PRACTICE QUESTIONS

Google Cloud

Professional Agentic Architect

PUBLISHED BY aiexamcerts

EDITION	Free Set 1
QUESTIONS	10
ISSUED	November 8, 2026
EXAM	Google Cloud Professional Agentic Architect

Shared by **aiexamcerts** on aiexamcerts

10 practice questions for the Google Cloud Professional Agentic Architect exam, each followed by its answer and a short explanation. More free sets and the full practice exam: **aiexamcerts.com/google-cloud/professional-agentic-architect**. Spotted a mistake? Email support@aiexamcerts.com with the exam name and the question number.

---

1. Workflow Builder in Gemini Enterprise was formerly known as Agent _____.

Answer: Designer

The exam guide lists Workflow Builder as the new name for Agent Designer. Questions may use either name when describing low-code agent building.

---

2. A regulated team must control the model weights and serving stack, and no prompt may leave its own network boundary. Which option fits?

A Call a proprietary hosted model through its public SaaS endpoint

B Use a consumer chat application with a company login

C Deploy an open-weight model from Model Garden onto infrastructure it controls, such as GKE

D Fine-tune a proprietary hosted model and then keep calling it as a service from the same provider

Answer: C

Self-hosting an open model gives control over weights, networking and data paths. A SaaS endpoint, even a fine-tuned one, leaves serving under the provider's control.

---

3. A request is to look up a customer's order status, which needs one database lookup. A colleague proposes a team of five coordinating agents. What is the best response?

A Use five agents anyway, because more agents always raise the accuracy of results

B Use a graph workflow with a human approval node placed before each of the lookups

C Use a LoopAgent so that the same lookup is repeated five times before replying

D Use a single agent with a lookup tool, since extra agents add latency and cost

Answer: D

Multi-agent designs pay off when tasks need specialization or parallelism, not for a single simple call. Extra coordination adds cost and more places for errors.

---

4. A main agent needs a specialist agent to compute a result and return it so the main agent can keep talking to the user. How should the specialist be attached in ADK?

A As a tool, wrapped as an agent tool, so control returns to the caller after the call

B As a sub-agent transfer, so that the specialist takes over the conversation for good

C As a ParallelAgent with the user's profile as its only input, running beside the main agent

D As a LoopAgent that repeats its work until the user ends the whole session

Answer: A

Calling a specialist as a tool returns its output to the caller, which stays in control. A transfer hands the conversation over, so the original agent would no longer speak to the user.

---

5. Model Context Protocol is a standard way for an agent client to discover and call external tools and data sources.

A True

B False

Answer: A (True)

MCP defines how clients connect to servers that expose tools, resources and prompts. This is why a single MCP server can be reused across different coding agents.

---

6. A company requires a linter to run every time the coding agent edits a source file, and a prompt instruction asking for this is sometimes ignored. Which mechanism makes it reliable?

A A longer rule that repeats the same instruction in capital letters near the top of every file

B A hook that executes the linter automatically at the defined point in the agent's lifecycle

C A skill that describes the linter but is loaded only on request

D A bigger model that follows prompt instructions more closely

Answer: B

Hooks run deterministic automation at defined events, so they do not depend on the model remembering an instruction. Rules and skills are guidance that the model may still skip.

---

7. A team wants its coding agent to read tickets from their issue tracker and query internal documentation through one standard interface that works across different agent clients. What should they configure?

A An A2A agent card that advertises the coding agent to the other agents it may work with

B A Cloud Run job that copies tickets into the agent's prompt each hour

C An MCP server that exposes the tracker and documentation as tools and resources

D A fine-tuned model that has memorized the tracker's contents

Answer: C

Model Context Protocol standardizes how an agent client reaches external tools and data. A2A is for communication between agents, so an agent card does not give the coding agent access to the tracker.

---

8. Users report that agent responses are slow. Traces show one request spends most of its time waiting on a downstream inventory tool. Which Google Cloud service gives this per-step view?

A Sensitive Data Protection, which scans content for personal data

B Memorystore for Redis, which stores in-memory key-value data for very fast lookups of values

C Cloud Storage, which stores the agent's output as objects

D Cloud Trace, which breaks a request into spans for each model and tool call

Answer: D

Distributed traces expose which step dominates latency, and Cloud Logging adds the detail behind each step. Neither Redis nor Storage provides a request timeline.

---

9. A refund agent built in CX Agent Studio must move the caller from a page that collects the order number to a page that confirms the refund, but only after the order number has been validated. Which construct defines that movement?

A A transition route on the page, conditioned on the order number being valid

B An in-console prompt template that spells out each refund step in ordinary sentences

C A Skill Registry entry that exposes the refund API to every agent in the project

D A few-shot example showing a complete refund conversation from start to finish

Answer: A

Pages hold the state of a flow, and transition routes carry the conditions that move the conversation to another page. A prompt template or few-shot example can shape wording, but it does not enforce an ordered, conditional path.

---

10. If the test suite passes, code produced by a coding agent can be merged without any human review.

A True

B False

Answer: B (False)

Passing tests are necessary but not sufficient, since tests may not cover the changed behavior or may themselves have been altered. Human review of agent changes remains part of a safe workflow.