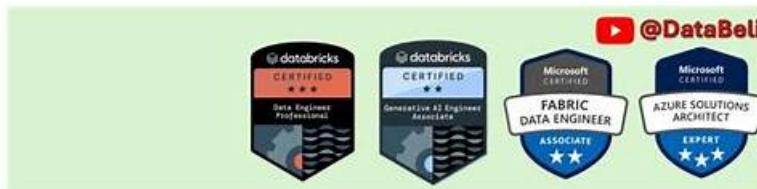


# Download Databricks-Certified-Professional-Data-Engineer Free Dumps | Databricks-Certified-Professional-Data-Engineer Exam Assessment



In order to facilitate the wide variety of users' needs the Databricks-Certified-Professional-Data-Engineer study guide have developed three models with the highest application rate in the present - PDF, software and online. No matter you are a student, a office staff or even a housewife, you can always find your most suitable way to study our Databricks-Certified-Professional-Data-Engineer Exam Q&A. Generally speaking, these three versions of our Databricks-Certified-Professional-Data-Engineer learning guide can support study on paper, computer and all kinds of eletronic devices. They are quite convenient.

Databricks Certified Professional Data Engineer certification exam consists of 60 multiple-choice questions that need to be answered within 90 minutes. The questions are designed to test the candidate's knowledge of Databricks and its features, as well as their ability to apply this knowledge to real-world scenarios. To pass the exam, candidates need to achieve a minimum score of 70%. Databricks Certified Professional Data Engineer Exam certification is valid for two years and can be renewed by retaking the exam or by completing a set of continuing education activities. Obtaining the Databricks Certified Professional Data Engineer certification can help data professionals demonstrate their expertise in building and deploying data engineering solutions using Databricks, which can be a valuable asset in today's data-driven business landscape.

[>> Download Databricks-Certified-Professional-Data-Engineer Free Dumps <<](#)

## Best Accurate Download Databricks-Certified-Professional-Data-Engineer Free Dumps by ITdumpsfree

With many advantages such as immediate download, simulation before the real test as well as high degree of privacy, our Databricks-Certified-Professional-Data-Engineer actual exam survives all the ordeals throughout its development and remains one of the best choices for those in preparation for exams. Many people have gained good grades after using our Databricks-Certified-Professional-Data-Engineer real test, so you will also enjoy the good results. Don't hesitate any more. Time and tide wait for no man. Now that using our Databricks-Certified-Professional-Data-Engineer practice materials have become an irresistible trend, why don't you accept it with pleasure?

Databricks Certified Professional Data Engineer certification is designed for data engineers who work with the Databricks platform and have a deep understanding of data engineering concepts. Databricks Certified Professional Data Engineer Exam certification exam tests the candidate's ability to design, build, and maintain data pipelines using Databricks, as well as their knowledge of data modeling, data warehousing, and data governance. Databricks Certified Professional Data Engineer Exam certification is recognized globally and indicates that the candidate has the skills and expertise needed to work with Databricks.

## Databricks Certified Professional Data Engineer Exam Sample Questions (Q82-Q87):

### NEW QUESTION # 82

How to determine if a table is a managed table vs external table?

- A. Run IS\_MANAGED('table\_name') function
- B. Run SQL command DESCRIBE EXTENDED table\_name and check type
- C. Run SQL command SHOW TABLES to see the type of the table
- D. All managed tables are stored in unity catalog
- E. All external tables are stored in data lake, managed tables are stored in DELTA lake

**Answer: B**

Explanation:

Explanation

The answer is Run SQL command DESCRIBE EXTENDED table\_name and check type Example of External table Graphical user interface, text, application Description automatically generated

□ Example of managed table

Graphical user interface, text, application, Teams Description automatically generated

□

### NEW QUESTION # 83

The data governance team is reviewing user for deleting records for compliance with GDPR. The following logic has been implemented to propagate deleted requests from the user\_lookup table to the user aggregate table.

□ Assuming that user\_id is a unique identifying key and that all users have requested deletion have been removed from the user\_lookup table, which statement describes whether successfully executing the above logic guarantees that the records to be deleted from the user\_aggregates table are no longer accessible and why?

- A. Yes: Delta Lake ACID guarantees provide assurance that the DELETE command succeeded fully and permanently purged these records.
- B. No: the Delta Lake DELETE command only provides ACID guarantees when combined with the MERGE INTO command
- C. No: the change data feed only tracks inserts and updates not deleted records.
- D. No: files containing deleted records may still be accessible with time travel until a VACUUM command is used to remove invalidated data files.

**Answer: D**

Explanation:

The DELETE operation in Delta Lake is ACID compliant, which means that once the operation is successful, the records are logically removed from the table. However, the underlying files that contained these records may still exist and be accessible via time travel to older versions of the table. To ensure that these records are physically removed and compliance with GDPR is maintained, a VACUUM command should be used to clean up these data files after a certain retention period. The VACUUM command will remove the files from the storage layer, and after this, the records will no longer be accessible.

### NEW QUESTION # 84

Where in the Spark UI can one diagnose a performance problem induced by not leveraging predicate push-down?

- A. In the Query Detail screen, by interpreting the Physical Plan
- B. In the Stage's Detail screen, in the Completed Stages table, by noting the size of data read from the Input column
- C. In the Executor's log file, by gripping for "predicate push-down"
- D. In the Storage Detail screen, by noting which RDDs are not stored on disk
- E. In the Delta Lake transaction log, by noting the column statistics

**Answer: A**

Explanation:

Explanation

This is the correct answer because it is where in the Spark UI one can diagnose a performance problem induced by not leveraging predicate push-down. Predicate push-down is an optimization technique that allows filtering data at the source before loading it into memory or processing it further. This can improve performance and reduce I/O costs by avoiding reading unnecessary data. To leverage predicate push-down, one should use supported data sources and formats, such as Delta Lake, Parquet, or JDBC, and use filter expressions that can be pushed down to the source. To diagnose a performance problem induced by not leveraging predicate push-down, one can use the Spark UI to access the Query Detail screen, which shows information about a SQL query executed on a Spark cluster. The Query Detail screen includes the Physical Plan, which is the actual plan executed by Spark to perform the query. The Physical Plan shows the physical operators used by Spark, such as Scan, Filter, Project, or Aggregate, and their input and output statistics, such as rows and bytes. By interpreting the Physical Plan, one can see if the filter expressions are pushed down to the source or not, and how much data is read or processed by each operator. Verified References: [Databricks Certified Data Engineer Professional], under "Spark Core" section; Databricks Documentation, under "Predicate pushdown" section; Databricks Documentation, under "Query detail page" section.

## NEW QUESTION # 85

Which of the following is a correct statement on how the data is organized in the storage when managing a DELTA table?

- A. All of the data is broken down into one or many parquet files, log file is removed once the transaction is committed.
- B. All of the data is broken down into one or many parquet files, but the log file is stored as a single json file, and every transaction creates a new data file(s) and log file gets appended.
- C. All of the data is stored into one parquet file, log files are broken down into one or many json files.
- D. All of the data and log are stored in a single parquet file
- E. All of the data is broken down into one or many parquet files, log files are broken down into one or many JSON files, and each transaction creates a new data file(s) and log file.

(Correct)

### Answer: E

Explanation:

Explanation

Answer is

All of the data is broken down into one or many parquet files, log files are broken down into one or many json files, and each transaction creates a new data file(s) and log file.

here is sample layout of how DELTA table might look,

□

## NEW QUESTION # 86

If you create a database sample\_db with the statement CREATE DATABASE sample\_db what will be the default location of the database in DBFS?

- A. Default location, DBFS/user/
- B. Default location, /user/db/
- C. Default Storage account
- D. Default Location, dbfs:/user/hive/warehouse
- E. Statement fails "Unable to create database without location"

### Answer: D

Explanation:

Explanation

The Answer is dbfs:/user/hive/warehouse this is the default location where spark stores user data-bases, the default can be changed using spark.sql.warehouse.dir a parameter. You can also provide a custom location using the LOCATION keyword.

Here is how this works,

Graphical user interface, text, application, email Description automatically generated

□ Default location

□ FYI, This can be changed used using cluster spark config or session config

Modify spark.sql.warehouse.dir location to change the default location

Graphical user interface, text, application Description automatically generated

□

## NEW QUESTION # 87

.....

**Databricks-Certified-Professional-Data-Engineer Exam Assessment:** <https://www.itdumpsf.com/Databricks-Certified-Professional-Data-Engineer-exam-passed.html>

- Databricks-Certified-Professional-Data-Engineer Valid Test Topics □ Dump Databricks-Certified-Professional-Data-Engineer Torrent □ Databricks-Certified-Professional-Data-Engineer Valid Exam Online □ Open website □ www.prep4sures.top □ and search for ➤ Databricks-Certified-Professional-Data-Engineer □ for free download □ Databricks-Certified-Professional-Data-Engineer Free Learning Cram
- Valid Download Databricks-Certified-Professional-Data-Engineer Free Dumps offer you accurate Exam Assessment | Databricks Certified Professional Data Engineer Exam □ Go to website ( www.pdfvce.com ) open and search for ➤ Databricks-Certified-Professional-Data-Engineer □ to download for free □ Databricks-Certified-Professional-Data-Engineer Free Learning Cram

