

# Databricks Certification Databricks-Certified-Professional-Data-Engineer exam pdf



As the captioned description said, our Databricks-Certified-Professional-Data-Engineer practice materials are filled with the newest points of knowledge about the exam. With many years of experience in this line, we not only compile real test content into our Databricks-Certified-Professional-Data-Engineer learning quiz, but the newest in to them. And our professionals always keep a close eye on the new changes of the subject and keep updating the Databricks-Certified-Professional-Data-Engineer study questions to the most accurate.

In general Dumpcollection Databricks-Certified-Professional-Data-Engineer exam simulator questions are practical, knowledge points are clear. According to candidates' replying, our exam questions contain most of real original test questions. You will not need to waste too much time on useless learning. Databricks-Certified-Professional-Data-Engineer Exam Simulator questions can help you understand key knowledge points and prepare easily and accordingly. Candidates should grasp this good opportunity to run into success clearly.

>> [Databricks-Certified-Professional-Data-Engineer Reliable Exam Answers](#) <<

## Pass4sure Databricks-Certified-Professional-Data-Engineer Exam Prep - Real Databricks-Certified-Professional-Data-Engineer Exam Dumps

Some candidates may considerate whether the Databricks-Certified-Professional-Data-Engineer exam guide is profession, but it can be sure that the contents of our study materials are compiled by industry experts after them refining the contents of textbooks, they have good knowledge of exam. Databricks-Certified-Professional-Data-Engineer test questions also has an automatic scoring function, giving you an objective rating after you take a mock exam to let you know your true level. With Databricks-Certified-Professional-Data-Engineer Exam Guide, you only need to spend 20-30 hours to study and you can successfully pass the exam. You will no longer worry about your exam because of bad study materials. If you decide to choose and practice our Databricks-

Certified-Professional-Data-Engineer test questions, our life will be even more exciting.

## **Databricks Certified Professional Data Engineer Exam Sample Questions (Q159-Q164):**

### **NEW QUESTION # 159**

A team of data engineer are adding tables to a DLT pipeline that contain repetitive expectations for many of the same data quality checks.

One member of the team suggests reusing these data quality rules across all tables defined for this pipeline.

What approach would allow them to do this?

- **A. Maintain data quality rules in a Delta table outside of this pipeline's target schema, providing the schema name as a pipeline parameter.**
- B. Use global Python variables to make expectations visible across DLT notebooks included in the same pipeline.
- C. Maintain data quality rules in a separate Databricks notebook that each DLT notebook of file.
- D. Add data quality constraints to tables in this pipeline using an external job with access to pipeline configuration files.

**Answer: A**

Explanation:

Maintaining data quality rules in a centralized Delta table allows for the reuse of these rules across multiple DLT (Delta Live Tables) pipelines. By storing these rules outside the pipeline's target schema and referencing the schema name as a pipeline parameter, the team can apply the same set of data quality checks to different tables within the pipeline. This approach ensures consistency in data quality validations and reduces redundancy in code by not having to replicate the same rules in each DLT notebook or file.

Reference:

Databricks Documentation on Delta Live Tables: [Delta Live Tables Guide](#)

### **NEW QUESTION # 160**

A member of the data engineering team has submitted a short notebook that they wish to schedule as part of a larger data pipeline. Assume that the commands provided below produce the logically correct results when run as presented.



```
Cmd 1
rawDF = spark.table("raw_data")

Cmd 2
rawDF.printSchema()

Cmd 3
flattenedDF = rawDF.select("*", "values.*")

Cmd 4
finalDF = flattenedDF.drop("values")

Cmd 5
finalDF.explain()

Cmd 6
display(finalDF)

Cmd 7
finalDF.write.mode("append").saveAsTable("flat_data")
```

Which command should be removed from the notebook before scheduling it as a job?

- A. Cmd 5
- **B. Cmd 6**
- C. Cmd 4
- D. Cmd 3
- E. Cmd 2

**Answer: B**

Explanation:

Cmd 6 is the command that should be removed from the notebook before scheduling it as a job. This command is selecting all the columns from the finalDF dataframe and displaying them in the notebook. This is not necessary for the job, as the finalDF dataframe is already written to a table in Cmd 7. Displaying the dataframe in the notebook will only consume resources and time, and it will not affect the output of the job. Therefore, Cmd 6 is redundant and should be removed.

The other commands are essential for the job, as they perform the following tasks:

Cmd 1: Reads the raw\_data table into a Spark dataframe called rawDF.

Cmd 2: Prints the schema of the rawDF dataframe, which is useful for debugging and understanding the data structure.

Cmd 3: Selects all the columns from the rawDF dataframe, as well as the nested columns from the values struct column, and creates a new dataframe called flattenedDF.

Cmd 4: Drops the values column from the flattenedDF dataframe, as it is no longer needed after flattening, and creates a new dataframe called finalDF.

Cmd 5: Explains the physical plan of the finalDF dataframe, which is useful for optimizing and tuning the performance of the job.

Cmd 7: Writes the finalDF dataframe to a table called flat\_data, using the append mode to add new data to the existing table.

#### NEW QUESTION # 161

Which of the following SQL statements can be used to update a transactions table, to set a flag on the table from Y to N

- A. UPDATE transactions SET active\_flag = 'N' WHERE active\_flag = 'Y'
- B. MERGE transactions SET active\_flag = 'N' WHERE active\_flag = 'Y'
- C. MODIFY transactions SET active\_flag = 'N' WHERE active\_flag = 'Y'
- D. REPLACE transactions SET active\_flag = 'N' WHERE active\_flag = 'Y'

**Answer: D**

Explanation:

Explanation

The answer is

UPDATE transactions SET active\_flag = 'N' WHERE active\_flag = 'Y'

Delta Lake supports UPDATE statements on the delta table, all of the changes as part of the update are ACID compliant.

### NEW QUESTION # 162

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 a job with at least one wide transformation, which of the following cluster configurations will result in maximum performance?

- A. \* Total VMs: 4  
\* 100 GB per Executor  
\* 40 Cores/Executor
- B. \* Total VMs; 1  
\* 400 GB per Executor  
\* 160 Cores / Executor
- C. \* Total VMs:2  
\* 200 GB per Executor  
\* 80 Cores / Executor
- D. \* Total VMs: 8  
\* 50 GB per Executor  
\* 20 Cores / Executor

**Answer: D**

Explanation:

This is the correct answer because it is the cluster configuration that will result in maximum performance for a job with at least one wide transformation. A wide transformation is a type of transformation that requires shuffling data across partitions, such as join, groupBy, or orderBy. Shuffling can be expensive and time-consuming, especially if there are too many or too few partitions. Therefore, it is important to choose a cluster configuration that can balance the trade-off between parallelism and network overhead. In this case, having 8 VMs with 50 GB per executor and 20 cores per executor will create 8 partitions, each with enough memory and CPU resources to handle the shuffling efficiently. Having fewer VMs with more memory and cores per executor will create fewer partitions, which will reduce parallelism and increase the size of each shuffle block. Having more VMs with less memory and cores per executor will create more partitions, which will increase parallelism but also increase the network overhead and the number of shuffle files. Verified References: [Databricks Certified Data Engineer Professional], under "Performance Tuning" section; Databricks Documentation, under "Cluster configurations" section.

### NEW QUESTION # 163

A production cluster has 3 executor nodes and uses the same virtual machine type for the driver and executor.

When evaluating the Ganglia Metrics for this cluster, which indicator would signal a bottleneck caused by code executing on the driver?

- A. Overall cluster CPU utilization is around 25%
- B. Bytes Received never exceeds 80 million bytes per second
- C. The five Minute Load Average remains consistent/flat
- D. Network I/O never spikes
- E. Total Disk Space remains constant

**Answer: A**

Explanation:

This is the correct answer because it indicates a bottleneck caused by code executing on the driver. A bottleneck is a situation where the performance or capacity of a system is limited by a single component or resource. A bottleneck can cause slow execution, high latency, or low throughput. A production cluster has 3 executor nodes and uses the same virtual machine type for the driver and executor. When evaluating the Ganglia Metrics for this cluster, one can look for indicators that show how the cluster resources are being utilized, such as CPU, memory, disk, or network. If the overall cluster CPU utilization is around 25%, it means that only one out of the four nodes (driver + 3 executors) is using its full CPU capacity, while the other three nodes are idle or underutilized. This suggests that the code executing on the driver is taking too long or consuming too much CPU resources, preventing the executors from receiving tasks or data to process. This can happen when the code has driver-side operations that are not parallelized or distributed, such as collecting large amounts of data to the driver, performing complex calculations on the driver, or using non-Spark libraries on the driver. Verified References: [Databricks Certified Data Engineer Professional], under "Spark Core" section; Databricks Documentation, under "View cluster status and event logs - Ganglia metrics" section; Databricks Documentation, under "Avoid collecting large RDDs" section.

In a Spark cluster, the driver node is responsible for managing the execution of the Spark application, including scheduling tasks, managing the execution plan, and interacting with the cluster manager. If the overall cluster CPU utilization is low (e.g., around 25%), it may indicate that the driver node is not utilizing the available resources effectively and might be a bottleneck.

## NEW QUESTION # 164

.....

In the major environment, people are facing more job pressure. So they want to get Databricks-Certified-Professional-Data-Engineer certification rise above the common herd. How to choose valid and efficient Databricks-Certified-Professional-Data-Engineer guide torrent should be the key topic most candidates may concern. So now, it is right, you come to us. Our company is famous for its high-quality in this field especially for Databricks-Certified-Professional-Data-Engineer Certification exams. It has been accepted by thousands of candidates who practice our study materials for their exam.

**Pass4sure Databricks-Certified-Professional-Data-Engineer Exam Prep:** [https://www.dumpcollection.com/Databricks-Certified-Professional-Data-Engineer\\_braindumps.html](https://www.dumpcollection.com/Databricks-Certified-Professional-Data-Engineer_braindumps.html)

When you search the Databricks-Certified-Professional-Data-Engineer study material on the internet, you will find many site which are related to Databricks-Certified-Professional-Data-Engineer actual test, Databricks-Certified-Professional-Data-Engineer exam cram is helpful for candidates who are urgent for Databricks-Certified-Professional-Data-Engineer certifications, Databricks Databricks-Certified-Professional-Data-Engineer Reliable Exam Answers Challenge is ubiquitous, only by constant and ceaseless effort, can you be the man you want to be, Databricks Databricks-Certified-Professional-Data-Engineer Reliable Exam Answers Or you can wait the updating or free change to other dumps if you have other test.

This lets Apple sell Intel Macs now, while giving developers Databricks-Certified-Professional-Data-Engineer time to gear up before they have pro customers demanding universal binaries at the top of their lungs.

An invasion of armies can be resisted, but not an idea whose time has come, When you search the Databricks-Certified-Professional-Data-Engineer Study Material on the internet, you will find many site which are related to Databricks-Certified-Professional-Data-Engineer actual test.

## **New Databricks-Certified-Professional-Data-Engineer Reliable Exam Answers | Reliable Databricks Pass4sure Databricks-Certified-Professional-Data-Engineer Exam Prep: Databricks Certified Professional Data Engineer Exam**

Databricks-Certified-Professional-Data-Engineer exam cram is helpful for candidates who are urgent for Databricks-Certified-Professional-Data-Engineer certifications, Challenge is ubiquitous, only by constant and ceaseless effort, can you be the man you want to be.

Or you can wait the updating or free change to other dumps if you have other Real Databricks-Certified-Professional-Data-Engineer Exam Dumps test, The Databricks Certified Professional Data Engineer Exam has three formats so that the students don't face any serious problems and prepare themselves with fully focused minds.

- Online Databricks-Certified-Professional-Data-Engineer Test  Latest Databricks-Certified-Professional-Data-Engineer Real Test  Exam Databricks-Certified-Professional-Data-Engineer Book  Immediately open  [www.practicevce.com](http://www.practicevce.com)  and search for  Databricks-Certified-Professional-Data-Engineer  to obtain a free download  Databricks-Certified-Professional-Data-Engineer Reliable Dumps Sheet

