

Databricks Databricks-Certified-Professional-Data-Engineer PDF Questions - Increase Your Exam Passing Chances



Love is precious and the price of freedom is higher. Do you think that learning day and night has deprived you of your freedom? Then let Our Databricks-Certified-Professional-Data-Engineer guide tests free you from the depths of pain. With Databricks-Certified-Professional-Data-Engineer guide tests, learning will no longer be a burden in your life. You can save much time and money to do other things what meaningful. You will no longer feel tired because of your studies, if you decide to choose and practice our Databricks-Certified-Professional-Data-Engineer Test Answers. Your life will be even more exciting.

Passing the Databricks Certified Professional Data Engineer exam is a significant achievement for any data engineer. It demonstrates that the candidate has a high level of expertise in working with Databricks and can design and manage complex data pipelines. Databricks Certified Professional Data Engineer Exam certification is also highly valued by employers and can lead to new career opportunities and higher salaries.

>> Practice Databricks-Certified-Professional-Data-Engineer Exam <<

New Databricks-Certified-Professional-Data-Engineer Exam Bootcamp | Exam Databricks-Certified-Professional-Data-Engineer Braindumps

If you study with our Databricks-Certified-Professional-Data-Engineer exam questions, then you are better than others, and of course you will get more opportunities. You will never be picked by others. You will become the target of business competition! This will be a happy event! You must understand what it means in this social opportunity. You can get your favorite project and get a higher salary! Our Databricks-Certified-Professional-Data-Engineer simulating exam can give you more than just the success of an exam, but also the various benefits that come along with successful Databricks-Certified-Professional-Data-Engineer exams.

Databricks Certified Professional Data Engineer Certification Exam is a rigorous and comprehensive certification exam that validates a professional's competence in Databricks data engineering principles and techniques. With the exponential growth of data management and processing, data professionals need certification to show that they are experts in handling data. With Databricks Certified Professional Data Engineer certification, data engineers can confidently build, implement and maintain effective and scalable data engineering solutions.

Databricks Certified Professional Data Engineer Exam Sample Questions (Q35-Q40):

NEW QUESTION # 35

Which configuration parameter directly affects the size of a spark-partition upon ingestion of data into Spark?

- A. spark.sql.files.openCostInBytes
- B. spark.sql.autoBroadcastJoinThreshold

- C. `spark.sql.files.maxPartitionBytes`
- D. `spark.sql.adaptive.coalescePartitions.minPartitionNum`
- E. `spark.sql.adaptive.advisoryPartitionSizeInBytes`

Answer: C

Explanation:

This is the correct answer because `spark.sql.files.maxPartitionBytes` is a configuration parameter that directly affects the size of a spark-partition upon ingestion of data into Spark. This parameter configures the maximum number of bytes to pack into a single partition when reading files from file-based sources such as Parquet, JSON and ORC. The default value is 128 MB, which means each partition will be roughly 128 MB in size, unless there are too many small files or only one large file. Verified Reference: [Databricks Certified Data Engineer Professional], under "Spark Configuration" section; Databricks Documentation, under "Available Properties - `spark.sql.files.maxPartitionBytes`" section.

NEW QUESTION # 36

Which statement describes integration testing?

- A. Validates interactions between subsystems of your application
- B. Validates an application use case
- C. Requires an automated testing framework
- D. Requires manual intervention
- E. Validates behavior of individual elements of your application

Answer: A

Explanation:

This is the correct answer because it describes integration testing. Integration testing is a type of testing that validates interactions between subsystems of your application, such as modules, components, or services.

Integration testing ensures that the subsystems work together as expected and produce the correct outputs or results. Integration testing can be done at different levels of granularity, such as component integration testing, system integration testing, or end-to-end testing. Integration testing can help detect errors or bugs that may not be found by unit testing, which only validates behavior of individual elements of your application. Verified References: [Databricks Certified Data Engineer Professional], under "Testing" section; Databricks Documentation, under "Integration testing" section.

NEW QUESTION # 37

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. 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.
- C. 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.
- D. 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.
- E. An incremental job will leverage running information in the state store to update aggregate values in the `gold_customer_lifetime_sales_summary` table.

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 # 38

A data engineer is designing an append-only pipeline that needs to handle both batch and streaming data in Delta Lake. The team wants to ensure that the streaming component can efficiently track which data has already been processed.

Which configuration should be set to enable this?

- A. `mergeSchema`
- B. `overwriteSchema`
- C. `partitionBy`

- **D. checkpointLocation**

Answer: D

Explanation:

Comprehensive and Detailed Explanation From Exact Extract of Databricks Data Engineer Documents:

When working with Delta Lake streaming ingestion, checkpointing is critical for maintaining fault tolerance and ensuring exactly-once data processing semantics.

The checkpointLocation parameter defines the directory where Spark Structured Streaming stores progress information, offsets, and metadata. This allows the engine to resume processing from the last committed offset without reprocessing previously ingested data. Without checkpointing, each stream restart would reprocess all data, leading to duplicates. Parameters like partitionBy or schema options (mergeSchema / overwriteSchema) affect table structure, not data lineage tracking. Therefore, the correct and required configuration for efficient streaming state management is checkpointLocation.

NEW QUESTION # 39

You are currently working on a production job failure with a job set up in job clusters due to a data issue, what cluster do you need to start to investigate and analyze the data?

- A. Databricks SQL Endpoint can be used to investigate the issue
- B. Existing job cluster can be used to investigate the issue
- C. A Job cluster can be used to analyze the problem
- **D. All-purpose cluster/ interactive cluster is the recommended way to run commands and view the data.**

Answer: D

Explanation:

Explanation

Answer is All-purpose cluster/ interactive cluster is the recommended way to run commands and view the data.

A job cluster can not provide a way for a user to interact with a notebook once the job is submitted, but an Interactive cluster allows to you display data, view visualizations write or edit queries, which makes it a perfect fit to investigate and analyze the data.

NEW QUESTION # 40

.....

New Databricks-Certified-Professional-Data-Engineer Exam Bootcamp: https://www.testsdumps.com/Databricks-Certified-Professional-Data-Engineer_real-exam-dumps.html

- Regular Databricks-Certified-Professional-Data-Engineer Update ☐ New Databricks-Certified-Professional-Data-Engineer Exam Vce ☐ Test Databricks-Certified-Professional-Data-Engineer Question ☐ Easily obtain ☼ Databricks-Certified-Professional-Data-Engineer ☐☼☐ for free download through 【 www.prep4away.com 】 ☐ Databricks-Certified-Professional-Data-Engineer Braindump Free
- Regular Databricks-Certified-Professional-Data-Engineer Update ☐ Databricks-Certified-Professional-Data-Engineer Training Material ☐ Databricks-Certified-Professional-Data-Engineer Paper ☐ Easily obtain free download of 《 Databricks-Certified-Professional-Data-Engineer 》 by searching on [www.pdfvce.com] ☐ Pdf Demo Databricks-Certified-Professional-Data-Engineer Download
- Practice Databricks-Certified-Professional-Data-Engineer Exams Free ☐ 100% Databricks-Certified-Professional-Data-Engineer Accuracy ☐ 100% Databricks-Certified-Professional-Data-Engineer Accuracy ↗ Easily obtain 【 Databricks-Certified-Professional-Data-Engineer 】 for free download through ➤ www.dumpsmaterials.com ☐ ☐ New Databricks-Certified-Professional-Data-Engineer Test Book
- High Quality and High Efficiency Databricks-Certified-Professional-Data-Engineer Study Braindumps - Pdfvce ☐ Easily obtain free download of { Databricks-Certified-Professional-Data-Engineer } by searching on ➡ www.pdfvce.com ☐ ☐ ☐ New Databricks-Certified-Professional-Data-Engineer Exam Vce
- HOT Practice Databricks-Certified-Professional-Data-Engineer Exam - Databricks Databricks Certified Professional Data Engineer Exam - High-quality New Databricks-Certified-Professional-Data-Engineer Exam Bootcamp ☐ Search for 「 Databricks-Certified-Professional-Data-Engineer 」 on ☐ www.prep4away.com ☐ immediately to obtain a free download ☐ Databricks-Certified-Professional-Data-Engineer Latest Dumps Files
- New Databricks-Certified-Professional-Data-Engineer Test Book ☐ Databricks-Certified-Professional-Data-Engineer Real Exam Answers ☐ Databricks-Certified-Professional-Data-Engineer Paper ☐ Immediately open ➤ www.pdfvce.com ☐ and search for ➡ Databricks-Certified-Professional-Data-Engineer ☐ to obtain a free download ☐ Databricks-

Certified-Professional-Data-Engineer Exam Lab Questions

- [illegible]