

Reliable Associate-Developer-Apache-Spark-3.5 Exam Guide, Associate-Developer-Apache-Spark-3.5 Test Braindumps



2026 Latest ITdumpsfree Associate-Developer-Apache-Spark-3.5 PDF Dumps and Associate-Developer-Apache-Spark-3.5 Exam Engine Free Share: <https://drive.google.com/open?id=1ZII9q1g4VvNa48lhvvMwoeUPQjOSmljj>

Our Associate-Developer-Apache-Spark-3.5 training materials provide 3 versions to the client and they include the PDF version, PC version, APP online version. Each version's using method and functions are different but the questions and answers of our Associate-Developer-Apache-Spark-3.5 study quiz is the same. The client can decide which Associate-Developer-Apache-Spark-3.5 version to choose according their hobbies and their practical conditions. You will be surprised by the convenient functions of our Associate-Developer-Apache-Spark-3.5 exam dumps.

There has been fierce and intensified competition going on in the practice materials market. As the leading commodity of the exam, our Associate-Developer-Apache-Spark-3.5 practice materials have get pressing requirements and steady demand from exam candidates all the time. So our Associate-Developer-Apache-Spark-3.5 practice materials have active demands than others with high passing rate of 98 to 100 percent. We are one of the largest and the most confessional dealer of practice materials. That is why our Associate-Developer-Apache-Spark-3.5 practice materials outreach others greatly among substantial suppliers of the exam.

>> Reliable Associate-Developer-Apache-Spark-3.5 Exam Guide <<

Pass Guaranteed Quiz Databricks - Fantastic Reliable Associate-Developer-Apache-Spark-3.5 Exam Guide

Before the clients purchase our Associate-Developer-Apache-Spark-3.5 study materials, they can have a free trial freely. The clients can log in our company's website and visit the pages of our products. The pages of our products lists many important information about our Associate-Developer-Apache-Spark-3.5 study materials and they include the price, version and updated time of our products, the exam name and code, the total amount of the questions and answers, the merits of our Associate-Developer-Apache-Spark-3.5 Study Materials and the discounts. You can have a comprehensive understanding of our Associate-Developer-Apache-Spark-3.5 study materials after you see this information. Then you can look at the free demos and try to answer them to see the value of our Associate-Developer-Apache-Spark-3.5 study materials and finally decide to buy them or not.

Databricks Certified Associate Developer for Apache Spark 3.5 - Python Sample Questions (Q68-Q73):

NEW QUESTION # 68

A data engineer wants to create an external table from a JSON file located at /data/input.json with the following requirements:

Create an external table named users

Automatically infer schema

Merge records with differing schemas

Which code snippet should the engineer use?

Options:

- A. CREATE TABLE users USING json OPTIONS (path '/data/input.json')
- B. CREATE EXTERNAL TABLE users USING json OPTIONS (path '/data/input.json')
- C. **CREATE EXTERNAL TABLE users USING json OPTIONS (path '/data/input.json', mergeSchema 'true')**
- D. CREATE EXTERNAL TABLE users USING json OPTIONS (path '/data/input.json', schemaMerge 'true')

Answer: C

Explanation:

To create an external table and enable schema merging, the correct syntax is:

CREATE EXTERNAL TABLE users

USING json

OPTIONS (

path '/data/input.json',

mergeSchema 'true'

)

mergeSchema is the correct option key (not schemaMerge)

EXTERNAL allows Spark to query files without managing their lifecycle

NEW QUESTION # 69

20 of 55.

What is the difference between df.cache() and df.persist() in Spark DataFrame?

- A. Both functions perform the same operation. The persist() function provides improved performance as its default storage level is DISK_ONLY.
- B. Both cache() and persist() can be used to set the default storage level (MEMORY_AND_DISK_DESER).
- C. persist() - Persists the DataFrame with the default storage level (MEMORY_AND_DISK_DESER), and cache() - Can be used to set different storage levels.
- D. **cache() - Persists the DataFrame with the default storage level (MEMORY_AND_DISK_DESER), and persist() - Can be used to set different storage levels to persist the contents of the DataFrame.**

Answer: D

Explanation:

Both cache() and persist() are Spark DataFrame storage operations that store computed results in memory (and optionally on disk) to speed up subsequent actions on the same DataFrame.

Key difference:

cache() is a shorthand for persist(StorageLevel.MEMORY_AND_DISK).

persist() allows specifying different storage levels, such as MEMORY_ONLY, DISK_ONLY, or MEMORY_AND_DISK_SER.

Example:

df.cache() # Uses MEMORY_AND_DISK by default

df.persist(StorageLevel.MEMORY_ONLY) # Custom storage level

Both trigger caching upon an action (e.g., count(), collect()).

Why the other options are incorrect:

A: persist() default is not DISK_ONLY; default storage level is MEMORY_AND_DISK.

B/C: cache() cannot set arbitrary levels; only persist() can.

Reference:

PySpark API Reference - DataFrame.cache() and DataFrame.persist().

Databricks Exam Guide (June 2025): Section "Developing Apache Spark DataFrame/DataSet API Applications" - caching, persistence, and storage levels.

NEW QUESTION # 70

21 of 55.

What is the behavior of the function date_sub(start, days) if a negative value is passed into the days parameter?

- A. The same start date will be returned.

- B. An error message of an invalid parameter will be returned.
- C. The number of days specified will be removed from the start date.
- D. The number of days specified will be added to the start date.

Answer: D

Explanation:

In Spark SQL, the function `date_sub(startDate, days)` returns the date that is days before `startDate`.

If the `days` parameter is negative, Spark interprets it as subtracting a negative number, which effectively adds days to the date.

Example:

```
from pyspark.sql.functions import date_sub, lit
df = spark.createDataFrame([( "2024-10-01", )], ["dt"])
df.select(date_sub("dt", -5).alias("new_date")).show()
```

Output:

```
+-----+
| new_date |
+-----+
|2024-10-06|
+-----+
```

Why the other options are incorrect:

B: No error occurs; negative values are supported.
 C: The start date changes if `days` ≠ 0.
 D: Subtracting days would move the date backward, not forward.

Reference:

Spark SQL Functions - `date_sub(startDate, days)` and `date_add(startDate, days)` behavior.

Databricks Exam Guide (June 2025): Section "Using Spark SQL" - working with date and timestamp functions.

NEW QUESTION # 71

A data engineer noticed improved performance after upgrading from Spark 3.0 to Spark 3.5. The engineer found that Adaptive Query Execution (AQE) was enabled.

Which operation is AQE implementing to improve performance?

- A. Improving the performance of single-stage Spark jobs
- B. Collecting persistent table statistics and storing them in the metastore for future use
- C. Optimizing the layout of Delta files on disk
- D. Dynamically switching join strategies

Answer: D

Explanation:

Adaptive Query Execution (AQE) is a Spark 3.x feature that dynamically optimizes query plans at runtime. One of its core features is:

Dynamically switching join strategies (e.g., from sort-merge to broadcast) based on runtime statistics.

Other AQE capabilities include:

Coalescing shuffle partitions

Skew join handling

Option A is correct.

Option B refers to statistics collection, which is not AQE's primary function.

Option C is too broad and not AQE-specific.

Option D refers to Delta Lake optimizations, unrelated to AQE.

Final answer: A

NEW QUESTION # 72

A data engineer is building an Apache Spark™ Structured Streaming application to process a stream of JSON events in real time. The engineer wants the application to be fault-tolerant and resume processing from the last successfully processed record in case of a failure. To achieve this, the data engineer decides to implement checkpoints.

Which code snippet should the data engineer use?

- A. `query = streaming_df.writeStream \`

```

    .format("console") \
    .outputMode("complete") \
    .start()
  • B. query = streaming_df.writeStream \
    .format("console") \
    .option("checkpoint", "/path/to/checkpoint") \
    .outputMode("append") \
    .start()
  • C. query = streaming_df.writeStream \
    .format("console") \
    .outputMode("append") \
    .start()
  • D. query = streaming_df.writeStream \
    .format("console") \
    .outputMode("append") \
    .option("checkpointLocation", "/path/to/checkpoint") \
    .start()

```

Answer: D

Explanation:

To enable fault tolerance and ensure that Spark can resume from the last committed offset after failure, you must configure a checkpoint location using the correct option key: "checkpointLocation".

From the official Spark Structured Streaming guide:

"To make a streaming query fault-tolerant and recoverable, a checkpoint directory must be specified using

.option("checkpointLocation", "/path/to/dir")." Explanation of options:

Option A uses an invalid option name: "checkpoint" (should be "checkpointLocation") Option B is correct: it sets checkpointLocation properly Option C lacks checkpointing and won't resume after failure Option D also lacks checkpointing configuration

NEW QUESTION # 73

.....

Associate-Developer-Apache-Spark-3.5 certification exam is a very import component Databricks certification exam. But passing Databricks certification Associate-Developer-Apache-Spark-3.5 exam is not so simple. In order to give to relieve pressure and save time and effort for candidates who take a preparation for the Associate-Developer-Apache-Spark-3.5 Certification Exam, ITdumpsfree specially produce a variety of training tools. So you can choose an appropriate quick training from ITdumpsfree to pass the exam.

Associate-Developer-Apache-Spark-3.5 Test Braindumps: <https://www.itdumpsfree.com/Associate-Developer-Apache-Spark-3.5-exam-passed.html>

Databricks Reliable Associate-Developer-Apache-Spark-3.5 Exam Guide Where can I download my product, As the top company in this field many companies regard Databricks Associate-Developer-Apache-Spark-3.5 Test Braindumps certification as one of products manages elite standards in most of countries, In the unlikely even if you fail the Associate-Developer-Apache-Spark-3.5 exam, we promise to give you full refund, Besides, there are Databricks Associate-Developer-Apache-Spark-3.5 reliable study vce that you can download to learn about our products.

When considering this method, remember that the hair that really gets noticed Associate-Developer-Apache-Spark-3.5 by the viewer is the hair at the edges, D is a programming language built to help programmers address the challenges of modern software development.

Pass Guaranteed Quiz Databricks - Associate-Developer-Apache-Spark-3.5 Useful Reliable Exam Guide

Where can I download my product, As the top company in this Associate-Developer-Apache-Spark-3.5 Vce Format field many companies regard Databricks certification as one of products manages elite standards in most of countries.

In the unlikely even if you fail the Associate-Developer-Apache-Spark-3.5 Exam, we promise to give you full refund, Besides, there are Databricks Associate-Developer-Apache-Spark-3.5 reliable study vce that you can download to learn about our products.

The Databricks Certified Associate Developer for Apache Spark 3.5 - Python training material is the right decision.

- Exam Dumps Associate-Developer-Apache-Spark-3.5 Pdf Associate-Developer-Apache-Spark-3.5 Latest Exam Papers Latest Associate-Developer-Apache-Spark-3.5 Exam Discount Open  www.testkingpass.com         <img alt="lightbulb icon" data-bbox="80