecture itself, not a statement about the architecture.

Questions Worth Asking

- If an agent in your Gemini Enterprise Agent Platform called api.openai.com right now, would your security team know? How quickly?
- If a prompt injection redirected an agent tool call to an attacker endpoint, what, architecturally, would prevent the socket from opening?
- If a compromised agent spoke tried to reach adjacent workloads in your cloud fabric, what limits the Blast Radius?
- When your compliance team asks for continuous evidence that AI egress is governed, what do you show them?
- If your agents run on more than one platform (Gemini, Bedrock, Foundry, self-hosted), how many policy engines, audit trails, and enforcement consoles does that require today?

What's Included

Deliverable	Detail
Insertion pattern	Architecture diagram and configuration notes showing exactly where Aviatrix enforcement goes in the Gemini Enterprise Agent Platform topology, for both the managed runtime and GKE shapes.
SmartGroup model	Tagging taxonomy keyed to Gemini Agent Platform's identity primitives: CIDR SmartGroup over the PSC-I network-attachment subnet, Kubernetes SmartGroup matching k8s_namespace=agents, and CIDR SmartGroup over the Gemini Agent Platform PSC endpoint.
Baseline DCF policy pack	Default-deny DCF rules pre-scoped to the destinations Gemini Enterprise Agents actually need: sanctioned tool WebGroups, the Gemini Agent Platform PSC endpoint, a named shadow-model deny, UDP/53 deny, and East-West isolation rules.
Bill of Materials	Gateway count and instance sizing, Controller version requirements, GCP prerequisites (PSC-I, Cloud DNS, Flow Logs), and deployment ordering.
GitHub repository	Full Terraform at aviatrix-blueprints/blueprints/gemini-enterprise-agents/. Includes a sample ADK agent and scenario probes. DCF policy lives in the same repo as the agent config. Clone it, configure it, deploy it.

Get Started

This Validated Containment Architecture is available now. Aviatrix Distributed Cloud Firewall customers on Controller 8.1 or later can deploy the SNI and domain baseline today. Selective TLS decryption, URL-path filtering, and egress IDS/IPS require Controller 9.0.

Request a 30-minute architecture review.

We walk through the policy model in your environment, identify the Gemini Enterprise AI agents you have running today, and show you the egress paths you currently cannot see.

About Aviatrix

Aviatrix® is pioneering the Cloud Native Security Fabric – the architecture the Containment Era requires. The Cloud Native Security Fabric governs every workload communication path across every cloud, every VPC, every Kubernetes cluster, and every serverless function, from a single policy plane. One rule. Universal propagation. Enforced at the workload, not at a chokepoint. Trusted by more than 500 of the world's leading enterprises. For more information, visit aviatrix.ai.