PRACTICE QUESTIONS

Oracle

1Z0-1154-26

Oracle AI Data Platform Professional

PUBLISHED BY aiexamcerts

EDITION	Free Set 1
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. A PySpark notebook defines df2 = `df.filter(...)` and nothing happens on the cluster. Why?

A The filter was rejected because DataFrames are immutable objects

B The notebook has not been attached to a catalog yet

C Transformations are lazy and run only when an action such as count() is called

D Spark needs a manual run job created before a filter executes

Answer: C

Spark transformations build a plan and execution starts at an action. Immutability of DataFrames is real but does not stop transformations from being defined, so it does not explain the lack of activity.

---

2. A data engineer queries a table in Oracle AI Data Platform using its fully qualified name. Which pattern does the platform use for that name?

A `workspace_name.catalog_name.table_name`

B `schema_name.catalog_name.table_name`

C `catalog_name.volume_name.table_name`

D `catalog_name.schema_name.table_name`

Answer: D

The platform uses a three-part namespace of catalog, schema and table. Workspaces hold notebooks, jobs and compute, not data assets, so a workspace name is not part of the table path, which rules out the workspace-first pattern.

---

3. Which two are default system roles in Oracle AI Data Platform? (Choose two.)

A AUDITOR

B WORKSPACE_OWNER

C AI_DATA_PLATFORM_ADMIN

D SHARE_RECIPIENT

E CATALOG_STEWARD

Answer: A, C

The two system roles are `AI_DATA_PLATFORM_ADMIN` and AUDITOR. `WORKSPACE_OWNER` is tempting because workspaces have owners, but workspace access is set with USER, `PRIVILEGED_USER` and ADMINISTRATOR permissions.

---

4. ADMIN permission on the Master Catalog confers ADMIN permission on all of its child objects.

A True

B False

Answer: A (True)

Permissions follow a hierarchy, and ADMIN at the top of the Master Catalog applies to the catalogs, schemas and tables below it.

---

5. Which PySpark call removes rows that are exact duplicates across the chosen key columns?

A distinct on a single column only

B dropDuplicates

C coalesce

D repartition

Answer: B

`dropDuplicates` removes repeated rows, optionally judging duplicates by a subset of columns. Repartition and coalesce only change how data is split across partitions and do not remove any rows.

---

6. A developer wants to define an agent in code rather than in the visual builder. Which third-party library does the platform documentation name?

A LangGraph

B Apache Airflow

C Delta Standalone

D Simba JDBC

Answer: A

The documentation says agents can be defined through the no-code flow builder or through code with third-party libraries such as LangGraph. Airflow is a workflow scheduler, not an agent framework.

---

7. A notebook writes to a table in an external catalog backed by Autonomous AI Lakehouse using Spark's `saveAsTable()` API. What restriction applies?

A Only append mode is supported, on any existing table

B The table is created automatically in the source if it is missing

C Only overwrite mode works, and the table must already exist in the source

D The write works only from SQL cells

Answer: C

For Autonomous AI Lakehouse based external catalogs, `saveAsTable()` supports overwrite mode only and requires an existing table. Append and auto-create are therefore not available.

---

8. For an external catalog, who manages the lifecycle of the metadata?

A The external source, with the Master Catalog keeping a copy

B The Master Catalog alone, with no copy kept at the source

C The workspace administrator who created the catalog

D The compute cluster

Answer: A

In an external catalog the metadata lifecycle is managed by the external source and the platform keeps a copy. A standard catalog is the case where the platform manages metadata lifecycle.

---

9. In a notebook, the default language applies to every cell and individual cells cannot use a different language.

A True

B False

Answer: B (False)

The language can be changed for the whole notebook or for individual cells, which allows Python and SQL to be combined in one notebook.

---

10. A job schedule can be created in a Paused state and activated later.

A True

B False

Answer: A (True)

When adding a schedule you choose whether it begins Active or Paused, and you can pause or activate it afterward.