

Course description

Build, deploy, secure and scale enterprise-grade AI agents using Google's agent development ecosystem. Explore the capabilities of Gemini models, Agent Studio, Agent Development Kit (ADK), and Google Cloud agent technologies to develop intelligent applications and automate enterprise workflows.

Learn how to design and optimize prompts using Agent Studio, build AI agents using ADK, integrate external tools and enterprise data through Model Context Protocol (MCP), implement multi-agent architectures using Agent-to-Agent (A2A) communication, and deploy production-ready agents on Google Cloud.

Gain practical experience in agent development, workflow orchestration, enterprise data integration, agent evaluation, security, governance and scalable AI agent architectures.

- Course level: Intermediate to Advanced
- Duration: 3 days (24 hours)
- Format: Instructor-Led Training (ILT)

Activities

This course includes instructor-led training, technical demonstrations, Agent Studio exercises, prompt engineering activities, hands-on Gemini application development, guided ADK agent development, MCP integration activities, RAG implementation exercises, multi-agent development workflows, agent deployment exercises, security and governance scenarios, enterprise architecture discussions, an enterprise capstone project, knowledge checks, and a final quiz to reinforce learning.

Course objectives

In this course, you will learn to:

- Explain the fundamentals of Generative AI and Agentic AI.
- Understand Google’s Gemini models and agent development ecosystem.
- Describe the architecture and capabilities of Agent Studio.
- Design, test and optimize prompts using Agent Studio.
- Create generative AI applications using Gemini models.
- Develop intelligent agents using Google’s Agent Development Kit (ADK).
- Configure agent instructions, tools, sessions and workflows.
- Integrate agents with external APIs and enterprise applications.
- Implement grounding and RAG for enterprise knowledge retrieval.
- Connect agents with external tools using Model Context Protocol (MCP).
- Develop multi-agent applications using ADK.
- Implement agent delegation and orchestration patterns.
- Understand Agent-to-Agent (A2A) communication.
- Deploy AI agents on Google Cloud.
- Configure authentication, authorization and agent security.
- Implement agent evaluation, monitoring and troubleshooting.
- Apply responsible AI practices and enterprise guardrails.
- Design scalable and production-ready enterprise AI agent solutions.

Key features

- Introduction to Generative AI and Agentic AI
- Google’s Gemini Models and Agent Ecosystem
- Agent Studio and Prompt Engineering
- Prompt Design and Optimization
- Multimodal Generative AI Applications
- Agent Development Kit (ADK)
- AI Agent Architecture and Development
- Tool Integration and Function Calling
- Enterprise Data Integration and Grounding
- Retrieval-Augmented Generation (RAG)
- Model Context Protocol (MCP)
- Agent-to-Agent (A2A) Communication
- Multi-Agent Architecture and Orchestration
- Production Agent Deployment on Google Cloud

- Agent Evaluation and Monitoring
- Enterprise AI Security and Governance
- Scalable Enterprise Agent Solutions

Intended audience

This course is intended for:

- Software Developers and Engineers
- AI/ML Engineers
- Generative AI Developers
- Cloud Engineers
- Google Cloud Architects
- Solutions Architects
- Enterprise Architects
- Application Developers
- Technical Consultants
- Platform Engineers
- DevOps Engineers
- Professionals building AI-powered applications
- Professionals interested in enterprise agent development

Prerequisites

We recommend that attendees of this course have:

- Basic understanding of Google Cloud Platform
- Familiarity with Generative AI and Large Language Models (LLMs)
- Basic Python programming knowledge
- Understanding of APIs and REST services
- Basic knowledge of cloud computing concepts
- Familiarity with AI application development concepts is beneficial
- Basic understanding of authentication and authorization is recommended

Products covered

- Gemini Models
- Agent Studio
- Gemini Enterprise Agent Platform
- Agent Development Kit (ADK)
- Agent Runtime
- Agent Engine

- Google Cloud
- Cloud Run
- Vertex AI
- Google Cloud IAM
- Model Context Protocol (MCP)
- Agent-to-Agent (A2A) Protocol
- Google Cloud APIs
- Google Cloud Logging
- Google Cloud Monitoring
- Enterprise Search and RAG Technologies

Course outline

Introduction to Generative AI and Google's Agent Ecosystem

- Generative AI and Large Language Models (LLMs)
- Agentic AI and Autonomous AI Agents
- Evolution from Chatbots to AI Agents
- Core Components of an AI Agent: Models, Instructions, Tools and Memory
- Reasoning, Planning and Decision-Making in AI Agents
- Google's Gemini Family of Models
- Google's AI Agent Ecosystem
- Gemini Enterprise Agent Platform
- Gemini Enterprise Agent Ready (GEAR) Programme
- AI Agent Development Lifecycle
- Enterprise Use Cases for AI Agents

Introduction to Agent Studio and Prompt Engineering

- Agent Studio Interface, Environment and Workspace
- Agent Studio Capabilities and Workflows
- Creating Generative AI Applications with Natural-Language Prompts
- Prompt Design Principles
- System Instructions and User Prompts
- Zero-Shot and Few-Shot Prompting
- Contextual and Instruction-Based Prompting
- Configuring Gemini Models and Parameters
- Refining, Managing and Evaluating Prompts
- Multimodal Prompting with Text, Images and Documents
- Grounding Prompts with External Knowledge Sources

- Generating Structured Outputs
- Testing and Comparing Model Responses
- Prototyping AI-Powered Applications

Building Generative AI Applications with Gemini

- Gemini Models and APIs
- Accessing Gemini Models through Google Cloud
- Integrating Gemini APIs into Applications
- Model Parameters and Generation Configurations
- Chat and Multi-Turn Conversations
- Managing Conversation History and Context
- Function Calling and Tool Integration
- Grounding and Retrieval-Augmented Generation (RAG)
- Connecting Gemini Applications with External Data Sources
- Contextual Question Answering
- Responsible AI and Safety Settings
- Application Integration Patterns

Developing AI Agents with Agent Development Kit (ADK)

- ADK Architecture and Core Components
- Setting Up the ADK Development Environment
- Creating and Configuring an AI Agent
- Integrating Gemini Models with ADK
- Agent Instructions and Capabilities
- LLM Agents and Workflow Agents
- Sequential, Parallel and Loop Agents
- Tools and Function Calling
- Custom Python Functions
- Agent Sessions and State
- Conversation Context and Memory
- Callbacks and Lifecycle Events
- Testing and Debugging ADK Agents
- Agent Execution Workflows

Agent Tools, Grounding and Enterprise Data Integration

- Tools and Capabilities of AI Agents
- Integrating External APIs and Services
- Custom Tools Using Python
- Model Context Protocol (MCP) Architecture and Components

- Connecting AI Agents with MCP Servers
- Enterprise Data Integration
- Grounding and Enterprise Knowledge Retrieval
- Google Cloud Search and RAG Capabilities
- Connecting Agents with Enterprise Knowledge Sources
- Retrieval-Based Question Answering
- Integrating Agents with Enterprise Applications
- Tool Execution and Response Handling

Deploying Production-Ready AI Agents

- AI Agent Deployment Architectures
- Production-Ready Agent Development
- Agent Runtime and Agent Engine Capabilities
- Preparing ADK Agents for Deployment
- Deploying Agents and AI Applications on Google Cloud and Cloud Run
- Managing Agent Configurations and Environments
- Authentication and Authorization
- Session and State Persistence
- Exposing Agents through APIs
- Logging and Monitoring
- Troubleshooting Deployed Agents
- Agent Evaluation and Performance Testing
- Production Deployment Best Practices

Building High-Performance Multi-Agent Systems

- Single-Agent versus Multi-Agent Architectures
- Multi-Agent Design Patterns and Orchestration
- Parent and Sub-Agent Relationships
- Sequential, Parallel and Loop-Based Agent Workflows
- Agent Delegation and Coordination
- Agent-to-Agent (A2A) Protocol
- Communication between Distributed Agents
- Integrating Multiple Agents with External Tools
- Collaborative Agent Workflows
- Designing Scalable Multi-Agent Applications
- Managing Multi-Agent Execution and Coordination

Securing and Governing Enterprise AI Agents

- Enterprise AI Agent Governance
- Security Considerations for Agentic AI Applications
- Agent Identity and Access Management
- Authentication and Authorization using Google Cloud IAM
- Security Controls for Agent Tools
- Agent Guardrails and Safety Mechanisms
- Callbacks to Control Agent Execution
- Responsible AI and Safety Considerations
- Protecting Agents against Prompt Injection
- Securing External Integrations and Enterprise Data
- Agent Evaluation and Quality Assessment
- Monitoring Agent Behaviour and Execution
- Logging, Auditing and Compliance
- Managing Risks of Autonomous Agents
- Enterprise AI Governance Best Practices

Scaling AI Agents Across the Enterprise

- Enterprise Architecture for AI Agents
- Designing Scalable Agentic AI Solutions
- Integrating Agents with Enterprise Applications
- Enterprise Workflow Automation
- Multi-Agent Orchestration for Business Processes
- Connecting Agents with Enterprise Knowledge Sources
- Optimizing Agent Performance and Reliability
- Monitoring and Managing Agent Execution
- Cost Optimization and Resource Management
- Enterprise AI Agent Deployment Strategies
- Industry-Specific Use Cases for Agentic AI
- Best Practices for Production-Ready AI Agent Solutions
- Enterprise Adoption and Implementation Considerations
- Capstone: Build an Enterprise Multi-Agent Solution

Course Review and Completion

- Course Review
- Google’s Agent Ecosystem Overview
- Agent Studio and Prompt Engineering Recap
- ADK Agent Development Recap

- MCP and Enterprise Integration Recap
- Multi-Agent Architecture and Orchestration Recap
- Agent Deployment and Security Recap
- Knowledge Validation
- Final Quiz

Certification details

Participants who successfully complete the training program receive a CloudThat course completion certificate.

The course is aligned with Google's Gemini Enterprise Agent Ready (GEAR) learning programme and relevant agent development learning pathways.

GEAR is a learning and skilling programme, not a standalone Google Cloud professional certification. Participants may pursue relevant Google Cloud skill badges and certification pathways separately.

FAQs

Who should attend this course?
This course is designed for developers, AI/ML engineers, cloud professionals, technical consultants and enterprise architects interested in building, deploying and managing AI agents using Google Cloud.

What is GEAR?
GEAR stands for Gemini Enterprise Agent Ready. It is Google's skilling programme focused on helping professionals develop the capabilities required to build, deploy, secure and scale enterprise AI agents.

Does this course cover Agent Studio?
Yes. The course includes dedicated coverage of Agent Studio, including its environment, prompt design, prompt engineering, multimodal prompting and generative AI application development.

Will I learn prompt engineering?
Yes. Participants will learn zero-shot prompting, few-shot prompting, contextual prompting, system instructions, prompt optimization and multimodal prompting using Agent Studio.

What is Google's Agent Development Kit (ADK)?
ADK is Google's framework for developing AI agents. The course covers ADK architecture, agent creation, tools, workflows, sessions, state management and agent orchestration.

Does this course cover multi-agent systems?
Yes. Participants will learn how to create multi-agent applications, configure parent and sub-agents, implement orchestration and understand A2A communication.

Will I learn MCP integration?
Yes. The course introduces Model Context Protocol (MCP) and demonstrates how AI agents can connect with external tools, services and enterprise data sources.

Does the course cover production deployment?
Yes. Participants will explore production agent architectures, deployment options on Google Cloud, authentication, monitoring and troubleshooting.

Do I need programming experience?
Basic Python programming knowledge is recommended, especially for hands-on ADK agent development and custom tool integration.

Does this course prepare me for Google Cloud certification?
This course supports agent development skills aligned with Google's GEAR learning programme. It can complement relevant Google Cloud certification preparation but does not independently confer a Google Cloud certification.

Are hands-on labs included?
Yes. The course includes practical activities covering Agent Studio, Gemini applications, ADK, MCP, multi-agent development, deployment, security and an enterprise capstone project.