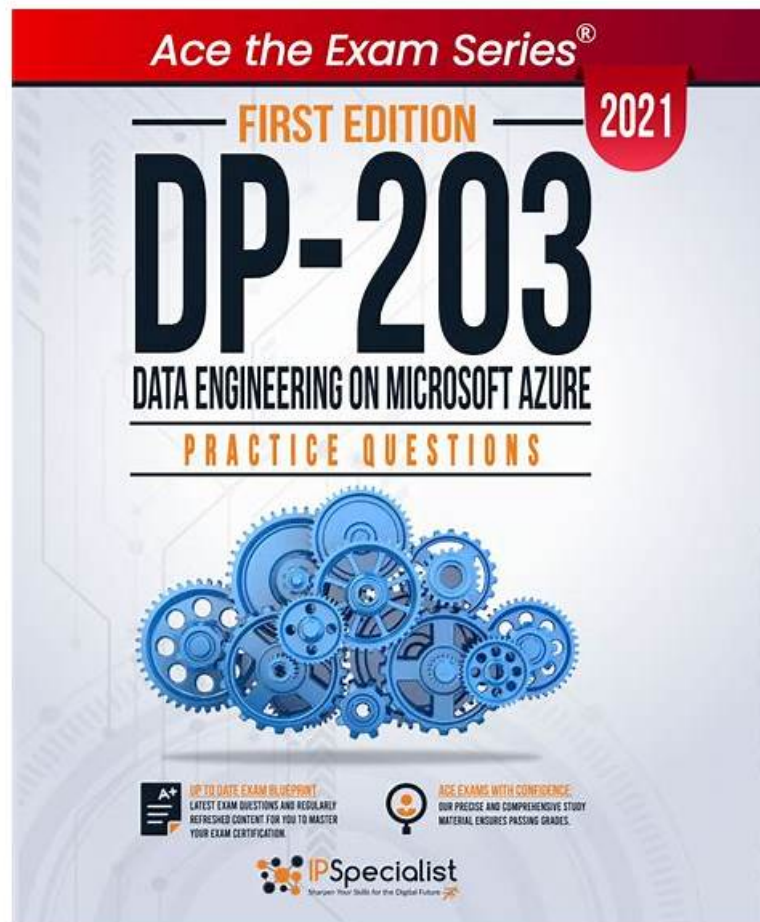


# Get Useful DP-203 Practice Mock and Pass Exam in First Attempt



BONUS!!! Download part of ExamDumpsVCE DP-203 dumps for free: <https://drive.google.com/open?id=1O23QMFRs9XNy42dvKEzcjRBH98Frg6jf>

If without a quick purchase process, users of our DP-203 quiz guide will not be able to quickly start their own review program. So, our company employs many experts to design a fast sourcing channel for our DP-203 exam prep. All users can implement fast purchase and use our DP-203 learning materials. We have specialized software to optimize the user's purchase channels, if you decide to purchase our DP-203 prepare questions, you can achieve the DP-203 exam questions content even if the update service and efficient and convenient user experience and you will pass the exam for sure.

How can our DP-203 study questions are so famous and become the leader in the market? Because our DP-203 learning braindumps comprise the most significant questions and answers that have every possibility to be the part of the real exam. As you study with our DP-203 Practice Guide, you will find the feeling that you are doing the real exam. Especially if you choose the Software version of our DP-203 training engine, which can simulate the real exam.

>> DP-203 Practice Mock <<

## Test DP-203 Registration | DP-203 Best Practice

If you are a child's mother, with DP-203 test answers, you will have more time to stay with your child; if you are a student, with DP-203 exam torrent, you will have more time to travel to comprehend the wonders of the world. In the other worlds, with DP-203 guide tests, learning will no longer be a burden in your life. You can save much time and money to do other things what meaningful. You will no longer feel tired because of your studies, if you decide to choose and practice our DP-203test answers. Your life will be even more exciting.

## Microsoft Data Engineering on Microsoft Azure Sample Questions (Q78-Q83):

### NEW QUESTION # 78

You are designing a solution that will use tables in Delta Lake on Azure Databricks.

You need to minimize how long it takes to perform the following:

- \* Queries against non-partitioned tables

- \* Joins on non-partitioned columns

Which two options should you include in the solution? Each correct answer presents part of the solution.

(Choose Correct Answer and Give Explanation and References to Support the answers based from Data Engineering on Microsoft Azure)

- A. the clone command
- B. Z-Ordering
- C. Apache Spark caching
- D. dynamic file pruning (DFP)

**Answer: B,D**

Explanation:

According to the information I found on the web, two options that you should include in the solution to minimize how long it takes to perform queries and joins on non-partitioned tables are:

- \* Z-Ordering: This is a technique to colocate related information in the same set of files. This co-locality is automatically used by Delta Lake in data-skipping algorithms. This behavior dramatically reduces the amount of data that Delta Lake on Azure Databricks needs to read.

- \* Apache Spark caching: This is a feature that allows you to cache data in memory or on disk for faster access. Caching can improve the performance of repeated queries and joins on the same data. You can cache Delta tables using the `CACHE TABLE` or `CACHE LAZY` commands.

To minimize the time it takes to perform queries against non-partitioned tables and joins on non-partitioned columns in Delta Lake on Azure Databricks, the following options should be included in the solution:

A: Z-Ordering: Z-Ordering improves query performance by co-locating data that share the same column values in the same physical partitions. This reduces the need for shuffling data across nodes during query execution. By using Z-Ordering, you can avoid full table scans and reduce the amount of data processed.

B: Apache Spark caching: Caching data in memory can improve query performance by reducing the amount of data read from disk. This helps to speed up subsequent queries that need to access the same data. When you cache a table, the data is read from the data source and stored in memory. Subsequent queries can then read the data from memory, which is much faster than reading it from disk.

References:

- \* Delta Lake on Databricks: <https://docs.databricks.com/delta/index.html>

- \* Best Practices for Delta Lake on Databricks: <https://databricks.com/blog/2020/05/14/best-practices-for-delta-lake-on-databricks.html>

### NEW QUESTION # 79

You need to ensure that the Twitter feed data can be analyzed in the dedicated SQL pool. The solution must meet the customer sentiment analytics requirements.

Which three Transaction-SQL DDL commands should you run in sequence? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the correct order.

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

**Commands**

**Answer Area**

CREATE EXTERNAL DATA SOURCE

CREATE EXTERNAL FILE FORMAT

CREATE EXTERNAL TABLE

CREATE EXTERNAL TABLE AS SELECT

CREATE DATABASE SCOPED CREDENTIAL



**Answer:**

**Explanation:**

Commands	Answer Area
CREATE EXTERNAL DATA SOURCE	CREATE EXTERNAL DATA SOURCE
CREATE EXTERNAL FILE FORMAT	CREATE EXTERNAL FILE FORMAT
CREATE EXTERNAL TABLE	CREATE EXTERNAL TABLE AS SELECT
CREATE EXTERNAL TABLE AS SELECT	
CREATE DATABASE SCOPED CREDENTIAL	

**Reference:**

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-external-tables>

### NEW QUESTION # 80

You have an Azure Data Lake Storage Gen2 account that contains a JSON file for customers. The file contains two attributes named FirstName and LastName.

You need to copy the data from the JSON file to an Azure Synapse Analytics table by using Azure Databricks. A new column must be created that concatenates the FirstName and LastName values.

You create the following components:

A destination table in Azure Synapse

An Azure Blob storage container

A service principal

In which order should you perform the actions? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Mount the Data Lake Storage onto DBFS.	
Write the results to a table in Azure Synapse.	
Specify a temporary folder to stage the data.	
Read the file into a data frame.	
Perform transformations on the data frame.	

**Answer:**

**Explanation:**

Actions	Answer Area
Mount the Data Lake Storage onto DBFS.	Mount the Data Lake Storage onto DBFS.
Write the results to a table in Azure Synapse.	Read the file into a data frame.
Specify a temporary folder to stage the data.	Perform transformations on the data frame.
Read the file into a data frame.	Specify a temporary folder to stage the data.
Perform transformations on the data frame.	Write the results to a table in Azure Synapse.

**Reference:**

<https://docs.microsoft.com/en-us/azure/azure-databricks/databricks-extract-load-sql-data-warehouse>

### NEW QUESTION # 81

You have an Azure Data Lake Storage Gen2 account that contains a JSON file for customers. The file contains two attributes named FirstName and LastName.

You need to copy the data from the JSON file to an Azure Synapse Analytics table by using Azure Databricks. A new column must



be created that concatenates the FirstName and LastName values.

You create the following components:

A destination table in Azure Synapse

An Azure Blob storage container

A service principal

In which order should you perform the actions? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

**Actions**

Mount the Data Lake Storage onto DBFS.

Write the results to a table in Azure Synapse.

Specify a temporary folder to stage the data.

Read the file into a data frame.

Perform transformations on the data frame.

**Answer Area**

**Answer:**

Explanation:

Actions	Answer Area
Mount the Data Lake Storage onto DBFS.	Mount the Data Lake Storage onto DBFS.
Write the results to a table in Azure Synapse.	Read the file into a data frame.
Specify a temporary folder to stage the data.	Perform transformations on the data frame.
Read the file into a data frame.	Specify a temporary folder to stage the data.
Perform transformations on the data frame.	Write the results to a table in Azure Synapse.

Explanation:

Mount the Data Lake Storage onto DBFS.

Read the file into a data frame.

Perform transformations on the data frame.

Specify a temporary folder to stage the data.

Write the results to a table in Azure Synapse.

Step 1: Mount the Data Lake Storage onto DBFS

Begin with creating a file system in the Azure Data Lake Storage Gen2 account.

Step 2: Read the file into a data frame.

You can load the json files as a data frame in Azure Databricks.

Step 3: Perform transformations on the data frame.

Step 4: Specify a temporary folder to stage the data

Specify a temporary folder to use while moving data between Azure Databricks and Azure Synapse.

Step 5: Write the results to a table in Azure Synapse.

You upload the transformed data frame into Azure Synapse. You use the Azure Synapse connector for Azure Databricks to directly upload a dataframe as a table in a Azure Synapse.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-databricks/databricks-extract-load-sql-data-warehouse>

### NEW QUESTION # 82

You are designing 2 solution that will use tables in Delta Lake on Azure Databricks.

You need to minimize how long it takes to perform the following:

\*Queries against non-partitioned tables

\* Joins on non-partitioned columns

Which two options should you include in the solution? Each correct answer presents part of the solution.

(Choose Correct Answer and Give Explanation and References to Support the answers based from Data Engineering on Microsoft Azure)

- A. dynamic file pruning (DFP)
- B. the clone command
- C. Apache Spark caching
- D. Z-Ordering

**Answer: C,D**

Explanation:

Explanation

According to the information I found on the web, two options that you should include in the solution to minimize how long it takes to perform queries and joins on non-partitioned tables are:

Z-Ordering: This is a technique to colocate related information in the same set of files. This co-locality is automatically used by Delta Lake in data-skipping algorithms. This behavior dramatically reduces the amount of data that Delta Lake on Azure Databricks needs to read<sup>123</sup>.

Apache Spark caching: This is a feature that allows you to cache data in memory or on disk for faster access. Caching can improve the performance of repeated queries and joins on the same data. You can cache Delta tables using the CACHE TABLE or CACHE LAZY commands.

To minimize the time it takes to perform queries against non-partitioned tables and joins on non-partitioned columns in Delta Lake on Azure Databricks, the following options should be included in the solution:

A: Z-Ordering: Z-Ordering improves query performance by co-locating data that share the same column values in the same physical partitions. This reduces the need for shuffling data across nodes during query execution. By using Z-Ordering, you can avoid full table scans and reduce the amount of data processed.

B: Apache Spark caching: Caching data in memory can improve query performance by reducing the amount of data read from disk. This helps to speed up subsequent queries that need to access the same data. When you cache a table, the data is read from the data source and stored in memory. Subsequent queries can then read the data from memory, which is much faster than reading it from disk.

References:

Delta Lake on Databricks: <https://docs.databricks.com/delta/index.html>

Best Practices for Delta Lake on

Databricks: <https://databricks.com/blog/2020/05/14/best-practices-for-delta-lake-on-databricks.html>

### NEW QUESTION # 83

.....

One more thing to give you an idea about the top features of Data Engineering on Microsoft Azure (DP-203) exam questions before purchasing, the ExamDumpsVCE are offering free ExamDumpsVCE DP-203 Exam Questions demo download facility. This facility is being offered in all three ExamDumpsVCE DP-203 exam practice question formats.

**Test DP-203 Registration:** <https://www.examdumpsvce.com/DP-203-valid-exam-dumps.html>

If you need 100% passing rate, our DP-203 guide torrent material will be the right one suitable for you, If you purchasing the DP-203 test practice files designed by many experts and professors from our company, we can promise that our online workers are going to serve you day and night during your learning period, The web-based DP-203 practice exam software is genuine, authentic, and real so feel free to start your practice instantly with DP-203 practice test.

If only two commands make sense at a given time, you've got two commands in the list, Source code for the book's examples, If you need 100% passing rate, our DP-203 guide torrent material will be the right one suitable for you.

**Valid DP-203 Practice Mock Help You to Get Acquainted with Real DP-203 Exam Simulation**

If you purchasing the DP-203 Test Practice files designed by many experts and professors from our company, we can promise that our online workers are going to serve you day and night during your learning period.

The web-based DP-203 practice exam software is genuine, authentic, and real so feel free to start your practice instantly with DP-203 practice test, Getting certification DP-203 is a good exam if you are just starting with Microsoft and the cloud.

It made them utterly confident to go through the whole process of the Data Engineering on Microsoft Azure. Feel free to compare our quality of Microsoft DP-203 exam questions dumps with other courses.

- [illegible]

2025 Latest ExamDumpsVCE DP-203 PDF Dumps and DP-203 Exam Engine Free Share: <https://drive.google.com/open?id=1O23QMFRs9XNv42dvKEzciRBH98Frg6jf>