

Agentic Infrastructure for the Autonomous Enterprise on Google Cloud

This course provides a technical guide to enable Solution Architects to shift from building isolated chatbots to deploying persistent, Gemini Enterprise-enabled AI workers on Google Cloud. Participants will master agentic memory design, API-driven tool orchestration, and infrastructure governance using the Google Cloud Agent Platform—including the Agent Platform Reasoning Engine for persistent state management and Agent Extensions for departmental integration. Learners will move beyond "Instructional Hope" to technical enforcement, building the "Paved Road" required to orchestrate multi-agent fleets and secure non-human identities.

 **DURATION**
2-3 hours

 **LEVEL**
200 - Intermediate

 **FORMAT**
Instructor-led

 **SCALABLE**
1-50, scaled to 1000

What you'll learn

- Diagnose Agency Gaps: Strip out "Track A" conversational habits and map structural frictions.
- Architect the Engine: Deploy persistent, context-aware reasoning loops inside the Build Tier.
- Secure the Perimeter: Enforce Zero Trust patterns and runtime protections via the Govern Tier.
- Scale the Fleet: Mitigate reasoning drift and measure ROI using the Optimize & Scale Tiers.

Overview	5 modules
Who this course is for	- Enterprise Architects: Who need to build a scalable, secure backbone for persistent, Gemini Enterprise-enabled AI workers. - Systems Integrators: Tasked with connecting AI to legacy ERP/CRM data via secure, cross-departmental tool orchestration using the Google Cloud Agent Platform. - IT Directors: Who need to move from "Project Thinking" to "Platform Thinking" to manage a global fleet of autonomous agents while maintaining governance and scalability.
Products	- Agent Platform - Model Armor - AlloyDB AI - Access Context Manager - AI Perimeter Policy - Agent Studio on Gemini Enterprise Agent Platform - AlloyDB for PostgreSQL - VPC Service Controls
Prerequisites	- Completion of Organizing for AI Success. - Foundational knowledge of Google Cloud (VPC, IAM, Cloud Run). - Helpful to be familiar with Python basics. - Helpful to be familiar with REST API structures. - Helpful to be familiar with RAG concepts.
Not covered	- Basic Prompt Engineering: This is an architecture course, not a "prompt hacks" class. - Model Fine-Tuning: We focus on leveraging frontier models through high-leverage infrastructure. - Standard UI Chatbots: We focus strictly on the "Track B" autonomous worker.

Module 01

The Shift to Agency

Topics	• The Pillars & Governance • The Architectural Friction Forces • The Autonomous Maturity Scale
Objectives	• Analyze the "Agency Gap" by diagnosing the functional intersection of Reasoning, Memory, and Tools with Data Governance to move from Track A (Chat) to Track B (Workers). • Diagnose the four technical frictions (Integration, Statelessness, Latency, and Governance) that prevent AI pilots from scaling into production. • Evaluate infrastructure readiness using the L1–L5 Maturity Scale to prioritize "Paved Road" investments for Level 4+ autonomy.
Activities	• 1 Use Case • 2 Case Studies • 1 Demo

Module 02

Building the "Paved Road"

Topics	• Reference Stack & Tool Archetypes • The Memory Decision Guide • Multi-Agent Orchestration Patterns • The Paved Road Lifecycle
Objectives	• Apply the Agent Platform SDK and Agent Platform Reasoning Engine to standardize agent deployment and manage persistent conversation state. • Evaluate the trade-offs between AlloyDB and Agent Search to select the optimal storage layer for metadata-heavy vs. high-scale agents. • Apply specific orchestration patterns (Hub-and-Spoke, Linear Relay, or Parallel Critic) to manage complex, multi-departmental goals. • Design an agentic deployment arc from Sandbox to Certified production to ensure infrastructure precedes autonomous action.
Activities	• 4 Demos

Module 03

The Autonomous Perimeter

Topics	• Threat Modeling for Agentic Systems • Identity Hierarchy & Credentials • Defending the Boundary: Model Armor • Responsible AI & Human-in-the-Loop
Objectives	• Apply agentic threat modeling to identify and mitigate risks like Indirect Prompt Injection and Tool-Chaining exploits. • Apply a three-layer identity model using Workload Identity Federation to ensure "Least Privilege" for autonomous workers. • Apply Model Armor as a real-time security proxy to filter malicious inputs and redact sensitive output data. • Analyze Responsible AI production requirements to embed accountability, traceability, and "Human-in-the-Loop" checkpoints within the Autonomous Perimeter.
Activities	• 1 Use Case • 4 Demos

Module 04

Sustaining Autonomy

Topics	• Infrastructure ROI • The GenAIOps Lifecycle • The Innovation Harvest
Objectives	• Analyze platform ROI by shifting from vanity metrics to Infrastructure Leverage Ratios and Component Reusability to prove the value of the "Paved Road." • Apply a continuous feedback loop using Golden Datasets and reasoning traces to detect and remediate "Reasoning Drift." • Apply the "Innovation Harvest" methodology to scale successful siloed tools into global, certified Gemini Enterprise assets.
Activities	• 1 Use Case • 1 Demo

Module 05

Summary and Quiz

Topics	• Review of Core Concepts
Objectives	• Evaluate understanding of core course concepts through scenario-based questions.
Activities	• 5 scenario-based multiple choice questions