

# Training Databricks-Certified-Professional-Data-Engineer Tools | Exam Databricks-Certified-Professional-Data-Engineer Consultant



Obtaining a Databricks-Certified-Professional-Data-Engineer certificate can prove your ability so that you can enhance your market value. When you want to correct the answer after you finish learning, the correct answer for our Databricks-Certified-Professional-Data-Engineer test prep is below each question, and you can correct it based on the answer. In addition, we design small buttons, which can also show or hide the Databricks-Certified-Professional-Data-Engineer Exam Torrent, and you can flexibly and freely choose these two modes according to your habit. In short, you will find the convenience and practicality of our Databricks-Certified-Professional-Data-Engineer quiz guide in the process of learning. We will also continue to innovate and improve functions to provide you with better services.

"PrepAwayExam" created a demo version for customer satisfaction so candidates can evaluate the Databricks-Certified-Professional-Data-Engineer exam questions before purchasing. Also, "PrepAwayExam" has made this Databricks Databricks-Certified-Professional-Data-Engineer practice exam material budget-friendly with many benefits that make it the best choice. Our team of experts who designed this Databricks-Certified-Professional-Data-Engineer Exam Questions assures that whoever prepares with it adequately, there is no doubt of failure and they will pass the Databricks CERTIFICATION EXAM on the first attempt. Purchase our "PrepAwayExam" study material now and get free updates for up to 1 year.

>> Training Databricks-Certified-Professional-Data-Engineer Tools <<

## Exam Databricks-Certified-Professional-Data-Engineer Consultant - Databricks-Certified-Professional-Data-Engineer Exam Certification Cost

It is well acknowledged that people who have a chance to participate in the simulation for the real test, they must have a fantastic advantage over other people to get good grade in the exam. Now, it is so lucky for you to meet this opportunity once in a blue. We offer you the simulation test with the App version of our Databricks-Certified-Professional-Data-Engineer preparation test, in order to let you be familiar with the environment of test as soon as possible. Under the help of the real Databricks-Certified-Professional-Data-Engineer test simulation, you can have a good command of key points which are more likely to be tested in the real test. Therefore that adds more confidence for you to make a full preparation of the upcoming Databricks-Certified-Professional-Data-Engineer Exam. In addition, since you can experience the process of Databricks-Certified-Professional-Data-Engineer the simulation test, you will feel less pressure about the approaching exam. It sounds wonderful, right? Of course, it is. So why not have a try? We promise you will enjoy this study.

## Databricks Certified Professional Data Engineer Exam Sample Questions (Q90-Q95):

### NEW QUESTION # 90

When scheduling Structured Streaming jobs for production, which configuration automatically recovers from query failures and

keeps costs low?

- A. Cluster: New Job Cluster;  
Retries: Unlimited;  
Maximum Concurrent Runs: Unlimited
- B. Cluster: New Job Cluster;  
Retries: None;  
Maximum Concurrent Runs: 1
- C. Cluster: Existing All-Purpose Cluster;  
Retries: Unlimited;  
Maximum Concurrent Runs: 1
- D. Cluster: Existing All-Purpose Cluster;  
Retries: Unlimited;  
Maximum Concurrent Runs: 1
- E. Cluster: Existing All-Purpose Cluster;  
Retries: None;  
Maximum Concurrent Runs: 1

**Answer: B**

Explanation:

Explanation

This is the best configuration for scheduling Structured Streaming jobs for production, as it automatically recovers from query failures and keeps costs low. A new job cluster is created for each run of the job and terminated when the job completes, which saves costs and avoids resource contention. Retries are not needed for Structured Streaming jobs, as they can automatically recover from failures using checkpointing and write-ahead logs. Maximum concurrent runs should be set to 1 to avoid duplicate output or data loss. Verified References: Databricks Certified Data Engineer Professional, under "Monitoring & Logging" section; Databricks Documentation, under "Schedule streaming jobs" section.

### NEW QUESTION # 91

A data pipeline uses Structured Streaming to ingest data from kafka to Delta Lake. Data is being stored in a bronze table, and includes the Kafka\_generated timestamp, key, and value. Three months after the pipeline is deployed the data engineering team has noticed some latency issued during certain times of the day.

A senior data engineer updates the Delta Table's schema and ingestion logic to include the current timestamp (as recoded by Apache Spark) as well the Kafka topic and partition. The team plans to use the additional metadata fields to diagnose the transient processing delays:

Which limitation will the team face while diagnosing this problem?

- A. Spark cannot capture the topic partition fields from the kafka source.
- B. Updating the table schema requires a default value provided for each file added.
- C. Updating the table schema will invalidate the Delta transaction log metadata.
- D. New fields not be computed for historic records.

**Answer: D**

Explanation:

When adding new fields to a Delta table's schema, these fields will not be retrospectively applied to historical records that were ingested before the schema change. Consequently, while the team can use the new metadata fields to investigate transient processing delays moving forward, they will be unable to apply this diagnostic approach to past data that lacks these fields.

:

Databricks documentation on Delta Lake schema management: <https://docs.databricks.com/delta/delta-batch.html#schema-management>

### NEW QUESTION # 92

A data engineer is creating a data ingestion pipeline to understand where customers are taking their rented bicycles during use. The engineer noticed that, over time, data being transmitted from the bicycle sensors fail to include key details like latitude and longitude. Downstream analysts need both the clean records and the quarantined records available for separate processing.

The data engineer already has this code:

```
import dlt
```

```

from pyspark.sql.functions import expr
rules = {
    "valid_lat": "(lat IS NOT NULL)",
    "valid_long": "(long IS NOT NULL)"
}
quarantine_rules = "NOT({})".format(" AND ".join(rules.values()))
@dlt.view
def raw_trips_data():
    return spark.readStream.table("ride_and_go.telemetry.trips")

```

How should the data engineer meet the requirements to capture good and bad data?

- A. `@dlt.table(name="trips_data_quarantine")`  
`def trips_data_quarantine():`  
 `return (`  
 `spark.readStream.table("raw_trips_data")`  
 `.filter(expr(quarantine_rules))`  
 `)`
- B. `@dlt.table`  
`@dlt.expect_all_or_drop(rules)`  
`def trips_data_quarantine():`  
 `return spark.readStream.table("raw_trips_data")`
- C. `@dlt.table(partition_cols=["is_quarantined", ])`  
`@dlt.expect_all(rules)`  
`def trips_data_quarantine():`  
 `return (`  
 `spark.readStream.table("raw_trips_data")`  
 `.withColumn("is_quarantined", expr(quarantine_rules))`  
 `)`
- D. `@dlt.view`  
`@dlt.expect_or_drop("lat_long_present", "(lat IS NOT NULL AND long IS NOT NULL)")`  
`def trips_data_quarantine():`  
 `return spark.readStream.table("ride_and_go.telemetry.trips")`

#### Answer: A

Explanation:

Comprehensive and Detailed

The requirement is that both valid (good) and invalid (bad) records must be captured and available separately for downstream processing. Invalid records should not simply be dropped; they must be quarantined in a dedicated table.

In Databricks Lakeflow Declarative Pipelines (DLT), this is achieved by creating separate output tables:

One table for valid records (Silver table) that pass the expectations.

Another quarantine table that explicitly captures records failing the expectations.

Option A correctly implements this by:

Declaring a DLT table `trips_data_quarantine`.

Using `.filter(expr(quarantine_rules))` to isolate invalid records (records where latitude or longitude is NULL).

This ensures analysts can query both good records (from the main Silver pipeline table) and bad records (from the quarantine table).

Why not the others?

B: Uses `@dlt.expect_or_drop`, which drops invalid records instead of quarantining them. This violates the requirement that quarantined data should be available.

C: Same as B, but applies expectations in bulk with `expect_all_or_drop`. Again, bad data is dropped, not quarantined.

D: Adds an `is_quarantined` flag in the same table. While it marks bad records, it does not separate them into a distinct quarantine table as required by the business use case.

Therefore, Option A is the only solution aligned with Databricks documentation for quarantining invalid data into a dedicated table while keeping valid data in the main pipeline.

#### NEW QUESTION # 93

A data engineering team is in the process of converting their existing data pipeline to utilize Auto Loader for incremental processing in the ingestion of JSON files. One data engineer comes across the following code block in the Auto Loader documentation:

1. `(streaming_df = spark.readStream.format("cloudFiles")`
2. `.option("cloudFiles.format", "json")`

3. .option("cloudFiles.schemaLocation", schemaLocation)

4. .load(sourcePath))

Assuming that schemaLocation and sourcePath have been set correctly, which of the following changes does the data engineer need to make to convert this code block to use Auto Loader to ingest the data?

- A. The data engineer needs to change the format("cloudFiles") line to format("autoLoader")
- B. There is no change required. The data engineer needs to ask their administrator to turn on Auto Loader
- **C. There is no change required. The inclusion of format("cloudFiles") enables the use of Auto Loader**
- D. There is no change required. Databricks automatically uses Auto Loader for streaming reads
- E. The data engineer needs to add the .autoLoader line before the .load(sourcePath) line

**Answer: C**

#### **NEW QUESTION # 94**

Which method is used to solve for coefficients b0, b1, ... bn in your linear regression model:

- A. Integer programming
- B. Apriori Algorithm
- **C. Ordinary Least squares**
- D. Ridge and Lasso

**Answer: C**

Explanation:

Explanation :  $RY = b0 + b1x1 + b2x2 + \dots + bnxn$

In the linear model, the bi's represent the unknown p parameters. The estimates for these unknown parameters are chosen so that, on average, the model provides a reasonable estimate of a person's income based on age and education. In other words, the fitted model should minimize the overall error between the linear model and the actual observations. Ordinary Least Squares (OLS) is a common technique to estimate the parameters

#### **NEW QUESTION # 95**

.....

Our top priority is to help every customer in cracking the Databricks Certified Professional Data Engineer Exam (Databricks-Certified-Professional-Data-Engineer) test. Therefore, we have created these formats so that every applicant can prepare successfully for the Databricks-Certified-Professional-Data-Engineer exam on the first attempt. We are aware that the cost for the registration of the Databricks Databricks-Certified-Professional-Data-Engineer examination is not what everyone can pay. After paying the hefty Databricks-Certified-Professional-Data-Engineer test registration fee, applicants usually run on a tight budget. This is why PrepAwayExam provides you with the Databricks-Certified-Professional-Data-Engineer real questions with up to 90 days of free updates.

#### **Exam Databricks-Certified-Professional-Data-Engineer Consultant:**

<https://www.preawayexam.com/Databricks/braindumps.Databricks-Certified-Professional-Data-Engineer.ete.file.html>

To understand our Databricks-Certified-Professional-Data-Engineer learning questions in detail please look at the introduction of our product on the website pages, Databricks Training Databricks-Certified-Professional-Data-Engineer Tools as soon as i opened it, i got lost, Choose the nay type of Databricks-Certified-Professional-Data-Engineer Databricks Certified Professional Data Engineer Exam practice exam questions that fit your Databricks-Certified-Professional-Data-Engineer exam preparation requirement and budget and start preparation without wasting further time, Databricks Training Databricks-Certified-Professional-Data-Engineer Tools There is no doubt that the answer is yes.

That game didn't need a tutorial, One 'TagID' is used for exactly one 'ElementName', To understand our Databricks-Certified-Professional-Data-Engineer learning questions in detail please look at the introduction of our product on the website pages.

#### **Training Databricks-Certified-Professional-Data-Engineer Tools - Pass Guaranteed Quiz Databricks First-grade Exam Databricks-Certified- Professional-Data-Engineer Consultant**

as soon as i opened it, i got lost, Choose the nay type of Databricks-Certified-Professional-Data-Engineer Databricks Certified Professional Data Engineer Exam practice exam questions that fit your Databricks-Certified-Professional-Data-Engineer exam preparation requirement and budget and start preparation without wasting further time.

There is no doubt that the answer is yes, We promise Databricks-Certified-Professional-Data-Engineer exam cram all we sold is the latest and valid version.