Agentic AI on Google Cloud Platform

- 4 Days
- Lecture and Hands-on Labs

Course Overview

This course teaches engineers how to design, reason about, and implement agentic AI systems, with explicit application to Google Cloud environments. The first three days establish the core architectural foundations of agentic systems, covering planning, tool use, memory, orchestration, reliability, and governance. These concepts are then mapped to Google Cloud using industry-standard tooling, including Vertex AI, the Agent Development Kit (ADK), and LangChain. The fourth day is a hands-on build day, where participants implement a complete agentic workflow on Google Cloud, focusing on real integration points, failure handling, and deployment shape rather than abstract demos. By the end of the course, participants will be able to design and build agentic AI workflows that are realistic, operable, and aligned with Google Cloud execution.

Review this course online at https://www.alta3.com/courses/ai-google-agent

Who Should Attend

- Application Developers building AI-powered systems on Google Cloud
- Platform and Cloud Engineers supporting AI workloads
- ML Engineers and MLOps practitioners
- Technical Architects and Technical Leads

What You'll Learn

- Understand the architectural invariants of agentic AI systems.
- Design agents using tools and memory for planning and adaptation.
- Apply agent roles and orchestration patterns with ADK and LangChain.
- Reason about agent failure modes and reliability in cloud environments.

Outline

Day 1 - 3: Agentic AI Core

1. Foundations of Agentic Systems
2. Anatomy of an Agent
3. Tool Use and Action Design
4. Memory, State, and Retrieval
5. Multi-Agent Decomposition and Orchestration
6. Reliability, Observability, and Evaluation
7. Deployment Readiness and Governance on Google Cloud #### Day 4: Hands-on Build Day – Agentic AI on Google Cloud
8. Building an End-to-End Agentic Workflow
9. Memory, Failure Handling, and Control
10. Deployment Shape and Operational Readiness

Labs

- Decomposing a real problem into agent responsibilities
- Designing an agent interface and decision loop
- Modeling tool contracts and failure scenarios
- Designing a memory strategy for an agentic workflow
- Designing a multi-agent architecture for a complex task
- Defining metrics and traces for agent reliability
- Designing a deployment-ready agentic architecture
- Implementing a simple agentic workflow on Vertex AI
- Injecting failures and validating recovery behavior
- Reviewing and validating the final deployment architecture

Prerequisites

- Familiarity with software systems and APIs
- General understanding of AI or LLM concepts
- Prior Google Cloud experience is helpful but not required