

# Reliable Databricks Databricks-Certified-Professional-Data-Engineer Braindumps Book, Online Databricks-Certified-Professional-Data-Engineer Version



Life is beset with all different obstacles that are not easily overcome. For instance, Databricks-Certified-Professional-Data-Engineer exams may be insurmountable barriers for the majority of population. However, with the help of our exam test, exams are no longer problems for you. The reason why our Databricks-Certified-Professional-Data-Engineer training materials outweigh other study prep can be attributed to three aspects, namely free renewal in one year, immediate download after payment and simulation for the software version. Now that using our Databricks-Certified-Professional-Data-Engineer practice materials have become an irresistible trend, why don't you accept Databricks-Certified-Professional-Data-Engineer learning guide with pleasure?

Databricks Certified Professional Data Engineer exam measures a candidate's ability to design, build, and manage data pipelines using Databricks. It covers a wide range of topics, including data ingestion, transformation, storage, and analysis. Candidates must demonstrate their proficiency in using Databricks tools and techniques to solve real-world data engineering problems. Databricks Certified Professional Data Engineer Exam certification exam is ideal for data engineers who want to validate their skills and expertise in using Databricks to build and manage data pipelines.

>> **Reliable Databricks Databricks-Certified-Professional-Data-Engineer Braindumps Book** <<

## Quiz Valid Databricks - Reliable Databricks-Certified-Professional-Data-Engineer Braindumps Book

The social environment is constantly changing, and our Databricks-Certified-Professional-Data-Engineer guide quiz is also advancing with the times. We have all kinds of experiences on the Databricks-Certified-Professional-Data-Engineer study braindumps for many years, so we know that the content of the exam is related to real-time information. The content of Databricks-Certified-Professional-Data-Engineer Exam Materials is constantly updated. Our professional experts have been specializing in this career for over ten years. And we can always provide with you the most accurate and valid Databricks-Certified-Professional-Data-Engineer learning guide.

## Databricks Certified Professional Data Engineer Exam Sample Questions (Q83-Q88):

### NEW QUESTION # 83

A Delta table of weather records is partitioned by date and has the below schema:

date DATE, device\_id INT, temp FLOAT, latitude FLOAT, longitude FLOAT

To find all the records from within the Arctic Circle, you execute a query with the below filter:

```
latitude > 66.3
```

Which statement describes how the Delta engine identifies which files to load?

- A. The Delta log is scanned for min and max statistics for the latitude column
- B. The Parquet file footers are scanned for min and max statistics for the latitude column

- C. The Hive metastore is scanned for min and max statistics for the latitude column
- D. All records are cached to attached storage and then the filter is applied
- E. All records are cached to an operational database and then the filter is applied

**Answer: A**

Explanation:

This is the correct answer because Delta Lake uses a transaction log to store metadata about each table, including min and max statistics for each column in each data file. The Delta engine can use this information to quickly identify which files to load based on a filter condition, without scanning the entire table or the file footers. This is called data skipping and it can improve query performance significantly. Verified References:

[Databricks Certified Data Engineer Professional], under "Delta Lake" section; [Databricks Documentation], under "Optimizations - Data Skipping" section.

In the Transaction log, Delta Lake captures statistics for each data file of the table. These statistics indicate per file:

- Total number of records
  - Minimum value in each column of the first 32 columns of the table
  - Maximum value in each column of the first 32 columns of the table
  - Null value counts for in each column of the first 32 columns of the table
- When a query with a selective filter is executed against the table, the query optimizer uses these statistics to generate the query result. It leverages them to identify data files that may contain records matching the conditional filter.

For the SELECT query in the question, The transaction log is scanned for min and max statistics for the price column

#### NEW QUESTION # 84

The data science team has created and logged a production model using MLflow. The following code correctly imports and applies the production model to output the predictions as a new DataFrame named preds with the schema "customer\_id LONG, predictions DOUBLE, date DATE".

```
from pyspark.sql.functions import current_date

model = mlflow.pyfunc.spark_udf(spark, model_uri="models:/churn/prod")
df = spark.table("customers")
columns = ["account_age", "time_since_last_seen", "app_rating"]
preds = (df.select(
    "customer_id",
    model(*columns).alias("predictions"),
    current_date().alias("date")
))
```



The data science team would like predictions saved to a Delta Lake table with the ability to compare all predictions across time. Churn predictions will be made at most once per day.

Which code block accomplishes this task while minimizing potential compute costs?

- A. `preds.write.format("delta").save("/preds/churn_preds")`
- **B. `preds.write.mode("append").saveAsTable("churn_preds")`**

```
(preds.write
  .format("delta")
  .mode("overwrite")
  .saveAsTable("churn_preds"))
```

- C.

```
(preds.writeStream
  .outputMode("append")
  .option("checkpointPath", "databricks/_checkpoints/churn_preds")
  .table("churn_preds")
)
```

- D.

```

(preds.writeStream
  .outputMode("overwrite") databricks
  .option("checkpointLocation", "checkpoints/churn_preds")
  .start("/preds/churn_preds")
)

```

- E.

Answer: B

**NEW QUESTION # 85**

When you drop a managed table using SQL syntax DROP TABLE table\_name how does it impact metadata, history, and data stored in the table?

- A. Drops table from meta store, drops metadata, history, and data in storage.
- B. Drops table from meta store, meta data and history but keeps the data in storage
- C. Drops table but keeps meta data, history and data in storage
- D. Drops table from meta store and data from storage but keeps metadata and history in storage
- E. Drops table and history but keeps meta data and data in storage

Answer: A

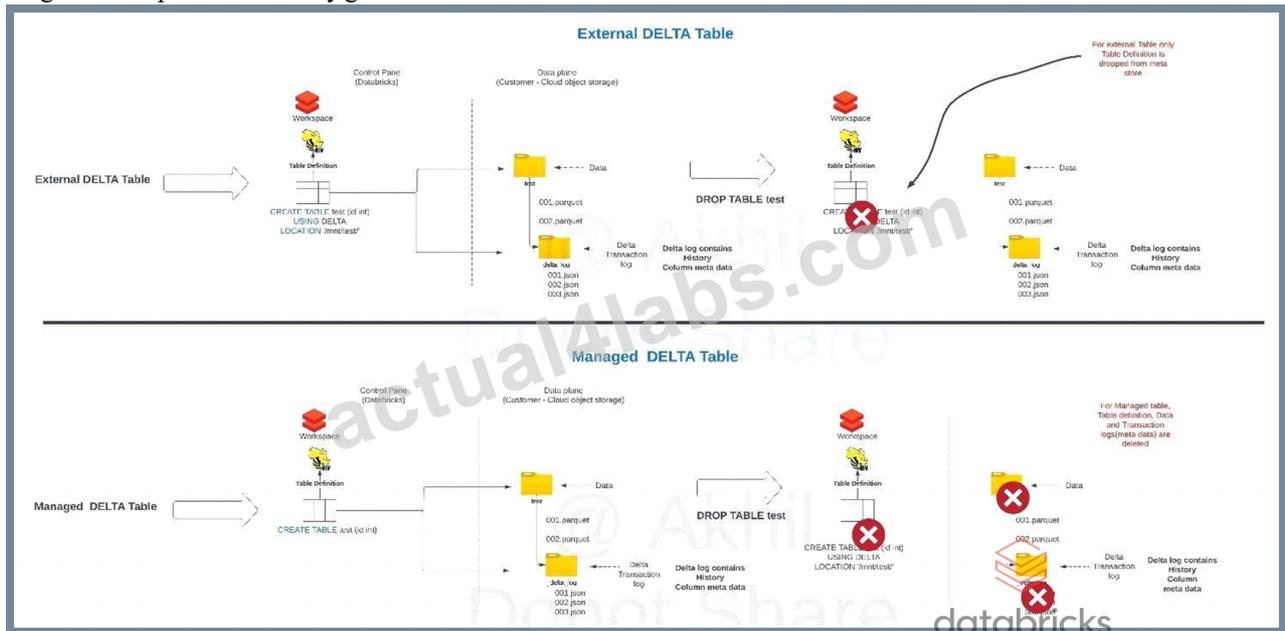
Explanation:

Explanation

For a managed table, a drop command will drop everything from metastore and storage.

See the below image to understand the differences between dropping an external table.

Diagram Description automatically generated



**NEW QUESTION # 86**

The data engineering team maintains the following code:

```

import pyspark.sql.functions as F

(spark.table("silver_customer_sales")
 .groupBy("customer_id")
 .agg(
   F.min("sale_date").alias("first_transaction_date"),
   F.max("sale_date").alias("last_transaction_date"),
   F.mean("sale_total").alias("average_sales"),
   F.countDistinct("order_id").alias("total_orders"),
   F.sum("sale_total").alias("lifetime_value")
 ).write
 .mode("overwrite")
 .table("gold_customer_lifetime_sales_summary")
)

```

Assuming that this code produces logically correct results and the data in the source table has been de-duplicated and validated, which statement describes what will occur when this code is executed?

- A. The `gold_customer_lifetime_sales_summary` table will be overwritten by aggregated values calculated from all records in the `silver_customer_sales` table as a batch job.
- B. The `silver_customer_sales` table will be overwritten by aggregated values calculated from all records in the `gold_customer_lifetime_sales_summary` table as a batch job.
- C. An incremental job will detect if new rows have been written to the `silver_customer_sales` table; if new rows are detected, all aggregates will be recalculated and used to overwrite the `gold_customer_lifetime_sales_summary` table.
- D. An incremental job will leverage running information in the state store to update aggregate values in the `gold_customer_lifetime_sales_summary` table.
- E. A batch job will update the `gold_customer_lifetime_sales_summary` table, replacing only those rows that have different values than the current version of the table, using `customer_id` as the primary key.

**Answer: A**

Explanation:

Explanation

This code is using the `pyspark.sql.functions` library to group the `silver_customer_sales` table by `customer_id` and then aggregate the data using the minimum sale date, maximum sale total, and sum of distinct order ids.

The resulting aggregated data is then written to the `gold_customer_lifetime_sales_summary` table, overwriting any existing data in that table. This is a batch job that does not use any incremental or streaming logic, and does not perform any merge or update operations. Therefore, the code will overwrite the `gold` table with the aggregated values from the `silver` table every time it is executed. References:

<https://docs.databricks.com/spark/latest/dataframes-datasets/introduction-to-dataframes-python.html>

<https://docs.databricks.com/spark/latest/dataframes-datasets/transforming-data-with-dataframes.html>

<https://docs.databricks.com/spark/latest/dataframes-datasets/aggregating-data-with-dataframes.html>

### NEW QUESTION # 87

A table named `user_ltv` is being used to create a view that will be used by data analysts on various teams. Users in the workspace are configured into groups, which are used for setting up data access using ACLs.

The `user_ltv` table has the following schema:

`email STRING, age INT, ltv INT`

The following view definition is executed:

```
CREATE VIEW email_ltv AS
SELECT
CASE WHEN
    is_member('marketing') THEN email
    ELSE 'REDACTED'
END AS email,
age,
ltv
FROM user_ltv
```



databricks

An analyst who is not a member of the marketing group executes the following query:

```
SELECT * FROM email_ltv
```

Which statement describes the results returned by this query?

- A. Three columns will be returned, but one column will be named "redacted" and contain only null values.
- B. The email, age, and ltv columns will be returned with the values in user\_ltv.
- C. Only the email and ltv columns will be returned; the email column will contain all null values.
- **D. Only the email and ltv columns will be returned; the email column will contain the string "REDACTED" in each row.**
- E. The email and ltv columns will be returned with the values in user\_ltv.

**Answer: D**

Explanation:

The code creates a view called email\_ltv that selects the email and ltv columns from a table called user\_ltv, which has the following schema: email STRING, age INT, ltv INT. The code also uses the CASE WHEN expression to replace the email values with the string "REDACTED" if the user is not a member of the marketing group. The user who executes the query is not a member of the marketing group, so they will only see the email and ltv columns, and the email column will contain the string "REDACTED" in each row. Verified Reference: [Databricks Certified Data Engineer Professional], under "Lakehouse" section; Databricks Documentation, under "CASE expression" section.

## NEW QUESTION # 88

.....

No matter how good the product is users will encounter some difficult problems in the process of use. Our Databricks-Certified-Professional-Data-Engineer real exam materials are not exceptional also, in order to enjoy the best product experience, as long as the user is in use process found any problem, can timely feedback to us, for the first time you check our Databricks-Certified-Professional-Data-Engineer Exam Question performance, professional maintenance staff to help users solve problems. Our Databricks-Certified-Professional-Data-Engineer learning reference files have a high efficient product maintenance team, and they can send the Databricks-Certified-Professional-Data-Engineer exam questions to you in a few minutes.

**Online Databricks-Certified-Professional-Data-Engineer Version:** <https://www.actual4labs.com/Databricks/Databricks-Certified-Professional-Data-Engineer-actual-exam-dumps.html>

- 100% Pass Quiz 2026 Databricks-Certified-Professional-Data-Engineer: Pass-Sure Reliable Databricks Certified Professional Data Engineer Exam Braindumps Book  Simply search for **>** Databricks-Certified-Professional-Data-Engineer  for free download on 《 [www.examcollectionpass.com](http://www.examcollectionpass.com) 》  Databricks-Certified-Professional-Data-Engineer New Guide Files
- Valid Databricks-Certified-Professional-Data-Engineer Exam Answers  Databricks-Certified-Professional-Data-Engineer Certification Dump  Download Databricks-Certified-Professional-Data-Engineer Demo  Open website **>** [www.pdfvce.com](http://www.pdfvce.com)  and search for **⇒** Databricks-Certified-Professional-Data-Engineer **⇐** for free download  Latest Databricks-Certified-Professional-Data-Engineer Exam Dumps
- Databricks-Certified-Professional-Data-Engineer Exam Braindumps - Databricks-Certified-Professional-Data-Engineer Exam Simulation - Databricks-Certified-Professional-Data-Engineer Reliable Questions and Answers  Search on  [www.exam4labs.com](http://www.exam4labs.com)  for **✓** Databricks-Certified-Professional-Data-Engineer  **✓**  to obtain exam materials for free download  Databricks-Certified-Professional-Data-Engineer Valid Exam Pass4sure

