PRACTICE QUESTIONS

Oracle

1Z0-1154-26

Oracle AI Data Platform Professional

PUBLISHED BY aiexamcerts

EDITION	Free Set 2
QUESTIONS	10
ISSUED	November 8, 2026
EXAM	Oracle 1Z0-1154-26, Oracle AI Data Platform Professional

Shared by **aiexamcerts** on aiexamcerts

10 practice questions for the Oracle 1Z0-1154-26 exam, Oracle AI Data Platform Professional, each followed by its answer and a short explanation. More free sets and the full practice exam: **aiexamcerts.com/oracle/1z0-1154-26**. Spotted a mistake? Email support@aiexamcerts.com with the exam code 1Z0-1154-26 and the question number.

---

1. What is the main benefit of cross-tenancy external tables and volumes?

A Faster cluster startup for every query in the tenancy

B Accessing live data in another tenancy without copying it

C Automatic conversion of all remote files to Delta format

D Removal of the need for any IAM policies in either tenancy

Answer: B

Cross-tenancy access lets users query data where it lives, avoiding ETL and duplicate storage. It still depends on IAM policies such as dynamic groups and admit statements, so it does not remove them.

---

2. Which open table technology does Oracle AI Data Platform use to provide ACID transactions and schema evolution for its tables?

A Plain CSV files in Object Storage

B Apache Kafka topics

C Delta Lake

D JDBC result sets

Answer: C

Delta Lake adds a transaction log over data files, which gives ACID guarantees and schema evolution. CSV files are a supported format but have no transaction log, so they cannot provide those guarantees on their own.

---

3. An agent must answer from the company's policy PDFs, DOCX and HTML files by finding relevant passages. Which platform component supports this?

A A managed volume without indexing

B An if/else task

C A knowledge base

D A Delta Sharing recipient

Answer: C

Knowledge bases store vector embeddings of documents such as PDF, DOCX and HTML and support semantic retrieval for agents. A volume stores the files but does not provide semantic search.

---

4. Which protocol does Oracle AI Data Platform use to share data with recipients?

A Iceberg REST catalog

B Delta Sharing

C GoldenGate replication

D Simba ODBC

Answer: B

Data sharing in the platform is built on the Delta Sharing protocol and avoids copying data. Iceberg compatibility is provided separately by Delta Uniform for table formats.

---

5. A user is granted SELECT on a standard catalog. What does that mean for tables in schemas under that catalog?

A Nothing, because each table needs its own separate grant

B The user can also insert and update rows in the tables

C The permission cascades, so the user can query those tables

D The user can list the schemas but cannot query tables in them

Answer: C

Permissions granted on a parent object apply to the objects it contains. SELECT allows queries, not data changes, so the insert option confuses it with MANAGE.

---

6. How are agents on Oracle AI Data Platform defined?

A As a graph of steps made of nodes such as agents and tools

B As a single prompt string stored in a catalog table

C As a Spark SQL view defined over a managed volume

D As a cron expression attached to a scheduled notebook

Answer: A

Agents are end-to-end applications defined as a graph whose nodes are agents and tools. A prompt alone does not describe the orchestration of tools and steps.

---

7. A user holds the `AI_DATA_PLATFORM_ADMIN` role but cannot open the Oracle AI Data Platform instance from the OCI console. What is the most likely cause?

A The user lacks the required OCI IAM policy permissions for the instance

B The user has no SELECT permission on any catalog yet

C The user has not been added to the AUDITOR system role

D The workspace has no running compute cluster at the moment

Answer: A

Access has two layers, OCI IAM policies first and then platform object permissions, and the IAM policies are needed even for the admin role. Catalog permissions and clusters only matter after the instance can be opened.

---

8. Task A must pass a computed value to Task B in the same job run. Why can't it simply change a job parameter?

A Job parameters can only hold numeric values, never text or JSON

B Job parameters are immutable in tasks, so use output parameters or storage

C Tasks cannot read job parameters, only task-level ones

D Only Python tasks may use parameters

Answer: B

A job parameter keeps its resolved value for the whole run, so tasks cannot overwrite it. Output parameters or intermediate storage are the documented ways to share information.

---

9. A data steward deletes an external table. What happens to the underlying data files?

A They are deleted together with the table definition

B They remain in place, because only the table definition is removed

C They are moved into a managed volume automatically

D They are archived for 30 days

Answer: B

An external table points at data the platform does not manage, so deletion removes only the metadata. Deleting a managed table is different, because the platform owns that data and removes it too.

---

10. Parameters resolve with the precedence job run, then task, then _____.

Answer: job (also accepted: job level, job-level)

Job-level parameters are defaults used when no more specific value is provided.