



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 1

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data organization leader is upset about the data analysis team's reports being different from the data engineering team's reports. The leader believes the siloed nature of their organization's data engineering and data analysis architectures is to blame.

Which of the following describes how a data lakehouse could alleviate this issue?

- A. Both teams would autoscale their work as data size evolves
- B. Both teams would use the same source of truth for their work
- C. Both teams would reorganize to report to the same department
- D. Both teams would be able to collaborate on projects in real-time
- E. Both teams would respond more quickly to ad-hoc requests

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 2

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following describes a scenario in which a data team will want to utilize cluster pools?

- A. An automated report needs to be refreshed as quickly as possible.
- B. An automated report needs to be made reproducible.
- C. An automated report needs to be tested to identify errors.
- D. An automated report needs to be version-controlled across multiple collaborators.
- E. An automated report needs to be runnable by all stakeholders.

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 3

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following is hosted completely in the control plane of the classic Databricks architecture?

- A. Worker node
- B. JDBC data source
- C. Databricks web application
- D. Databricks Filesystem
- E. Driver node

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 4

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following benefits of using the Databricks Lakehouse Platform is provided by Delta Lake?

- A. The ability to manipulate the same data using a variety of languages
- B. The ability to collaborate in real time on a single notebook
- C. The ability to set up alerts for query failures
- D. The ability to support batch and streaming workloads
- E. The ability to distribute complex data operations

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 5

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following describes the storage organization of a Delta table?

- A. Delta tables are stored in a single file that contains data, history, metadata, and other attributes.
- B. Delta tables store their data in a single file and all metadata in a collection of files in a separate location.
- C. Delta tables are stored in a collection of files that contain data, history, metadata, and other attributes.
- D. Delta tables are stored in a collection of files that contain only the data stored within the table.
- E. Delta tables are stored in a single file that contains only the data stored within the table.

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 6

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following code blocks will remove the rows where the value in column age is greater than 25 from the existing Delta table my_table and save the updated table?

- A. `SELECT * FROM my_table WHERE age > 25;`
- B. `UPDATE my_table WHERE age > 25;`
- C. `DELETE FROM my_table WHERE age > 25;`
- D. `UPDATE my_table WHERE age <= 25;`
- E. `DELETE FROM my_table WHERE age <= 25;`

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 7

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer has realized that they made a mistake when making a daily update to a table. They need to use Delta time travel to restore the table to a version that is 3 days old. However, when the data engineer attempts to time travel to the older version, they are unable to restore the data because the data files have been deleted.

Which of the following explains why the data files are no longer present?

- A. The VACUUM command was run on the table
- B. The TIME TRAVEL command was run on the table
- C. The DELETE HISTORY command was run on the table
- D. The OPTIMIZE command was run on the table
- E. The HISTORY command was run on the table

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 8

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following Git operations must be performed outside of Databricks Repos?

- A. Commit
- B. Pull
- C. Push
- D. Clone
- E. Merge

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 9

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following data lakehouse features results in improved data quality over a traditional data lake?

- A. A data lakehouse provides storage solutions for structured and unstructured data.
- B. A data lakehouse supports ACID-compliant transactions.
- C. A data lakehouse allows the use of SQL queries to examine data.
- D. A data lakehouse stores data in open formats.
- E. A data lakehouse enables machine learning and artificial Intelligence workloads.

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 10

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer needs to determine whether to use the built-in Databricks Notebooks versioning or version their project using Databricks Repos. Which of the following is an advantage of using Databricks Repos over the Databricks Notebooks versioning?

- A. Databricks Repos automatically saves development progress
- B. Databricks Repos supports the use of multiple branches
- C. Databricks Repos allows users to revert to previous versions of a notebook
- D. Databricks Repos provides the ability to comment on specific changes
- E. Databricks Repos is wholly housed within the Databricks Lakehouse Platform

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 11

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer has left the organization. The data team needs to transfer ownership of the data engineer's Delta tables to a new data engineer. The new data engineer is the lead engineer on the data team.

Assuming the original data engineer no longer has access, which of the following individuals must be the one to transfer ownership of the Delta tables in Data Explorer?

- A. Databricks account representative
- B. This transfer is not possible
- C. Workspace administrator
- D. New lead data engineer
- E. Original data engineer

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 12

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data analyst has created a Delta table sales that is used by the entire data analysis team. They want help from the data engineering team to implement a series of tests to ensure the data is clean. However, the data engineering team uses Python for its tests rather than SQL.

Which of the following commands could the data engineering team use to access sales in PySpark?

- A. `SELECT * FROM sales`
- B. There is no way to share data between PySpark and SQL.
- C. `spark.sql("sales")`
- D. `spark.delta.table("sales")`
- E. `spark.table("sales")`

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 13

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following commands will return the location of database customer360?

- A. DESCRIBE LOCATION customer360;
- B. DROP DATABASE customer360;
- C. DESCRIBE DATABASE customer360;
- D. ALTER DATABASE customer360 SET DBPROPERTIES ('location' = '/user');
- E. USE DATABASE customer360;

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 14

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer wants to create a new table containing the names of customers that live in France.

They have written the following command:

```
CREATE TABLE customersInFrance
  AS
SELECT id,
       firstName,
       lastName,
FROM customerLocations
WHERE country = 'FRANCE';
```

A senior data engineer mentions that it is organization policy to include a table property indicating that the new table includes personally identifiable information (PII).

Which of the following lines of code fills in the above blank to successfully complete the task?

- A. There is no way to indicate whether a table contains PII.
- B. "COMMENT PII"
- C. TBLPROPERTIES PII
- D. COMMENT "Contains PII"
- E. PII

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 15

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following benefits is provided by the array functions from Spark SQL?

- A. An ability to work with data in a variety of types at once
- B. An ability to work with data within certain partitions and windows
- C. An ability to work with time-related data in specified intervals
- D. An ability to work with complex, nested data ingested from JSON files
- E. An ability to work with an array of tables for procedural automation

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 16

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following commands can be used to write data into a Delta table while avoiding the writing of duplicate records?

- A. DROP
- B. IGNORE
- C. MERGE
- D. APPEND
- E. INSERT

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 17

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer needs to apply custom logic to string column city in table stores for a specific use case. In order to apply this custom logic at scale, the data engineer wants to create a SQL user-defined function (UDF).

Which of the following code blocks creates this SQL UDF?

A.

```
CREATE FUNCTION combine_nyc(city STRING)
  RETURNS STRING
  RETURN CASE
    WHEN city = "brooklyn" THEN "new york"
    ELSE city
  END;
```

B.

```
CREATE UDF combine_nyc(city STRING)
  RETURNS STRING
  CASE
    WHEN city = "brooklyn" THEN "new york"
    ELSE city
  END;
```

C.

```
CREATE UDF combine_nyc(city STRING)
  RETURN CASE
    WHEN city = "brooklyn" THEN "new york"
    ELSE city
  END;
```

D.

```
CREATE FUNCTION combine_nyc(city STRING)
  RETURN CASE
    WHEN city = "brooklyn" THEN "new york"
    ELSE city
  END;
```

E.

```
CREATE UDF combine_nyc(city STRING)
  RETURNS STRING
  RETURN CASE
    WHEN city = "brooklyn" THEN "new york"
    ELSE city
  END;
```

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 18

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data analyst has a series of queries in a SQL program. The data analyst wants this program to run every day. They only want the final query in the program to run on Sundays. They ask for help from the data engineering team to complete this task.

Which of the following approaches could be used by the data engineering team to complete this task?

- A. They could submit a feature request with Databricks to add this functionality.
- B. They could wrap the queries using PySpark and use Python's control flow system to determine when to run the final query.
- C. They could only run the entire program on Sundays.
- D. They could automatically restrict access to the source table in the final query so that it is only accessible on Sundays.
- E. They could redesign the data model to separate the data used in the final query into a new table.

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 19

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer runs a statement every day to copy the previous day's sales into the table transactions. Each day's sales are in their own file in the location `"/transactions/raw"`.

Today, the data engineer runs the following command to complete this task:

```
COPY INTO transactions
FROM "/transactions/raw"
FILEFORMAT = PARQUET;
```

After running the command today, the data engineer notices that the number of records in table transactions has not changed.

Which of the following describes why the statement might not have copied any new records into the table?

- A. The format of the files to be copied were not included with the `FORMAT_OPTIONS` keyword.
- B. The names of the files to be copied were not included with the `FILES` keyword.
- C. The previous day's file has already been copied into the table.
- D. The `PARQUET` file format does not support `COPY INTO`.
- E. The `COPY INTO` statement requires the table to be refreshed to view the copied rows.

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 20

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer needs to create a table in Databricks using data from their organization's existing SQLite database.

They run the following command:

```
CREATE TABLE jdbc_customer360
USING _____
OPTIONS (
  url "jdbc:sqlite:/customers.db",
  dbtable "customer360"
)
```

Which of the following lines of code fills in the above blank to successfully complete the task?

- A. org.apache.spark.sql.jdbc
- B. autoloader
- C. DELTA
- D. sqlite
- E. org.apache.spark.sql.sqlite

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 21

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineering team has two tables. The first table `march_transactions` is a collection of all retail transactions in the month of March. The second table `april_transactions` is a collection of all retail transactions in the month of April. There are no duplicate records between the tables.

Which of the following commands should be run to create a new table `all_transactions` that contains all records from `march_transactions` and `april_transactions` without duplicate records?

- A.

```
CREATE TABLE all_transactions AS
SELECT * FROM march_transactions
INNER JOIN SELECT * FROM april_transactions;
```
- B.

```
CREATE TABLE all_transactions AS
SELECT * FROM march_transactions
UNION SELECT * FROM april_transactions;
```
- C.

```
CREATE TABLE all_transactions AS
SELECT * FROM march_transactions
OUTER JOIN SELECT * FROM april_transactions;
```
- D.

```
CREATE TABLE all_transactions AS
SELECT * FROM march_transactions
INTERSECT SELECT * FROM april_transactions;
```
- E.

```
CREATE TABLE all_transactions AS
SELECT * FROM march_transactions
MERGE SELECT * FROM april_transactions;
```

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 22

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer only wants to execute the final block of a Python program if the Python variable `day_of_week` is equal to 1 and the Python variable `review_period` is True. Which of the following control flow statements should the data engineer use to begin this conditionally executed code block?

- A. `if day_of_week = 1 and review_period:`
- B. `if day_of_week = 1 and review_period = "True":`
- C. `if day_of_week == 1 and review_period == "True":`
- D. `if day_of_week == 1 and review_period:`
- E. `if day_of_week = 1 & review_period: = "True":`

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 23

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer is attempting to drop a Spark SQL table `my_table`. The data engineer wants to delete all table metadata and data.

They run the following command:

```
DROP TABLE IF EXISTS my_table -
```

While the object no longer appears when they run `SHOW TABLES`, the data files still exist.

Which of the following describes why the data files still exist and the metadata files were deleted?

- A. The table's data was larger than 10 GB
- B. The table's data was smaller than 10 GB
- C. The table was external
- D. The table did not have a location
- E. The table was managed

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 24

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer wants to create a data entity from a couple of tables. The data entity must be used by other data engineers in other sessions. It also must be saved to a physical location.

Which of the following data entities should the data engineer create?

- A. Database
- B. Function
- C. View
- D. Temporary view
- E. Table

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 25

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer is maintaining a data pipeline. Upon data ingestion, the data engineer notices that the source data is starting to have a lower level of quality. The data engineer would like to automate the process of monitoring the quality level.

Which of the following tools can the data engineer use to solve this problem?

- A. Unity Catalog
- B. Data Explorer
- C. Delta Lake
- D. Delta Live Tables
- E. Auto Loader

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 26

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A Delta Live Table pipeline includes two datasets defined using STREAMING LIVE TABLE. Three datasets are defined against Delta Lake table sources using LIVE TABLE. The table is configured to run in Production mode using the Continuous Pipeline Mode.

Assuming previously unprocessed data exists and all definitions are valid, what is the expected outcome after clicking Start to update the pipeline?

- A. All datasets will be updated at set intervals until the pipeline is shut down. The compute resources will persist to allow for additional testing.
- B. All datasets will be updated once and the pipeline will persist without any processing. The compute resources will persist but go unused.
- C. All datasets will be updated at set intervals until the pipeline is shut down. The compute resources will be deployed for the update and terminated when the pipeline is stopped.
- D. All datasets will be updated once and the pipeline will shut down. The compute resources will be terminated.
- E. All datasets will be updated once and the pipeline will shut down. The compute resources will persist to allow for additional testing.

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 27

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

In order for Structured Streaming to reliably track the exact progress of the processing so that it can handle any kind of failure by restarting and/or reprocessing, which of the following two approaches is used by Spark to record the offset range of the data being processed in each trigger?

- A. Checkpointing and Write-ahead Logs
- B. Structured Streaming cannot record the offset range of the data being processed in each trigger.
- C. Replayable Sources and Idempotent Sinks
- D. Write-ahead Logs and Idempotent Sinks
- E. Checkpointing and Idempotent Sinks

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 28

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following describes the relationship between Gold tables and Silver tables?

- A. Gold tables are more likely to contain aggregations than Silver tables.
- B. Gold tables are more likely to contain valuable data than Silver tables.
- C. Gold tables are more likely to contain a less refined view of data than Silver tables.
- D. Gold tables are more likely to contain more data than Silver tables.
- E. Gold tables are more likely to contain truthful data than Silver tables.

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 29

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following describes the relationship between Bronze tables and raw data?

- A. Bronze tables contain less data than raw data files.
- B. Bronze tables contain more truthful data than raw data.
- C. Bronze tables contain aggregates while raw data is unaggregated.
- D. Bronze tables contain a less refined view of data than raw data.
- E. Bronze tables contain raw data with a schema applied.

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 30

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following tools is used by Auto Loader process data incrementally?

- A. Checkpointing
- B. Spark Structured Streaming
- C. Data Explorer
- D. Unity Catalog
- E. Databricks SQL

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 31

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer has configured a Structured Streaming job to read from a table, manipulate the data, and then perform a streaming write into a new table. The code block used by the data engineer is below:

```
(spark.table("sales")
  .withColumn("avg_price", col("sales") / col("units"))
  .writeStream
  .option("checkpointLocation", checkpointPath)
  .outputMode("complete")
  ._____
  .table("new_sales")
)
```

If the data engineer only wants the query to execute a micro-batch to process data every 5 seconds, which of the following lines of code should the data engineer use to fill in the blank?

- A. trigger("5 seconds")
- B. trigger()
- C. trigger(once="5 seconds")
- D. trigger(processingTime="5 seconds")
- E. trigger(continuous="5 seconds")

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 32

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A dataset has been defined using Delta Live Tables and includes an expectations clause:

```
CONSTRAINT valid_timestamp EXPECT (timestamp > '2020-01-01') ON VIOLATION DROP ROW
```

What is the expected behavior when a batch of data containing data that violates these constraints is processed?

- A. Records that violate the expectation are dropped from the target dataset and loaded into a quarantine table.
- B. Records that violate the expectation are added to the target dataset and flagged as invalid in a field added to the target dataset.
- C. Records that violate the expectation are dropped from the target dataset and recorded as invalid in the event log.
- D. Records that violate the expectation are added to the target dataset and recorded as invalid in the event log.
- E. Records that violate the expectation cause the job to fail.

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 33

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following describes when to use the CREATE STREAMING LIVE TABLE (formerly CREATE INCREMENTAL LIVE TABLE) syntax over the CREATE LIVE TABLE syntax when creating Delta Live Tables (DLT) tables using SQL?

- A. CREATE STREAMING LIVE TABLE should be used when the subsequent step in the DLT pipeline is static.
- B. CREATE STREAMING LIVE TABLE should be used when data needs to be processed incrementally.
- C. CREATE STREAMING LIVE TABLE is redundant for DLT and it does not need to be used.
- D. CREATE STREAMING LIVE TABLE should be used when data needs to be processed through complicated aggregations.
- E. CREATE STREAMING LIVE TABLE should be used when the previous step in the DLT pipeline is static.

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 34

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer is designing a data pipeline. The source system generates files in a shared directory that is also used by other processes. As a result, the files should be kept as is and will accumulate in the directory. The data engineer needs to identify which files are new since the previous run in the pipeline, and set up the pipeline to only ingest those new files with each run.

Which of the following tools can the data engineer use to solve this problem?

- A. Unity Catalog
- B. Delta Lake
- C. Databricks SQL
- D. Data Explorer
- E. Auto Loader

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 35

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following Structured Streaming queries is performing a hop from a Silver table to a Gold table?

- A.

```
(spark.readStream.load(rawSalesLocation)
  .writeStream
  .option("checkpointLocation", checkpointPath)
  .outputMode("append")
  .table("newSales")
)
```
- B.

```
(spark.read.load(rawSalesLocation)
  .writeStream
  .option("checkpointLocation", checkpointPath)
  .outputMode("append")
  .table("newSales")
)
```
- C.

```
(spark.table("sales")
  .withColumn("avgPrice", col("sales") / col("units"))
  .writeStream
  .option("checkpointLocation", checkpointPath)
  .outputMode("append")
  .table("newSales")
)
```
- D.

```
(spark.table("sales")
  .filter(col("units") > 0)
  .writeStream
  .option("checkpointLocation", checkpointPath)
  .outputMode("append")
  .table("newSales")
)
```
- E.

```
(spark.table("sales")
  .groupBy("store")
  .agg(sum("sales"))
  .writeStream
  .option("checkpointLocation", checkpointPath)
  .outputMode("complete")
  .table("newSales")
)
```

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 36

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer has three tables in a Delta Live Tables (DLT) pipeline. They have configured the pipeline to drop invalid records at each table. They notice that some data is being dropped due to quality concerns at some point in the DLT pipeline. They would like to determine at which table in their pipeline the data is being dropped.

Which of the following approaches can the data engineer take to identify the table that is dropping the records?

- A. They can set up separate expectations for each table when developing their DLT pipeline.
- B. They cannot determine which table is dropping the records.
- C. They can set up DLT to notify them via email when records are dropped.
- D. They can navigate to the DLT pipeline page, click on each table, and view the data quality statistics.
- E. They can navigate to the DLT pipeline page, click on the "Error" button, and review the present errors.

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 37

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer has a single-task Job that runs each morning before they begin working. After identifying an upstream data issue, they need to set up another task to run a new notebook prior to the original task.

Which of the following approaches can the data engineer use to set up the new task?

- A. They can clone the existing task in the existing Job and update it to run the new notebook.
- B. They can create a new task in the existing Job and then add it as a dependency of the original task.
- C. They can create a new task in the existing Job and then add the original task as a dependency of the new task.
- D. They can create a new job from scratch and add both tasks to run concurrently.
- E. They can clone the existing task to a new Job and then edit it to run the new notebook.

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 38

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

An engineering manager wants to monitor the performance of a recent project using a Databricks SQL query. For the first week following the project's release, the manager wants the query results to be updated every minute. However, the manager is concerned that the compute resources used for the query will be left running and cost the organization a lot of money beyond the first week of the project's release.

Which of the following approaches can the engineering team use to ensure the query does not cost the organization any money beyond the first week of the project's release?

- A. They can set a limit to the number of DBUs that are consumed by the SQL Endpoint.
- B. They can set the query's refresh schedule to end after a certain number of refreshes.
- C. They cannot ensure the query does not cost the organization money beyond the first week of the project's release.
- D. They can set a limit to the number of individuals that are able to manage the query's refresh schedule.
- E. They can set the query's refresh schedule to end on a certain date in the query scheduler.

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 39

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data analysis team has noticed that their Databricks SQL queries are running too slowly when connected to their always-on SQL endpoint. They claim that this issue is present when many members of the team are running small queries simultaneously. They ask the data engineering team for help. The data engineering team notices that each of the team's queries uses the same SQL endpoint.

Which of the following approaches can the data engineering team use to improve the latency of the team's queries?

- A. They can increase the cluster size of the SQL endpoint.
- B. They can increase the maximum bound of the SQL endpoint's scaling range.
- C. They can turn on the Auto Stop feature for the SQL endpoint.
- D. They can turn on the Serverless feature for the SQL endpoint.
- E. They can turn on the Serverless feature for the SQL endpoint and change the Spot Instance Policy to "Reliability Optimized."

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 40

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer wants to schedule their Databricks SQL dashboard to refresh once per day, but they only want the associated SQL endpoint to be running when it is necessary.

Which of the following approaches can the data engineer use to minimize the total running time of the SQL endpoint used in the refresh schedule of their dashboard?

- A. They can ensure the dashboard's SQL endpoint matches each of the queries' SQL endpoints.
- B. They can set up the dashboard's SQL endpoint to be serverless.
- C. They can turn on the Auto Stop feature for the SQL endpoint.
- D. They can reduce the cluster size of the SQL endpoint.
- E. They can ensure the dashboard's SQL endpoint is not one of the included query's SQL endpoint.

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 41

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer has been using a Databricks SQL dashboard to monitor the cleanliness of the input data to an ELT job. The ELT job has its Databricks SQL query that returns the number of input records containing unexpected NULL values. The data engineer wants their entire team to be notified via a messaging webhook whenever this value reaches 100.

Which of the following approaches can the data engineer use to notify their entire team via a messaging webhook whenever the number of NULL values reaches 100?

- A. They can set up an Alert with a custom template.
- B. They can set up an Alert with a new email alert destination.
- C. They can set up an Alert with a new webhook alert destination.
- D. They can set up an Alert with one-time notifications.
- E. They can set up an Alert without notifications.

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 42

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A single Job runs two notebooks as two separate tasks. A data engineer has noticed that one of the notebooks is running slowly in the Job's current run. The data engineer asks a tech lead for help in identifying why this might be the case.

Which of the following approaches can the tech lead use to identify why the notebook is running slowly as part of the Job?

- A. They can navigate to the Runs tab in the Jobs UI to immediately review the processing notebook.
- B. They can navigate to the Tasks tab in the Jobs UI and click on the active run to review the processing notebook.
- C. They can navigate to the Runs tab in the Jobs UI and click on the active run to review the processing notebook.
- D. There is no way to determine why a Job task is running slowly.
- E. They can navigate to the Tasks tab in the Jobs UI to immediately review the processing notebook.

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 43

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer has a Job with multiple tasks that runs nightly. Each of the tasks runs slowly because the clusters take a long time to start. Which of the following actions can the data engineer perform to improve the start up time for the clusters used for the Job?

- A. They can use endpoints available in Databricks SQL
- B. They can use jobs clusters instead of all-purpose clusters
- C. They can configure the clusters to be single-node
- D. They can use clusters that are from a cluster pool
- E. They can configure the clusters to autoscale for larger data sizes

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 44

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A new data engineering team team. has been assigned to an ELT project. The new data engineering team will need full privileges on the database customers to fully manage the project.

Which of the following commands can be used to grant full permissions on the database to the new data engineering team?

- A. GRANT USAGE ON DATABASE customers TO team;
- B. GRANT ALL PRIVILEGES ON DATABASE team TO customers;
- C. GRANT SELECT PRIVILEGES ON DATABASE customers TO teams;
- D. GRANT SELECT CREATE MODIFY USAGE PRIVILEGES ON DATABASE customers TO team;
- E. GRANT ALL PRIVILEGES ON DATABASE customers TO team;

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 45

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A new data engineering team has been assigned to work on a project. The team will need access to database customers in order to see what tables already exist. The team has its own group team.

Which of the following commands can be used to grant the necessary permission on the entire database to the new team?

- A. GRANT VIEW ON CATALOG customers TO team;
- B. GRANT CREATE ON DATABASE customers TO team;
- C. GRANT USAGE ON CATALOG team TO customers;
- D. GRANT CREATE ON DATABASE team TO customers;
- E. GRANT USAGE ON DATABASE customers TO team;

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 46

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer is running code in a Databricks Repo that is cloned from a central Git repository. A colleague of the data engineer informs them that changes have been made and synced to the central Git repository. The data engineer now needs to sync their Databricks Repo to get the changes from the central Git repository.

Which of the following Git operations does the data engineer need to run to accomplish this task?

- A. Merge
- B. Push
- C. Pull
- D. Commit
- E. Clone

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 47

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following is a benefit of the Databricks Lakehouse Platform embracing open source technologies?

- A. Cloud-specific integrations
- B. Simplified governance
- C. Ability to scale storage
- D. Ability to scale workloads
- E. Avoiding vendor lock-in

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 48

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer needs to use a Delta table as part of a data pipeline, but they do not know if they have the appropriate permissions.

In which of the following locations can the data engineer review their permissions on the table?

- A. Databricks Filesystem
- B. Jobs
- C. Dashboards
- D. Repos
- E. Data Explorer

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 49

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following describes a scenario in which a data engineer will want to use a single-node cluster?

- A. When they are working interactively with a small amount of data
- B. When they are running automated reports to be refreshed as quickly as possible
- C. When they are working with SQL within Databricks SQL
- D. When they are concerned about the ability to automatically scale with larger data
- E. When they are manually running reports with a large amount of data

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 50

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer has been given a new record of data:

id STRING = 'a1'

rank INTEGER = 6

rating FLOAT = 9.4

Which of the following SQL commands can be used to append the new record to an existing Delta table my_table?

- A. INSERT INTO my_table VALUES ('a1', 6, 9.4)
- B. my_table UNION VALUES ('a1', 6, 9.4)
- C. INSERT VALUES ('a1' , 6, 9.4) INTO my_table
- D. UPDATE my_table VALUES ('a1', 6, 9.4)
- E. UPDATE VALUES ('a1', 6, 9.4) my_table

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 51

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer has realized that the data files associated with a Delta table are incredibly small. They want to compact the small files to form larger files to improve performance.

Which of the following keywords can be used to compact the small files?

- A. REDUCE
- B. OPTIMIZE
- C. COMPACTION
- D. REPARTITION
- E. VACUUM

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 52

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

In which of the following file formats is data from Delta Lake tables primarily stored?

- A. Delta
- B. CSV
- C. Parquet
- D. JSON
- E. A proprietary, optimized format specific to Databricks

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 53

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following is stored in the Databricks customer's cloud account?

- A. Databricks web application
- B. Cluster management metadata
- C. Repos
- D. Data
- E. Notebooks

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 54

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

Which of the following can be used to simplify and unify siloed data architectures that are specialized for specific use cases?

- A. None of these
- B. Data lake
- C. Data warehouse
- D. All of these
- E. Data lakehouse

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 55

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data architect has determined that a table of the following format is necessary:

employeeId	startDate	avgRating
a1	2009-01-06	5.5
a2	2018-11-21	7.1
...

Which of the following code blocks uses SQL DDL commands to create an empty Delta table in the above format regardless of whether a table already exists with this name?

- ```
CREATE TABLE IF NOT EXISTS table_name (
 employeeId STRING,
 startDate DATE,
 avgRating FLOAT
)

CREATE OR REPLACE TABLE table_name AS
SELECT
 employeeId STRING,
 startDate DATE,
 avgRating FLOAT
USING DELTA

CREATE OR REPLACE TABLE table_name WITH COLUMNS (
 employeeId STRING,
 startDate DATE,
 avgRating FLOAT
) USING DELTA

CREATE TABLE table_name AS
SELECT
 employeeId STRING,
 startDate DATE,
 avgRating FLOAT

CREATE OR REPLACE TABLE table_name (
 employeeId STRING,
 startDate DATE,
 avgRating FLOAT
)
```
- A.
- B.
- C.
- D.
- E.

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 56

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

---

A data engineer has a Python notebook in Databricks, but they need to use SQL to accomplish a specific task within a cell. They still want all of the other cells to use Python without making any changes to those cells.

Which of the following describes how the data engineer can use SQL within a cell of their Python notebook?

- A. It is not possible to use SQL in a Python notebook
- B. They can attach the cell to a SQL endpoint rather than a Databricks cluster
- C. They can simply write SQL syntax in the cell
- D. They can add %sql to the first line of the cell
- E. They can change the default language of the notebook to SQL

Show Suggested Answer







Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 57

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

---

Which of the following SQL keywords can be used to convert a table from a long format to a wide format?

- A. TRANSFORM
- B. PIVOT
- C. SUM
- D. CONVERT
- E. WHERE

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 58

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

---

Which of the following describes a benefit of creating an external table from Parquet rather than CSV when using a CREATE TABLE AS SELECT statement?

- A. Parquet files can be partitioned
- B. CREATE TABLE AS SELECT statements cannot be used on files
- C. Parquet files have a well-defined schema
- D. Parquet files have the ability to be optimized
- E. Parquet files will become Delta tables

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 59

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

---

A data engineer wants to create a relational object by pulling data from two tables. The relational object does not need to be used by other data engineers in other sessions. In order to save on storage costs, the data engineer wants to avoid copying and storing physical data.

Which of the following relational objects should the data engineer create?

- A. Spark SQL Table
- B. View
- C. Database
- D. Temporary view
- E. Delta Table

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 60

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

---

A data analyst has developed a query that runs against Delta table. They want help from the data engineering team to implement a series of tests to ensure the data returned by the query is clean. However, the data engineering team uses Python for its tests rather than SQL.

Which of the following operations could the data engineering team use to run the query and operate with the results in PySpark?

- A. `SELECT * FROM sales`
- B. `spark.delta.table`
- C. `spark.sql`
- D. There is no way to share data between PySpark and SQL.
- E. `spark.table`

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 61

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

---

Which of the following commands will return the number of null values in the member\_id column?

- A. `SELECT count(member_id) FROM my_table;`
- B. `SELECT count(member_id) - count_null(member_id) FROM my_table;`
- C. `SELECT count_if(member_id IS NULL) FROM my_table;`
- D. `SELECT null(member_id) FROM my_table;`
- E. `SELECT count_null(member_id) FROM my_table;`

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 62

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer needs to apply custom logic to identify employees with more than 5 years of experience in array column employees in table stores. The custom logic should create a new column exp\_employees that is an array of all of the employees with more than 5 years of experience for each row. In order to apply this custom logic at scale, the data engineer wants to use the FILTER higher-order function.

Which of the following code blocks successfully completes this task?

- ```
SELECT
  store_id,
  employees,
  FILTER (employees, i -> i.years_exp > 5) AS exp_employees
FROM stores;
```
- ```
SELECT
 store_id,
 employees,
 FILTER (exp_employees, years_exp > 5) AS exp_employees
FROM stores;
```
- ```
SELECT
  store_id,
  employees,
  FILTER (employees, years_exp > 5) AS exp_employees
FROM stores;
```
- ```
SELECT
 store_id,
 employees,
 CASE WHEN employees.years_exp > 5 THEN employees
 ELSE NULL
 END AS exp_employees
FROM stores;
```
- ```
SELECT
  store_id,
  employees,
  FILTER (exp_employees, i -> i.years_exp > 5) AS exp_employees
FROM stores;
```

Show Suggested Answer

Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 63

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer has a Python variable `table_name` that they would like to use in a SQL query. They want to construct a Python code block that will run the query using `table_name`.

They have the following incomplete code block:

```
____(f"SELECT customer_id, spend FROM {table_name}")
```

Which of the following can be used to fill in the blank to successfully complete the task?

- A. `spark.delta.sql`
- B. `spark.delta.table`
- C. `spark.table`
- D. `dbutils.sql`
- E. `spark.sql`

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 64

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer has created a new database using the following command:

```
CREATE DATABASE IF NOT EXISTS customer360;
```

In which of the following locations will the customer360 database be located?

- A. dbfs:/user/hive/database/customer360
- B. dbfs:/user/hive/warehouse
- C. dbfs:/user/hive/customer360
- D. More information is needed to determine the correct response
- E. dbfs:/user/hive/database

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 65

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer is attempting to drop a Spark SQL table `my_table` and runs the following command:

```
DROP TABLE IF EXISTS my_table;
```

After running this command, the engineer notices that the data files and metadata files have been deleted from the file system.

Which of the following describes why all of these files were deleted?

- A. The table was managed
- B. The table's data was smaller than 10 GB
- C. The table's data was larger than 10 GB
- D. The table was external
- E. The table did not have a location

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 66

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer that is new to using Python needs to create a Python function to add two integers together and return the sum?

Which of the following code blocks can the data engineer use to complete this task?

A.

```
function add_integers(x, y):  
    return x + y
```

B.

```
function add_integers(x, y):  
    x + y
```

C.

```
def add_integers(x, y):  
    print(x + y)
```

D.

```
def add_integers(x, y):  
    return x + y
```

E.

```
def add_integers(x, y):  
    x + y
```

Show Suggested Answer





Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 67

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

In which of the following scenarios should a data engineer use the MERGE INTO command instead of the INSERT INTO command?

- A. When the location of the data needs to be changed
- B. When the target table is an external table
- C. When the source table can be deleted
- D. When the target table cannot contain duplicate records
- E. When the source is not a Delta table

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 68

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer is working with two tables. Each of these tables is displayed below in its entirety.

sales

customer_id	spend	units
a1	28.94	7
a3	874.12	23
a4	8.99	1

favorite_stores

customer_id	store_id
a1	s1
a2	s1
a4	s2

The data engineer runs the following query to join these tables together:

```
SELECT
  sales.customer_id,
  sales.spend,
  favorite_stores.store_id
FROM sales
LEFT JOIN favorite_stores
ON sales.customer_id = favorite_stores.customer_id;
```

Which of the following will be returned by the above query?

A.

customer_id	spend	store_id
a1	28.94	s1
a4	8.99	s2

B.

customer_id	spend	units	store_id
a1	28.94	7	s1
a4	8.99	1	s2

C.

customer_id	spend	store_id
a1	28.94	s1
a3	874.12	NULL
a4	8.99	s2

D.

customer_id	spend	store_id
a1	28.94	s1
a2	NULL	s1
a3	874.12	NULL
a4	8.99	s2

E.

customer_id	spend	store_id
a1	28.94	s1
a2	NULL	s1
a4	8.99	s2

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 69

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer needs to create a table in Databricks using data from a CSV file at location /path/to/csv.

They run the following command:

```
CREATE TABLE new_table
_____
OPTIONS (
  header = "true",
  delimiter = "|"
)
LOCATION "path/to/csv"
```

Which of the following lines of code fills in the above blank to successfully complete the task?

- A. None of these lines of code are needed to successfully complete the task
- B. USING CSV
- C. FROM CSV
- D. USING DELTA
- E. FROM "path/to/csv"

Show Suggested Answer

Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 70

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer has configured a Structured Streaming job to read from a table, manipulate the data, and then perform a streaming write into a new table.

The code block used by the data engineer is below:

```
(spark.readStream
  .table("sales")
  .withColumn("avg_price", col("sales") / col("units"))
  .writeStream
  .option("checkpointLocation", checkpointPath)
  .outputMode("complete")
  ._____
  .table("new_sales")
)
```

If the data engineer only wants the query to process all of the available data in as many batches as required, which of the following lines of code should the data engineer use to fill in the blank?

- A. processingTime(1)
- B. trigger(availableNow=True)
- C. trigger(parallelBatch=True)
- D. trigger(processingTime="once")
- E. trigger(continuous="once")

Show Suggested Answer

Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 71

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A data engineer has developed a data pipeline to ingest data from a JSON source using Auto Loader, but the engineer has not provided any type inference or schema hints in their pipeline. Upon reviewing the data, the data engineer has noticed that all of the columns in the target table are of the string type despite some of the fields only including float or boolean values.

Which of the following describes why Auto Loader inferred all of the columns to be of the string type?

- A. There was a type mismatch between the specific schema and the inferred schema
- B. JSON data is a text-based format
- C. Auto Loader only works with string data
- D. All of the fields had at least one null value
- E. Auto Loader cannot infer the schema of ingested data

Show Suggested Answer



Actual exam question from Databricks's Certified Data Engineer Associate

Question #: 72

Topic #: 1

[\[All Certified Data Engineer Associate Questions\]](#)

A Delta Live Table pipeline includes two datasets defined using STREAMING LIVE TABLE. Three datasets are defined against Delta Lake table sources using LIVE TABLE.

The table is configured to run in Development mode using the Continuous Pipeline Mode.

Assuming previously unprocessed data exists and all definitions are valid, what is the expected outcome after clicking Start to update the pipeline?

- A. All datasets will be updated once and the pipeline will shut down. The compute resources will be terminated.
- B. All datasets will be updated at set intervals until the pipeline is shut down. The compute resources will persist until the pipeline is shut down.
- C. All datasets will be updated once and the pipeline will persist without any processing. The compute resources will persist but go unused.
- D. All datasets will be updated once and the pipeline will shut down. The compute resources will persist to allow for additional testing.
- E. All datasets will be updated at set intervals until the pipeline is shut down. The compute resources will persist to allow for additional testing.

Show Suggested Answer

