

# Databricks Databricks-Certified-Professional-Data-Engineer PDF Questions [2026] To Gain Brilliant Result



What's more, part of that ExamsLabs Databricks-Certified-Professional-Data-Engineer dumps now are free:  
<https://drive.google.com/open?id=1ONtgrg4F6cwX9ErWr4KVxsmpPQiWJoyx>

If you would like to use all kinds of electronic devices to prepare for the Databricks-Certified-Professional-Data-Engineer exam, then I am glad to tell you that our online app version is definitely your perfect choice. With the online app version of our study materials, you can just feel free to practice the questions in our Databricks-Certified-Professional-Data-Engineer Training Materials no matter you are using your mobile phone, personal computer, or tablet PC. In addition, another strong point of the online app version of our Databricks-Certified-Professional-Data-Engineer learning guide is that it is convenient for you to use even though you are in offline environment.

Databricks Certified Professional Data Engineer exam is a hands-on exam that requires the candidate to complete a set of tasks using Databricks. Databricks-Certified-Professional-Data-Engineer exam evaluates the candidate's ability to design and implement data pipelines, work with data sources and sinks, and perform transformations using Databricks. Databricks-Certified-Professional-Data-Engineer Exam also tests the candidate's ability to optimize and tune data pipelines for performance and reliability.

>> [Test Databricks-Certified-Professional-Data-Engineer Simulator Online](#) <<

## Test Databricks-Certified-Professional-Data-Engineer Simulator Online offer you accurate Exam Introduction to pass Databricks Databricks Certified Professional Data Engineer Exam exam

It is our consistent aim to serve our customers wholeheartedly. Our Databricks-Certified-Professional-Data-Engineer study materials try to ensure that every customer is satisfied, which can be embodied in the convenient and quick refund process. Although the passing rate of our Databricks-Certified-Professional-Data-Engineer Study Materials is close to 100 %, if you are still worried, we can give you another guarantee: if you don't pass the exam, you can get a full refund. Yes, this is the truth.

## Databricks Certified Professional Data Engineer Exam Sample Questions (Q95-Q100):

### NEW QUESTION # 95

You are currently asked to work on building a data pipeline, you have noticed that you are currently working with a data source that has a lot of data quality issues and you need to monitor data quality and enforce it as part of the data ingestion process, which of the following tools can be used to address this problem?

- A. DELTA LIVE TABLES
- B. STRUCTURED STREAMING with MULTI HOP
- C. UNITY Catalog and Data Governance
- D. JOBS and TASKS
- E. AUTO LOADER

**Answer: A**

Explanation:

Explanation

The answer is, DELTA LIVE TABLES

Delta live tables expectations can be used to identify and quarantine bad data, all of the data quality metrics are stored in the event logs which can be used to later analyze and monitor.

DELTA LIVE Tables expectations

Below are three types of expectations, make sure to pay attention differences between these three.

Retain invalid records:

Use the expect operator when you want to keep records that violate the expectation. Records that violate the expectation are added to the target dataset along with valid records:

Python

```
1.@dlt.expect("valid timestamp", "col('timestamp') > '2012-01-01'")
```

SQL

```
1.CONSTRAINT valid_timestamp EXPECT (timestamp > '2012-01-01')
```

Drop invalid records:

Use the expect or drop operator to prevent the processing of invalid records. Records that violate the expectation are dropped from the target dataset:

Python

```
1.@dlt.expect_or_drop("valid_current_page", "current_page_id IS NOT NULL AND current_page_title IS NOT NULL") SQL
```

```
1.CONSTRAINT valid_current_page EXPECT (current_page_id IS NOT NULL and current_page_title IS NOT NULL) ON
```

```
VIOLATION DROP ROW Fail on invalid records:
```

When invalid records are unacceptable, use the expect or fail operator to halt execution immediately when a record fails validation.

If the operation is a table update, the system atomically rolls back the transaction:

Python

```
1.@dlt.expect_or_fail("valid_count", "count > 0")
```

SQL

```
1.CONSTRAINT valid_count EXPECT (count > 0) ON VIOLATION FAIL UPDATE
```

## NEW QUESTION # 96

Which statement describes Delta Lake Auto Compaction?

- **A. An asynchronous job runs after the write completes to detect if files could be further compacted; if yes, an optimize job is executed toward a default of 128 MB.**
- B. Before a Jobs cluster terminates, optimize is executed on all tables modified during the most recent job.
- C. Data is queued in a messaging bus instead of committing data directly to memory; all data is committed from the messaging bus in one batch once the job is complete.
- D. Optimized writes use logical partitions instead of directory partitions; because partition boundaries are only represented in metadata, fewer small files are written.
- E. An asynchronous job runs after the write completes to detect if files could be further compacted; if yes, an optimize job is executed toward a default of 1 GB.

**Answer: A**

Explanation:

This is the correct answer because it describes the behavior of Delta Lake Auto Compaction, which is a feature that automatically optimizes the layout of Delta Lake tables by coalescing small files into larger ones.

Auto Compaction runs as an asynchronous job after a write to a table has succeeded and checks if files within a partition can be further compacted. If yes, it runs an optimize job with a default target file size of 128 MB.

Auto Compaction only compacts files that have not been compacted previously. Verified References:

[Databricks Certified Data Engineer Professional], under "Delta Lake" section; Databricks Documentation, under "Auto Compaction for Delta Lake on Databricks" section.

"Auto compaction occurs after a write to a table has succeeded and runs synchronously on the cluster that has performed the write.

Auto compaction only compacts files that haven't been compacted previously."

<https://learn.microsoft.com/en-us/azure/databricks/delta/tune-file-size>

## NEW QUESTION # 97

Assuming that the Databricks CLI has been installed and configured correctly, which Databricks CLI command can be used to upload a custom Python Wheel to object storage mounted with the DBFS for use with a production job?

- A. workspace
- B. fs
- C. configure
- **D. libraries**
- E. jobs

**Answer: D**

Explanation:

The libraries command group allows you to install, uninstall, and list libraries on Databricks clusters. You can use the libraries install command to install a custom Python Wheel on a cluster by specifying the --whl option and the path to the wheel file. For example, you can use the following command to install a custom Python Wheel named mylib-0.1-py3-none-any.whl on a cluster with the id 1234-567890-abcde123:

```
databricks libraries install --cluster-id 1234-567890-abcde123 --whl
dbfs/mnt/mylib/mylib-0.1-py3-none-any.whl
```

This will upload the custom Python Wheel to the cluster and make it available for use with a production job.

You can also use the libraries uninstall command to uninstall a library from a cluster, and the libraries list command to list the libraries installed on a cluster.

References:

- \* Libraries CLI (legacy): <https://docs.databricks.com/en/archive/dev-tools/cli/libraries-cli.html>
- \* Library operations: <https://docs.databricks.com/en/dev-tools/cli/commands.html#library-operations>
- \* Install or update the Databricks CLI: <https://docs.databricks.com/en/dev-tools/cli/install.html>

### NEW QUESTION # 98

The following code has been migrated to a Databricks notebook from a legacy workload:

The code executes successfully and provides the logically correct results, however, it takes over 20 minutes to extract and load around 1 GB of data.

Which statement is a possible explanation for this behavior?

- A. Instead of cloning, the code should use %sh pip install so that the Python code can get executed in parallel across all nodes in a cluster.
- B. %sh does not distribute file moving operations; the final line of code should be updated to use %fs instead.
- C. Python will always execute slower than Scala on Databricks. The run.py script should be refactored to Scala.
- **D. %sh executes shell code on the driver node. The code does not take advantage of the worker nodes or Databricks optimized Spark.**
- E. %sh triggers a cluster restart to collect and install Git. Most of the latency is related to cluster startup time.

**Answer: D**

Explanation:

<https://www.databricks.com/blog/2020/08/31/introducing-the-databricks-web-terminal.html> The code is using %sh to execute shell code on the driver node. This means that the code is not taking advantage of the worker nodes or Databricks optimized Spark. This is why the code is taking longer to execute. A better approach would be to use Databricks libraries and APIs to read and write data from Git and DBFS, and to leverage the parallelism and performance of Spark. For example, you can use the Databricks Connect feature to run your Python code on a remote Databricks cluster, or you can use the Spark Git Connector to read data from Git repositories as Spark DataFrames.

### NEW QUESTION # 99

The data engineering team maintains the following code:

Assuming that this code produces logically correct results and the data in the source tables has been de-duplicated and validated, which statement describes what will occur when this code is executed?

- A. No computation will occur until enriched\_itemized\_orders\_by\_account is queried; upon query materialization, results will be calculated using the current valid version of data in each of the three tables referenced in the join logic.
- **B. The enriched\_itemized\_orders\_by\_account table will be overwritten using the current valid version of data in each of the three tables referenced in the join logic.**

- C. An incremental job will leverage information in the state store to identify unjoined rows in the source tables and write these rows to the enriched\_itemized\_orders\_by\_account table.
- D. An incremental job will detect if new rows have been written to any of the source tables; if new rows are detected, all results will be recalculated and used to overwrite the enriched\_itemized\_orders\_by\_account table.
- E. A batch job will update the enriched\_itemized\_orders\_by\_account table, replacing only those rows that have different values than the current version of the table, using accountID as the primary key.

**Answer: B**

Explanation:

The provided PySpark code performs the following operations:

- \* Reads Data from silver\_customer\_sales Table:
- \* The code starts by accessing the silver\_customer\_sales table using the spark.table method.
- \* Groups Data by customer\_id:
- \* The .groupBy("customer\_id") function groups the data based on the customer\_id column.
- \* Aggregates Data:
- \* The .agg() function computes several aggregate metrics for each customer\_id:
- \* F.min("sale\_date").alias("first\_transaction\_date"): Determines the earliest sale date for the customer.
- \* F.max("sale\_date").alias("last\_transaction\_date"): Determines the latest sale date for the customer.
- \* F.mean("sale\_total").alias("average\_sales"): Calculates the average sale amount for the customer.
- \* F.countDistinct("order\_id").alias("total\_orders"): Counts the number of unique orders placed by the customer.
- \* F.sum("sale\_total").alias("lifetime\_value"): Calculates the total sales amount (lifetime value) for the customer.
- \* Writes Data to gold\_customer\_lifetime\_sales\_summary Table:
- \* The .write.mode("overwrite").table("gold\_customer\_lifetime\_sales\_summary") command writes the aggregated data to the gold\_customer\_lifetime\_sales\_summary table.
- \* The mode("overwrite") specifies that the existing data in the gold\_customer\_lifetime\_sales\_summary table will be completely replaced by the new aggregated data.

Conclusion:

When this code is executed, it reads all records from the silver\_customer\_sales table, performs the specified aggregations grouped by customer\_id, and then overwrites the entire gold\_customer\_lifetime\_sales\_summary table with the aggregated results. Therefore, option D accurately describes this process: "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." References:

PySpark DataFrame groupBy  
PySpark Basics

## NEW QUESTION # 100

.....

ExamsLabs is the preeminent platform, which offers Databricks-Certified-Professional-Data-Engineer exam materials duly equipped by experts. If you want you spend least time getting the best result, our exam materials must be your best choice. Our Databricks-Certified-Professional-Data-Engineer exam materials are best suited to busy specialized who can learn in their seemly timings. Our study materials have satisfied in PDF format which can certainly be retrieved on all the digital devices. You can install it in your smartphone, Laptop or Tables to use. What most useful is that PDF format of our Databricks-Certified-Professional-Data-Engineer Exam Materials can be printed easily, you can learn it everywhere and every time you like. It is really convenient for candidates who are busy to prepare the exam. You can save so much time and energy to do other things that you will make best use of you time.

**Exam Databricks-Certified-Professional-Data-Engineer Introduction:** <https://www.examslabs.com/Databricks/Databricks-Certification/best-Databricks-Certified-Professional-Data-Engineer-exam-dumps.html>

- Pass Guaranteed Quiz Accurate Databricks - Databricks-Certified-Professional-Data-Engineer - Test Databricks Certified Professional Data Engineer Exam Simulator Online  Search for **>** Databricks-Certified-Professional-Data-Engineer  on **➔** [www.dumpsmaterials.com](http://www.dumpsmaterials.com)  immediately to obtain a free download  Databricks-Certified-Professional-Data-Engineer New Question
- Pass Guaranteed Quiz Databricks - Databricks-Certified-Professional-Data-Engineer -Professional Test Simulator Online   Open **➔** [www.pdfvce.com](http://www.pdfvce.com)  and search for **▶** Databricks-Certified-Professional-Data-Engineer **◀** to download exam materials for free  Valid Databricks-Certified-Professional-Data-Engineer Test Questions
- Reliable Databricks-Certified-Professional-Data-Engineer Exam Prep  Databricks-Certified-Professional-Data-Engineer Latest Test Prep  Databricks-Certified-Professional-Data-Engineer Exam Quiz  Download  Databricks-Certified-Professional-Data-Engineer  for free by simply entering **➔** [www.examcollectionpass.com](http://www.examcollectionpass.com)  website  Valid Databricks-Certified-Professional-Data-Engineer Test Questions

