PRACTICE QUESTIONS

IBM

C1000-207

IBM watsonx Orchestrate AI Engineer v1 - Associate

PUBLISHED BY aiexamcerts

EDITION	Free Set 1
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. A reviewer finds an answer that reads well but cites a policy clause that is absent from the retrieved sources. Which quality dimension has failed?

A Fluency, meaning the sentences are grammatical, easy to read and well organized

B Groundedness, meaning each claim must be supported by the provided sources

C Latency, meaning the reply arrived more slowly than the response-time target set for the agent

D Throughput, meaning many requests were served in one second

Answer: B

Groundedness checks whether each claim is supported by the sources the model was given, and an invented clause fails it even when the prose is fluent. A common mistake is fluency, since the answer reads well, but fluency says nothing about whether the content is supported. Latency and throughput measure speed and volume, not support.

---

2. Which deliverable does the orchestrate agents export command produce for an agent?

A A compiled binary of the agent's Python code for Linux servers

B A screenshot of the agent's most recent conversation

C A ZIP archive that holds the agent's definition

D A signed token that grants the environment administrator access

Answer: C

The export command writes the agent out as a ZIP archive, which can be kept under version control or imported elsewhere. The other options describe artifacts the command does not produce. The option most likely to mislead is the compiled binary, since a Python-based agent might seem to ship as code, but the export is a ZIP of the definition.

---

3. A Loop node stops as soon as its body has run once, whatever its condition says.

A True

B False

Answer: B (False)

A Loop node keeps repeating while its condition holds, so the condition, not the first pass, decides when it stops. A common wrong answer is that one pass always ends the loop, which would make a Loop useless for polling a job until it finishes.

---

4. Given that an agent holds write access to an entire CRM, although it only needs to read account status, which principle is violated, and what is the fix?

A Availability, so add more replicas of the CRM to absorb peak load

B Data residency, so move the CRM to a different region

C Explainability, so add a longer explanation to every answer

D Least privilege, so grant only the read scope that account-status lookups need

Answer: D

Least privilege means giving an agent only the permissions its task needs, and limiting it to read access shrinks what a misbehaving agent can do. This is also the remedy for the excessive agency risk in the OWASP list for LLM applications. A common mistake is availability, since the CRM is central to the business, but the problem is the scope of permission.

---

5. Suppose a refund-dispute procedure runs to several pages and applies to only a small share of conversations. The team wants the agent to load it only when a dispute appears. Which concept fits?

A An agent skill, a reusable procedure in a SKILL.md file that loads only when relevant

B A permanent paragraph in every agent's instructions, sent on every turn of every single conversation

C A knowledge base that indexes the full refund policy PDF for keyword and semantic search

D An environment variable that switches the refund procedure on and off

Answer: A

An agent skill is a reusable procedure defined in a `SKILL.md` file, and the agent loads it only when it is relevant to the request. A permanent instruction paragraph would load on every turn and occupy context that is rarely needed. A common mistake is the knowledge base, since the policy is a document, but a knowledge base is built for grounding answers, not for holding a step-by-step procedure.

---

6. Our agent answers refund questions from memory and sometimes invents policy details. Which change most directly reduces this?

A Raise the temperature so the answers sound more confident, complete and fluent to every customer

B Attach a knowledge base holding the current refund policy and instruct the agent to use it

C Delete the instructions so the model has no rules to break

D Add more collaborator agents so that answers are checked by several models before the reply is sent

Answer: B

Grounding the agent in the current policy document, and instructing it to answer from that source, gives it facts to use instead of memory. Raising the temperature makes output less predictable, not more accurate. The option most likely to mislead is the extra collaborators, since a second opinion sounds like a check, but collaborators that rely on the same memory tend to repeat the same invention.

---

7. When a tool call to a payment API times out and the agent retries it automatically, which retry pattern keeps the retries from overwhelming the API during an outage?

A Immediate retries with no limit until one call succeeds

B Parallel retries from every collaborator at the same moment

C Exponential backoff with jitter and a maximum number of retries

D Retries that first raise the context window before each attempt

Answer: C

Exponential backoff with jitter spaces retries further apart and spreads them out, and a retry cap stops the loop from running forever. This pattern suits only operations that are safe to repeat. The option most likely to mislead is the immediate retry, since fast recovery sounds desirable, but it piles traffic onto a failing system.

---

8. A supervisor agent can route a request to a watsonx Assistant that has been registered as one of its collaborators.

A True

B False

Answer: A (True)

A registered watsonx Assistant instance can serve as a collaborator in an agent network, and supervisor agents use their collaborators to decide where a request goes. A common wrong answer is that only native agents can receive delegated work, but the assistant is a supported kind of collaborator.

---

9. For a customer-facing agent that must give the same answer to the same policy question across repeated runs, which decoding setting is the strongest lever?

A A high temperature, which spreads probability over more candidate tokens

B A larger maximum output length, which allows longer replies

C A longer system prompt with more descriptive adjectives

D A low temperature, which makes sampling more deterministic

Answer: D

Temperature controls how much randomness the sampling step adds, so a low value makes repeated answers more consistent. The maximum output length limits how long a reply can be, not how much it varies. The option most likely to mislead is the larger output limit, since it changes replies, but it does not make them more repeatable.

---

10. After an agent is imported into an active watsonx Orchestrate environment, what state is it in before anyone deploys it?

A A draft (undeployed) state, which end users cannot reach

B A live state, so end users can chat with it straight away

C A paused state that waits for its knowledge base to finish indexing

D A locked state that can only be edited through the connection store

Answer: A

An imported agent enters a draft (undeployed) state in the active environment, and draft agents are not available to end users. Running the deploy command moves the agent to the live state. The most likely trap is the live state, because an import feels like publishing, but deployment is the separate step that makes an agent live.