PRACTICE QUESTIONS

IBM

# C1000-207

IBM watsonx Orchestrate AI Engineer v1 - Associate

PUBLISHED BY aiexamcerts

| | |
|---|---|
| EDITION | Free Set 2 |
| QUESTIONS | 10 |
| ISSUED | November 8, 2026 |
| EXAM | IBM C1000-207, IBM watsonx Orchestrate AI Engineer v1 - Associate |

Shared by **aiexamcerts** on aiexamcerts

10 practice questions for the IBM C1000-207 exam, IBM watsonx Orchestrate AI Engineer v1 - Associate, each followed by its answer and a short explanation. More free sets and the full practice exam: **aiexamcerts.com/ibm/c1000-207**. Spotted a mistake? Email support@aiexamcerts.com with the exam code C1000-207 and the question number.

---

**1.** Inside a workflow, a loop polls a status tool until a job completes. Which design choice most reduces the risk that the loop never ends?

**A** Remove the End Node so the run stays open for later

**B** Give the loop an explicit exit condition and a maximum number of iterations

**C** Add a random delay to each iteration so that polling requests do not all line up

**D** Hide the loop from the observability dashboard so that it is never reported to anyone

**Answer: B**

A Loop node repeats while its condition holds, so an explicit exit condition plus an iteration cap guarantees the run ends even when the job stalls. Removing the End Node breaks the rule that every workflow ends with one. A common mistake is the random delay, which spreads load, but it does nothing to stop the loop.

---

**2.** Fixing an agent to one specific model version, so that a newer release is adopted only when the team chooses, is called _____.

**Answer: version pinning** (also accepted: pinning, pin)

Pinning keeps an agent on a known model version, so its behavior does not change when a provider updates the default. A common wrong answer is versioning, which describes the labels that exist but not the act of locking an agent to one of them.

---

**3.** Where an engineer wants an agent to call one existing operation of a REST API, which kind of tool fits that need most directly?

**A** A knowledge base built from the API's sample JSON responses

**B** An environment that stores the API's base URL and version

**C** A single OpenAPI operation imported as a tool

**D** A flow whose Start Node holds the API key in plain text

**Answer: C**

The ADK can import a single OpenAPI operation as a tool, so the agent calls exactly that endpoint with the parameters its schema defines. A knowledge base indexes documents for retrieval and cannot invoke an operation. The option most likely to mislead is the knowledge base, since sample responses look like data, but reading data is not the same as calling the endpoint.

---

**4.** If a tool receives an order ID that does not exist and returns a clear, structured error such as 'order not found' instead of crashing, what is the main benefit for the agent?

**A** The environment is redeployed automatically with the new error handler included

**B** The knowledge base reindexes the missing order so that later calls to it succeed without errors

**C** The model's temperature drops to zero for the rest of the conversation

**D** The agent can tell the user what failed and choose a sensible next step

**Answer: D**

A structured error gives the agent something it can explain or act on, such as asking the customer for a different order number. A crash leaves the agent with no usable result. The most likely trap is the automatic redeployment, since a new error handler sounds like something that should propagate, but a tool error does not trigger a redeploy.

---

**5.** In distributed tracing, one timed operation inside a trace, such as a single model call, is called a _____.

**Answer: span**

A span records one timed step inside a trace along with its metadata, and a trace groups the spans for one interaction. A common wrong answer is the trace itself, since the whole interaction is the larger unit that contains the spans.

---

**6.** Which TWO practices make an OpenAPI-based tool easier for an agent to call correctly? (Choose two.)

**A** Expose one generic operation that accepts any path and any request body

**B** Document required parameters and the values each one accepts in the schema

**C** Leave parameter descriptions empty to keep the specification short

**D** Give each operation a clear operationId and summary that describe its purpose

**E** Return HTML pages from every endpoint so the agent can read them

**Answer: B, D**

A clear `operationId` and summary tell the agent what each operation is for, and documented required parameters with their allowed values let it build valid requests. A single catch-all operation hides that information from the model. The most likely trap is the generic operation, since fewer operations sounds simpler, but it makes correct calls harder to construct.

---

**7.** When a request is repeated with the same key, a tool that gives the same result instead of acting twice is described as _____.

**Answer: idempotent**

An idempotent operation produces the same outcome however many times it is called with the same request, which is what stops a retried cancellation from canceling twice. Stateless is the tempting wrong answer, but a stateless tool can still act twice when retried, so statelessness alone does not provide the guarantee.

---

**8.** The per-request cutoff that stops a tool call from waiting forever is called a _____.

**Answer: timeout** (also accepted: time-out, time out)

A timeout ends a call that has not answered within a set time, so the agent can recover, retry or report a failure. A retry is the tempting wrong answer, since it is the usual next step after a timeout, but a retry is not the cutoff itself.

---

**9.** A flow needs a small calculation that is simpler to write as Python than as a chain of tool calls. Which node type runs inline Python?

**A** Script, which runs inline Python

**B** Prompt, which makes a single LLM call

**C** Decisions, which applies a decision table

**D** Doc Processing, which extracts text from a document

**Answer: A**

The Script node runs inline Python inside a flow, which suits small calculations and transformations. Prompt makes a model call, Decisions applies a decision table, and Doc Processing extracts text from a document. The option most likely to mislead is Prompt, since the calculation could be phrased as a request to the model, but a model call is not the same as running code.

---

**10.** Which TWO statements about choosing and managing models for agents are accurate? (Choose two.)

**A** Changing temperature permanently alters the model's weights

**B** A model with a larger context window is always cheaper per call

**C** Track deprecation dates so agents move to a supported model before support ends

**D** Once a model is selected, its outputs stay fixed, so no evaluation is needed

**E** Re-run the agent's test set after changing the model behind it

**Answer: C, E**

Swapping a model can change behavior, so the agent's test set should run again, and deprecation dates should be tracked so agents move before support ends. Temperature is a sampling setting rather than a change to the weights, and a larger context window does not make each call cheaper. A common mistake is the larger-window claim, since bigger often sounds better, but more tokens usually cost more.