Useful Databricks - Associate-Developer-Apache-Spark-3.5 - Databricks Certified Associate Developer for Apache Spark 3.5 - Python Practice Questions

BTW, DOWNLOAD part of SurePassExams Associate-Developer-Apache-Spark-3.5 dumps from Cloud Storage: https://drive.google.com/open?id=1XAfF48VqcVGZ8R9lGB06WScPpeb5M_pS

Our staff will be on-line service 24 hours a day. I believe that you have also contacted a lot of service personnel, but I still imagine you praise the staff of our Associate-Developer-Apache-Spark-3.5 study engine. They have the best skills and the most professional service attitude on the Associate-Developer-Apache-Spark-3.5 Practice Questions. He can solve any problems you have encountered while using Associate-Developer-Apache-Spark-3.5 exam simulating for all of our staffs are trained to be professional to help our customers. And they are kind and considerate.

If you are finding a study material to prepare your exam, our material will end your search. Our Associate-Developer-Apache-Spark-3.5 exam torrent has a high quality that you can't expect. I think our Databricks Certified Associate Developer for Apache Spark 3.5 - Python prep torrent will help you save much time, and you will have more free time to do what you like to do. I can guarantee that you will have no regrets about using our Associate-Developer-Apache-Spark-3.5 Test Braindumps When the time for action arrives, stop thinking and go in, try our Associate-Developer-Apache-Spark-3.5 exam torrent, you will find our products will be a very good choice for you to pass your exam and get you certificate in a short time.

>> Associate-Developer-Apache-Spark-3.5 Practice Questions <<

2026 First-grade Associate-Developer-Apache-Spark-3.5 Practice Questions Help You Pass Associate-Developer-Apache-Spark-3.5 Easily

To provide our users with the Databricks Certified Associate Developer for Apache Spark 3.5 - Python (Associate-Developer-Apache-Spark-3.5) latest questions based on the sections of the actual exam quesions, we regularly update our Associate-Developer-Apache-Spark-3.5 study material. Also, SurePassExams provides free updates of Databricks Associate-Developer-Apache-Spark-3.5 Exam Questions for up to 365 days. For customers who don't crack the Databricks Associate-Developer-Apache-Spark-3.5 test after using our product, SurePassExams will provides them a refund guarantee according to terms and conditions.

Databricks Certified Associate Developer for Apache Spark 3.5 - Python Sample Questions (Q130-Q135):

NEW QUESTION # 130
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. persist() - Persists the DataFrame with the default storage level (MEMORY_AND_DISK_DESER), and cache() - Can be used to set different storage levels.
- C. Both cache() and persist() can be used to set the default storage level (MEMORY_AND_DISK_DESER).
- 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 # 131
A data engineer wants to write a Spark job that creates a new managed table. If the table already exists, the job should fail and not modify anything.
Which save mode and method should be used?

- A. save with mode ErrorIfExists
- B. saveAsTable with mode Overwrite
- C. saveAsTable with mode ErrorIfExists
- D. save with mode Ignore

Answer: C

Explanation:
The method saveAsTable() creates a new table and optionally fails if the table exists.

From Spark documentation:
"The mode 'ErrorIfExists' (default) will throw an error if the table already exists." Thus:
Option A is correct.
Option B (Overwrite) would overwrite existing data - not acceptable here.
Option C and D use save(), which doesn't create a managed table with metadata in the metastore.
Final answer: A

NEW QUESTION # 132
33 of 55.
The data engineering team created a pipeline that extracts data from a transaction system.
The transaction system stores timestamps in UTC, and the data engineers must now transform the transaction_datetime field to the "America/New_York" timezone for reporting.
Which code should be used to convert the timestamp to the target timezone?

- A. raw.withColumn("transaction_datetime", convert_timezone(col("transaction_datetime"), "America/New_York"))
- B. raw.withColumn("transaction_datetime", from_utc_timestamp(col("transaction_datetime"), "America/New_York"))
- C. raw.withColumn("transaction_datetime", date_format(col("transaction_datetime"), "America/New_York"))
- D. raw.withColumn("transaction_datetime", to_utc_timestamp(col("transaction_datetime"), "America/New_York"))

Answer: B

Explanation:
In Spark SQL, to convert a UTC timestamp to another timezone, you use the function from_utc_timestamp().
Correct syntax:
from pyspark.sql.functions import from_utc_timestamp, col
df_converted = raw.withColumn(
"transaction_datetime",
from_utc_timestamp(col("transaction_datetime"), "America/New_York")
)
This adjusts the UTC time into the specified timezone using Spark's timezone database.
Why the other options are incorrect:
B: to_utc_timestamp() converts local time to UTC, not the other way around.
C: date_format() formats timestamps as strings but doesn't adjust timezones.
D: convert_timezone() is not a valid Spark SQL function.
Reference:
Spark SQL Functions - from_utc_timestamp() and to_utc_timestamp().
Databricks Exam Guide (June 2025): Section "Using Spark SQL" - working with timestamps and timezone conversions.

NEW QUESTION # 133
A Spark DataFramedfis cached using theMEMORY_AND_DISKstorage level, but the DataFrame is too large to fit entirely in memory.
What is the likely behavior when Spark runs out of memory to store the DataFrame?

- A. Spark duplicates the DataFrame in both memory and disk. If it doesn't fit in memory, the DataFrame is stored and retrieved from the disk entirely.
- B. Spark will store as much data as possible in memory and spill the rest to disk when memory is full, continuing processing with performance overhead.
- C. Spark splits the DataFrame evenly between memory and disk, ensuring balanced storage utilization.
- D. Spark stores the frequently accessed rows in memory and less frequently accessed rows on disk, utilizing both resources to offer balanced performance.

Answer: B

Explanation:
Comprehensive and Detailed Explanation From Exact Extract:
When using theMEMORY_AND_DISKstorage level, Spark attempts to cache as much of the DataFrame in memory as possible. If the DataFrame does not fit entirely in memory, Spark will store the remaining partitions on disk. This allows processing to continue, albeit with a performance overhead due to disk I/O.
As per the Spark documentation:

"MEMORY_AND_DISK: It stores partitions that do not fit in memory on disk and keeps the rest in memory.
This can be useful when working with datasets that are larger than the available memory."
- Perficient Blogs: Spark - StorageLevel
This behavior ensures that Spark can handle datasets larger than the available memory by spilling excess data to disk, thus preventing job failures due to memory constraints.

NEW QUESTION # 134
Which feature of Spark Connect is considered when designing an application to enable remote interaction with the Spark cluster?

- A. It is primarily used for data ingestion into Spark from external sources
- B. It allows for remote execution of Spark jobs
- C. It can be used to interact with any remote cluster using the REST API
- D. It provides a way to run Spark applications remotely in any programming language

Answer: B

Explanation:
Spark Connect introduces a decoupled client-server architecture. Its key feature is enabling Spark job submission and execution from remote clients - in Python, Java, etc.
From Databricks documentation:
"Spark Connect allows remote clients to connect to a Spark cluster and execute Spark jobs without being co-located with the Spark driver." A is close, but "any language" is overstated (currently supports Python, Java, etc., not literally all).
B refers to REST, which is not Spark Connect's mechanism.
D is incorrect; Spark Connect isn't focused on ingestion.
Final answer: C

NEW QUESTION # 135
.....

Associate-Developer-Apache-Spark-3.5 study material has a high quality service team. First of all, the authors of study materials are experts in the field. They have been engaged in research on the development of the industry for many years, and have a keen sense of smell for changes in the examination direction. During your installation, Associate-Developer-Apache-Spark-3.5 exam questions hired dedicated experts to provide you with free remote online guidance. During your studies, Associate-Developer-Apache-Spark-3.5 Exam Torrent also provides you with free online services for 24 hours, regardless of where and when you are, as long as an email, we will solve all the problems for you. At the same time, if you fail to pass the exam after you have purchased Associate-Developer-Apache-Spark-3.5 training materials, you just need to submit your transcript to our customer service staff and you will receive a full refund.

Associate-Developer-Apache-Spark-3.5 New Dumps: https://www.surepassexams.com/Associate-Developer-Apache-Spark-3.5-exam-bootcamp.html

Databricks Associate-Developer-Apache-Spark-3.5 Practice Questions If you want to know details about each exam materials, our service will be waiting for you 7*24 online, Databricks Associate-Developer-Apache-Spark-3.5 Practice Questions We provide online contact system 24 hours per day, 7 days a week to our customers, Databricks Associate-Developer-Apache-Spark-3.5 Practice Questions Because these exam dumps on our website are based on the real exam and edited by our IT experts with years of experience, their qualities are guaranteed and they have a 99% hit rate, As a professional website, SurePassExams Associate-Developer-Apache-Spark-3.5 New Dumps does not only guarantee you will receive a high score in your actual test, but also provide you with the most efficiency way to get success.

Whether youre in front of a live audience or Associate-Developer-Apache-Spark-3.5 Test Objectives Pdf a webcam, the way you present yourself verbally and non-verbally is crucial to your success, As you adjust the Feather slider, the Associate-Developer-Apache-Spark-3.5 Top Exam Dumps outer circle expands or contracts to indicate the hardness or softness of the brush.

2026 Associate-Developer-Apache-Spark-3.5 – 100% Free Practice Questions | Perfect Databricks Certified Associate Developer for Apache Spark 3.5 - Python New Dumps

If you want to know details about each exam materials, our service Associate-Developer-Apache-Spark-3.5 Top Exam Dumps will

be waiting for you 7*24 online, We provide online contact system 24 hours per day, 7 days a week to our customers.

Because these exam dumps on our website are based on the Real Associate-Developer-Apache-Spark-3.5 Exam and edited by our IT experts with years of experience, their qualities are guaranteed and they have a 99% hit rate.

As a professional website, SurePassExams does not only guarantee Latest Associate-Developer-Apache-Spark-3.5 Exam Price you will receive a high score in your actual test, but also provide you with the most efficiency way to get success.

SurePassExams provides Databricks Associate-Developer-Apache-Spark-3.5 desktop-based practice software for you to test your knowledge and abilities.

P.S. Free 2026 Databricks Associate-Developer-Apache-Spark-3.5 dumps are available on Google Drive shared by SurePassExams: https://drive.google.com/open?id=1XAfF48VqcVGZ8R9lGB06WScPpeb5M_pS