

Pdf Demo Databricks-Certified-Professional-Data-Engineer Download - Latest Databricks-Certified-Professional-Data-Engineer Training



Please don't worry about the purchase process because it's really simple for you. The first step is to select the Databricks-Certified-Professional-Data-Engineer test guide, choose your favorite version, the contents of different version are the same, but different in their ways of using. The second step: fill in with your email and make sure it is correct, because we send our Databricks Certified Professional Data Engineer Exam learn tool to you through the email. Later, if there is an update, our system will automatically send you the latest Databricks Certified Professional Data Engineer Exam version. At the same time, choose the appropriate payment method, such as SWREG, DHpay, etc. Next, enter the payment page, it is noteworthy that we only support credit card payment, do not support debit card. Generally, the system will send the Databricks-Certified-Professional-Data-Engineer Certification material to your mailbox within 10 minutes. If you don't receive it please contact our after-sale service timely.

So we can say that the Databricks-Certified-Professional-Data-Engineer practice questions are the top-notch Databricks Certified Professional Data Engineer Exam (Databricks-Certified-Professional-Data-Engineer) dumps that will provide you with everything that you must need for instant Databricks Databricks-Certified-Professional-Data-Engineer exam preparation. Take the right decision regarding your quick Databricks Certified Professional Data Engineer Exam (Databricks-Certified-Professional-Data-Engineer) exam questions preparation and download the real, valid, and updated Databricks-Certified-Professional-Data-Engineer exam dumps and start this journey.

>> Pdf Demo Databricks-Certified-Professional-Data-Engineer Download <<

Latest Databricks-Certified-Professional-Data-Engineer Training, Databricks-Certified-Professional-Data-Engineer Reliable Test Test

It is never too late to try new things no matter how old you are. Someone always give up their dream because of their ages, someone give up trying to overcome Databricks-Certified-Professional-Data-Engineer exam because it was difficult for them. Now, no matter what the reason you didn't pass the exam, our study materials will try our best to help you. If you are not sure what kinds of Databricks-Certified-Professional-Data-Engineer Exam Question is appropriate for you, you can try our free demo of the PDF version. There must be one that suits you best. Your life will become more meaningful because of your new change, and our Databricks-Certified-Professional-Data-Engineer question torrents will be your first step.

Databricks Certified Professional Data Engineer certification exam is a valuable credential for data engineers who work with Databricks. It demonstrates to potential employers and clients that the holder has the necessary skills and knowledge to work effectively with the platform. Additionally, the certification can enhance the holder's career prospects by providing opportunities for advancement and higher-paying positions.

Databricks Certified Professional Data Engineer Exam Sample Questions (Q133-Q138):

NEW QUESTION # 133

A Spark job is taking longer than expected. Using the Spark UI, a data engineer notes that the Min, Median, and Max Durations for tasks in a particular stage show the minimum and median time to complete a task as roughly the same, but the max duration for a task to be roughly 100 times as long as the minimum.

Which situation is causing increased duration of the overall job?

- A. Network latency due to some cluster nodes being in different regions from the source data
- B. **Skew caused by more data being assigned to a subset of spark-partitions.**
- C. Spill resulting from attached volume storage being too small.
- D. Credential validation errors while pulling data from an external system.
- E. Task queueing resulting from improper thread pool assignment.

Answer: B

Explanation:

This is the correct answer because skew is a common situation that causes increased duration of the overall job. Skew occurs when some partitions have more data than others, resulting in uneven distribution of work among tasks and executors. Skew can be caused by various factors, such as skewed data distribution, improper partitioning strategy, or join operations with skewed keys. Skew can lead to performance issues such as long-running tasks, wasted resources, or even task failures due to memory or disk spills. Verified Reference: [Databricks Certified Data Engineer Professional], under "Performance Tuning" section; Databricks Documentation, under "Skew" section.

NEW QUESTION # 134

Consider flipping a coin for which the probability of heads is p , where p is unknown, and our goal is to estimate p . The obvious approach is to count how many times the coin came up heads and divide by the total number of coin flips. If we flip the coin 1000 times and it comes up heads 367 times, it is very reasonable to estimate p as approximately 0.367. However, suppose we flip the coin only twice and we get heads both times. Is it reasonable to estimate p as 1.0? Intuitively, given that we only flipped the coin twice, it seems a bit rash to conclude that the coin will always come up heads, and _____ is a way of avoiding such rash conclusions.

- A. Linear Regression
- B. **Laplace Smoothing**
- C. Naive Bayes
- D. Logistic Regression

Answer: B

Explanation:

Explanation

Smooth the estimates: consider flipping a coin for which the probability of heads is p , where p is unknown, and

our goal is to estimate p. The obvious approach is to count how many times the coin came up heads and divide by the total number of coin flips. If we flip the coin 1000 times and it comes up heads 367 times, it is very reasonable to estimate p as approximately 0.367. However, suppose we flip the coin only twice and we get heads both times. Is it reasonable to estimate p as 1.0? Intuitively, given that we only flipped the coin twice, it seems a bit rash to conclude that the coin will always come up heads, and smoothing is a way of avoiding such rash conclusions. A simple smoothing method, called Laplace smoothing (or Laplace's law of succession or add-one smoothing in R&N), is to estimate p by (one plus the number of heads) / (two plus the total number of flips). Said differently, if we are keeping count of the number of heads and the number of tails, this rule is equivalent to starting each of our counts at one, rather than zero. Another advantage of Laplace smoothing is that it avoids estimating any probabilities to be zero, even for events never observed in the data. Laplace add-one smoothing now assigns too much probability to unseen words

NEW QUESTION # 135

The team has decided to take advantage of table properties to identify a business owner for each table, which of the following table DDL syntax allows you to populate a table property identifying the business owner of a table CREATE TABLE inventory (id INT, units FLOAT)

- A. CREATE TABLE inventory (id INT, units FLOAT)
SET TAG (business_owner = 'supply chain')
- B. **TBLPROPERTIES (business_owner = 'supply chain')**
- C. CREATE TABLE inventory (id INT, units FLOAT)
SET PROPERTY (business_owner = 'supply chain')
- D. SET TBLPROPERTIES business_owner = 'supply chain'
CREATE TABLE inventory (id INT, units FLOAT)
- E. CREATE TABLE inventory (id INT, units FLOAT)
SET (business_owner = 'supply chain')

Answer: B

Explanation:

Explanation

CREATE TABLE inventory (id INT, units FLOAT) TBLPROPERTIES (business_owner = 'supply chain') Table properties and table options (Databricks SQL) | Databricks on AWS Alter table command can be used to update the TBLPROPERTIES ALTER TABLE inventory SET TBLPROPERTIES(business_owner , 'operations')

NEW QUESTION # 136

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`
`@dlt.expect_all_or_drop(rules)`
`def trips_data_quarantine():`
 `return spark.readStream.table("raw_trips_data")`
- B. `@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")
• 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.table(name="trips_data_quarantine")
    def trips_data_quarantine():
        return (
            spark.readStream.table("raw_trips_data")
            .filter(expr(quarantine_rules))
        )

```

Answer: D

Explanation:

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 # 137

The data governance team has instituted a requirement that all tables containing Personal Identifiable Information (PH) must be clearly annotated. This includes adding column comments, table comments, and setting the custom table property "contains_pii" = true.

The following SQL DDL statement is executed to create a new table:

□ Which command allows manual confirmation that these three requirements have been met?

- A. SHOW TBLPROPERTIES dev.pii test
- B. SHOW TABLES dev
- C. DESCRIBE DETAIL dev.pii test
- D. DESCRIBE EXTENDED dev.pii test
- E. DESCRIBE HISTORY dev.pii test

Answer: D

Explanation:

This is the correct answer because it allows manual confirmation that these three requirements have been met.

The requirements are that all tables containing Personal Identifiable Information (PII) must be clearly annotated, which includes adding column comments, table comments, and setting the custom table property

"contains_pii" = true. The DESCRIBE EXTENDED command is used to display detailed information about a table, such as its schema, location, properties, and comments. By using this command on the dev.pii_test table, one can verify that the table has been created with the correct column comments, table comment, and custom table property as specified in the SQL DDL statement.

Verified References: [Databricks Certified Data Engineer Professional], under "Lakehouse" section; Databricks Documentation, under "DESCRIBE EXTENDED" section.

NEW QUESTION # 138

After clients pay for our Databricks-Certified-Professional-Data-Engineer exam torrent successfully, they will receive the mails sent by our system in 5-10 minutes. Then the client can click the links and download and then you can use our Databricks-Certified-Professional-Data-Engineer questions torrent to learn. Because time is very important for the people who prepare for the exam, the client can download immediately after paying is the great advantage of our Databricks-Certified-Professional-Data-Engineer Guide Torrent. So it is very convenient for the client to use.

Latest Databricks-Certified-Professional-Data-Engineer Training: <https://www.premiumvcedump.com/Databricks/valid-Databricks-Certified-Professional-Data-Engineer-premium-vce-exam-dumps.html>