


# Databricks-Certified-Data-Engineer-Professional Dumps Free Download, Accurate Databricks-Certified-Data-Engineer-Professional Test



Our Databricks-Certified-Data-Engineer-Professional practice materials are prepared for the diligent people craving for success. Almost all people pursue a promising career, the reality is not everyone acts quickly and persistently. That is the reason why success belongs to few people. Once you try our Databricks-Certified-Data-Engineer-Professional exam test, you will be motivated greatly and begin to make changes. Our study questions always update frequently to guarantee that you can get enough test banks and follow the trend in the theory and the practice. That is to say, our product boosts many advantages and to gain a better understanding of our Databricks-Certified-Data-Engineer-Professional question torrent.

NewPassLeader provides 24/7 customer support to answer any of your queries or concerns regarding the Databricks Certified Data Engineer Professional Exam (Databricks-Certified-Data-Engineer-Professional) certification exam. They have a team of highly skilled and experienced professionals who have a thorough knowledge of the Databricks Certified Data Engineer Professional Exam (Databricks-Certified-Data-Engineer-Professional) exam questions and format.

>> **Databricks-Certified-Data-Engineer-Professional Dumps Free Download** <<

## Reliable Databricks-Certified-Data-Engineer-Professional Dumps Free Download Supply you Verified Accurate Test for Databricks-Certified-Data-Engineer-Professional: Databricks Certified Data Engineer Professional Exam to Prepare easily

The NewPassLeader is one of the top-rated and trusted platforms that are committed to making the Databricks Certified Data Engineer Professional Exam (Databricks-Certified-Data-Engineer-Professional) certification exam journey successful. To achieve this objective NewPassLeader has hired a team of experienced and qualified Databricks-Certified-Data-Engineer-Professional Exam trainers. They work together and put all their expertise to maintain the top standard of Databricks Databricks-Certified-Data-Engineer-Professional practice test all the time.

## Databricks Certified Data Engineer Professional Exam Sample Questions (Q141-Q146):

### NEW QUESTION # 141

A data engineer is implementing Unity Catalog governance for a multi-team environment. Data scientists need interactive clusters for basic data exploration tasks, while automated ETL jobs require dedicated processing. How should the data engineer configure cluster isolation policies to enforce least privilege and ensure Unity Catalog compliance?

- A. Create compute policies with STANDARD access mode for interactive workloads and DEDICATED access mode for automated jobs.
- B. Allow all users to create any cluster type and rely on manual configuration to enable Unity Catalog access modes.
- C. Configure all clusters with NO ISOLATION\_SHARED access mode since Unity Catalog works with any cluster

configuration.

- D. Use only DEDICATED access mode for both interactive workloads and automated jobs to maximize security isolation.

**Answer: A**

Explanation:

Unity Catalog enforces governance and data isolation through cluster access modes and compute policies. According to Databricks documentation, "Interactive clusters that multiple users share should use Standard access mode, while automated jobs and production pipelines should use Dedicated access mode for stricter isolation." Standard access mode allows multiple users to share the same compute resources but still respects Unity Catalog permissions. Dedicated access mode isolates the job run's execution environment, ensuring that data access is limited to the job's identity. Configuring these modes within compute policies enforces least privilege and ensures all compute complies with Unity Catalog security standards. Options A and C are incorrect because using only Dedicated clusters reduces resource efficiency, while "No isolation" clusters are not Unity Catalog compliant.

### NEW QUESTION # 142

A data engineer is configuring a Databricks Asset Bundle to deploy a job with granular permissions.

The requirements are:

- Grant the data-engineers group CAN\_MANAGE access to the job.
- Ensure the auditors' group can view the job but not modify/run it.
- Avoid granting unintended permissions to other users/groups.

How should the data engineer deploy the job while meeting the requirements?

- A. resources:  
jobs:  
my-job:  
name: data-pipeline  
tasks: [...]  
job\_clusters: [...]  
permissions:  
- group\_name: data-engineers  
level: CAN\_MANAGE  
- group\_name: auditors  
level: CAN\_VIEW  
- group\_name: admin-team  
level: IS\_OWNER
- B. permissions:  
- group\_name: data-engineers  
level: CAN\_MANAGE  
- group\_name: auditors  
level: CAN\_VIEW  
resources:  
jobs:  
my-job:  
name: data-pipeline  
tasks: [...]  
job\_clusters: [...]
- C. resources:  
jobs:  
my-job:  
name: data-pipeline  
tasks: [...]  
job\_clusters: [...]  
permissions:  
- group\_name: data-engineers  
level: CAN\_MANAGE  
- group\_name: auditors  
level: CAN\_VIEW
- D. resources:  
jobs:  
my-job:

```
name: data-pipeline
tasks: [...]
job: [...]
permissions:
- group_name: data-engineers
level: CAN_MANAGE
permissions:
- group_name: auditors
level: CAN_VIEW
```

**Answer: C**

Explanation:

Databricks Asset Bundles (DABs) allow jobs, clusters, and permissions to be defined as code in YAML configuration files. According to the Databricks documentation on job permissions and bundle deployment, when defining permissions within a job resource, they must be scoped directly under that specific job's definition. This ensures that permissions are applied only to the intended job resource and not inadvertently propagated to other jobs or resources.

In this scenario, the data engineer must grant the data-engineers group CAN\_MANAGE access, allowing them to configure, edit, and manage the job, while the auditors group should only have CAN\_VIEW, giving them read-only access to see configurations and results without the ability to modify or execute. Importantly, no additional groups should be granted permissions, in order to follow the principle of least privilege.

Options A and B introduce unnecessary or unintended groups (like admin-team in A) or define permissions outside of the job scope (as in B). Option C improperly separates the permissions block outside the job resource, which is not aligned with Databricks bundle best practices.

Option D is the correct approach because it defines the job resource my-job with its name, tasks, clusters, and the exact intended permissions (CAN\_MANAGE for data-engineers and CAN\_VIEW for auditors). This aligns with Databricks' principle of least privilege and ensures compliance with governance standards in Unity Catalog-enabled workspaces.

#### NEW QUESTION # 143

Given the following PySpark code snippet in a Databricks notebook:

```
filtered_df = spark.read.format("delta").load("/mnt/data/large_table")
\
.filtered_df.filter("event_date > '2024-01-01'")
filtered_df.count()
```

The data engineer notices from the Query Profiler that the scan operator for filtered\_df is reading almost all files, despite the filter being applied.

What is the probable reason for poor data skipping?

- A. The filter is executed only after the full data scan, preventing data skipping.
- B. The filter condition involves a data type excluded from data skipping support.
- C. The event\_date column is outside the table's partitioning and Z-ordering scheme.
- D. The Delta table lacks optimization that enables dynamic file pruning.

**Answer: C**

Explanation:

Delta Lake's data skipping relies on partitioning and clustering (such as Z-ordering) on the filtered columns. If event\_date is neither a partition column nor included in the table's clustering strategy, Spark must scan most files because file-level statistics cannot be effectively used to prune irrelevant data.

#### NEW QUESTION # 144

The data engineering team has configured a Databricks SQL query and alert to monitor the values in a Delta Lake table. The recent\_sensor\_recordings table contains an identifying sensor\_id alongside the timestamp and temperature for the most recent 5 minutes of recordings.

The below query is used to create the alert:

The query is set to refresh each minute and always completes in less than 10 seconds. The alert is set to trigger when mean (temperature) > 120. Notifications are triggered to be sent at most Get Latest & Actual Certified-Data-Engineer-Professional Exam's Question and Answers from every 1 minute.

If this alert raises notifications for 3 consecutive minutes and then stops, which statement must be true?

- A. The total average temperature across all sensors exceeded 120 on three consecutive executions of the query
- B. The maximum temperature recording for at least one sensor exceeded 120 on three consecutive executions of the query
- C. The recent\_sensor\_recordingstable was unresponsive for three consecutive runs of the query
- **D. The average temperature recordings for at least one sensor exceeded 120 on three consecutive executions of the query**
- E. The source query failed to update properly for three consecutive minutes and then restarted

**Answer: D**

Explanation:

This is the correct answer because the query is using a GROUP BY clause on the sensor\_id column, which means it will calculate the mean temperature for each sensor separately. The alert will trigger when the mean temperature for any sensor is greater than 120, which means at least one sensor had an average temperature above 120 for three consecutive minutes. The alert will stop when the mean temperature for all sensors drops below 120.

#### NEW QUESTION # 145

The data architect has mandated that all tables in the Lakehouse should be configured as external Delta Lake tables. Which approach will ensure that this requirement is met?

- A. When the workspace is being configured, make sure that external cloud object storage has been mounted.
- **B. Whenever a table is being created, make sure that the location keyword is used.**
- C. Whenever a database is being created, make sure that the location keyword is used
- D. When configuring an external data warehouse for all table storage, leverage Databricks for all ELT.
- E. When tables are created, make sure that the external keyword is used in the create table statement.

**Answer: B**

Explanation:

This is the correct answer because it ensures that this requirement is met. The requirement is that all tables in the Lakehouse should be configured as external Delta Lake tables. An external table is a table that is stored outside of the default warehouse directory and whose metadata is not managed by Databricks. An external table can be created by using the location keyword to specify the path to an existing directory in a cloud storage system, such as DBFS or S3. By creating external tables, the data engineering team can avoid losing data if they drop or overwrite the table, as well as leverage existing data without moving or copying it.

#### NEW QUESTION # 146

.....

We are popular not only because our outstanding Databricks-Certified-Data-Engineer-Professional practice dumps, but also for our well-praised after-sales service. After purchasing our Databricks-Certified-Data-Engineer-Professional practice materials, the free updates will be sent to your mailbox for one year long if our experts make any of our Databricks-Certified-Data-Engineer-Professional Guide materials. They are also easily understood by exam candidates. Our Databricks-Certified-Data-Engineer-Professional actual exam can succeed you from tremendous materials with least time and quickest pace based on your own drive and practice to win.

**Accurate Databricks-Certified-Data-Engineer-Professional Test:** <https://www.newpassleader.com/Databricks/Databricks-Certified-Data-Engineer-Professional-exam-preparation-materials.html>

There is no doubt that it is very difficult for most people to pass the Databricks-Certified-Data-Engineer-Professional exam and have the certification easily. It is well known that Accurate Databricks-Certified-Data-Engineer-Professional Test certification training is experiencing a great demand in IT industry area, Databricks Databricks-Certified-Data-Engineer-Professional Dumps Free Download No matter what kind of problems you meet please don't be shy to let us know, it's our pleasure to help you in any way, Databricks Databricks-Certified-Data-Engineer-Professional Dumps Free Download In case you feel confuse and cannot decide which one to practice and remember.

Improve Your Debugging with Debugger Markup Language, Fine-tune shot length, There is no doubt that it is very difficult for most people to pass the Databricks-Certified-Data-Engineer-Professional Exam and have the certification easily.

## Practical Databricks-Certified-Data-Engineer-Professional Question Dumps

