

Databricks-Certified-Data-Engineer-Professional Vce File

| Databricks-Certified-Data-Engineer-Professional

Authentic Exam Questions



Our Databricks-Certified-Data-Engineer-Professional exam questions have a very high hit rate, of course, will have a very high pass rate. Before you select a product, you must have made a comparison of your own pass rates. Our Databricks-Certified-Data-Engineer-Professional study materials must appear at the top of your list. And our Databricks-Certified-Data-Engineer-Professional learning quiz has a 99% pass rate. This is the result of our efforts and the best gift to the user. Our Databricks-Certified-Data-Engineer-Professional Study Materials can have such a high pass rate, and it is the result of step by step that all members uphold the concept of customer first. If you use a trial version of Databricks-Certified-Data-Engineer-Professional training prep, you will want to buy it!

Databricks Databricks-Certified-Data-Engineer-Professional Exam Syllabus Topics:

Section	Objectives
Data Ingestion and Processing	- Batch and streaming ingestion with Auto Loader - Structured Streaming fundamentals - ETL pipeline design patterns
Production Pipelines and Orchestration	- Databricks Workflows - Error handling and recovery strategies - Job scheduling and monitoring
Delta Lake and Data Management	- Schema evolution and enforcement - Time travel and versioning - Delta Lake transactions and ACID properties
Data Modeling and Transformation	- Performance optimization techniques - Spark SQL transformations - Dimensional modeling concepts

Databricks Lakehouse Platform Architecture	- Medallion architecture (Bronze, Silver, Gold) - Data governance concepts (Unity Catalog basics) - Workspace and cluster architecture
--	--

>> **Databricks-Certified-Data-Engineer-Professional Vce File** <<

Databricks-Certified-Data-Engineer-Professional Authentic Exam Questions - Databricks-Certified-Data-Engineer-Professional Reliable Exam Preparation

If you want to get some achievement in the IT field Databricks certifications will be a stepping-stone. In fact high senior positions have a large demand. Databricks-Certified-Data-Engineer-Professional new test braindumps will pave the way for you to clear exam and obtain a certification. If you are an experienced IT test engine, owing one certification under the help of Databricks-Certified-Data-Engineer-Professional new test braindumps will improve your value; companies may have more cooperation opportunities.

Databricks Certified Data Engineer Professional Exam Sample Questions (Q50-Q55):

NEW QUESTION # 50

A data engineer is building a Lakeflow Declarative Pipelines pipeline to process healthcare claims data. A metadata JSON file defines data quality rules for multiple tables, including:

```
{
  "claims": [
    {"name": "valid_patient_id", "constraint": "patient_id IS NOT NULL"},
    {"name": "non_negative_amount", "constraint": "claim_amount >= 0"}
  ]
}
```

The pipeline must dynamically apply these rules to the claims table without hardcoding the rules. How should the data engineer achieve this?

- A. Load the JSON metadata, loop through its entries, and apply expectations using `dlt.expect_all`.
- B. Reference each expectation with `@dlt.expect` decorators in the table declaration.
- C. Invoke an external API to validate records against the metadata rules.
- D. Use a SQL CONSTRAINT block referencing the JSON file path.

Answer: A

Explanation:

Lakeflow Declarative Pipelines provide the `expect_all` method for programmatically applying multiple data quality expectations at once. The documentation explains that `@dlt.expect_all` accepts a dictionary of expectation names mapped to SQL constraints, allowing rules to be dynamically loaded from metadata such as JSON files. This ensures that pipelines remain maintainable and scalable without needing to hardcode individual `@dlt.expect` decorators. The event logs will track each expectation's pass and fail counts individually, making it auditable.

NEW QUESTION # 51

A junior data engineer seeks to leverage Delta Lake's Change Data Feed functionality to create a Type 1 table representing all of the values that have ever been valid for all rows in a bronze table created with the property `delta.enableChangeDataFeed = true`. They plan to execute the following code as a daily job:

```

from pyspark.sql.functions import col

(spark.read.format("delta")
 .option("readChangeFeed", "true")
 .option("startingVersion", 0)
 .table("bronze")
 .filter(col("_change_type").isin(["update_postimage", "insert"]))
 .write
 .mode("append")
 .table("bronze_history_type1")
)

```

Which statement describes the execution and results of running the above query multiple times?

- A. Each time the job is executed, the entire available history of inserted or updated records will be appended to the target table, resulting in many duplicate entries.
- B. Each time the job is executed, only those records that have been inserted or updated since the last execution will be appended to the target table giving the desired result.
- C. Each time the job is executed, newly updated records will be merged into the target table, overwriting previous values with the same primary keys.
- D. Each time the job is executed, the target table will be overwritten using the entire history of inserted or updated records, giving the desired result.
- E. Each time the job is executed, the differences between the original and current versions are calculated; this may result in duplicate entries for some records.

Answer: A

Explanation:

Reading table's changes, captured by CDF, using `spark.read` means that you are reading them as a static source. So, each time you run the query, all table's changes (starting from the specified `startingVersion`) will be read.

NEW QUESTION # 52

A table named `user_ltv` is being used to create a view that will be used by data analysts on Get Latest & Actual Certified-Data-Engineer-Professional Exam's Question and Answers from 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,
ltv
FROM user_ltv

```

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

Answer: C

Explanation:

The code creates a view called email_ltv that selects the email and ltv columns from a table called user_itv, 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.

NEW QUESTION # 53

Although the Databricks Utilities Secrets module provides tools to store sensitive credentials and avoid accidentally displaying them in plain text users should still be careful with which credentials are stored here and which users have access to using these secrets. Which statement describes a limitation of Databricks Secrets?

- A. Because the SHA256 hash is used to obfuscate stored secrets, reversing this hash will display Get Latest & Actual Certified-Data-Engineer-Professional Exam's Question and Answers from the value in plain text.
- B. Iterating through a stored secret and printing each character will display secret contents in plain text.
- C. Secrets are stored in an administrators-only table within the Hive Metastore; database administrators have permission to query this table by default.
- D. The Databricks REST API can be used to list secrets in plain text if the personal access token has proper credentials.
- E. Account administrators can see all secrets in plain text by logging on to the Databricks Accounts console.

Answer: D

Explanation:

This is the correct answer because it describes a limitation of Databricks Secrets. Databricks Secrets is a module that provides tools to store sensitive credentials and avoid accidentally displaying them in plain text. Databricks Secrets allows creating secret scopes, which are collections of secrets that can be accessed by users or groups. Databricks Secrets also allows creating and managing secrets using the Databricks CLI or the Databricks REST API. However, a limitation of Databricks Secrets is that the Databricks REST API can be used to list secrets in plain text if the personal access token has proper credentials. Therefore, users should still be careful with which credentials are stored in Databricks Secrets and which users have access to using these secrets.

NEW QUESTION # 54

Each configuration below is identical to the extent that each cluster has 400 GB total of RAM 160 total cores and only one Executor per VM.

Given an extremely long-running job for which completion must be guaranteed, which cluster configuration will be able to guarantee completion of the job in light of one or more VM failures?

- A. - Total VMs: 1
- 400 GB per Executor
- 160 Cores/Executor
- B. - Total VMs: 16
- 25 GB per Executor
- 10 Cores / Executor
- C. - Total VMs: 2
- 200 GB per Executor
- 80 Cores / Executor
- D. - Total VMs: 4
- 100 GB per Executor
- 40 Cores / Executor
- E. - Total VMs: 8
- 50 GB per Executor

Answer: B

• • • • •

Databricks-Certified-Data-Engineer-Professional Authentic Exam Questions: <https://www.exam4free.com/Databricks-Certified-Data-Engineer-Professional-valid-dumps.html>

- [illegible]

myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt,
myportal.utt.edu.tt, Disposable vapes