Microsoft DP-700 Vce Download - Test DP-700 Guide

2026 Latest Getcertkey DP-700 PDF Dumps and DP-700 Exam Engine Free Share: https://drive.google.com/open?id=1UJIwWblC3sQMh8V8xqaTntk3gQztOTpX

You should prepare with Getcertkey DP-700 Questions that are in compliance with DP-700 exam content. More than 90,000 professionals worldwide have provided their feedback, helping create and launch DP-700 questions in the market. So, if you're determined to pass the Microsoft exam and achieve DP-700 Certification to accelerate your career, it's time to build your knowledge and skills. You can try the demo version of Implementing Data Engineering Solutions Using Microsoft Fabric (DP-700) practice dumps before payment.

Microsoft DP-700 Exam Syllabus Topics:

Topic	Details
Topic 1	• Ingest and transform data: This section of the exam measures the skills of Data Engineers that cover designing and implementing data loading patterns. It emphasizes preparing data for loading into dimensional models, handling batch and streaming data ingestion, and transforming data using various methods. A skill to be measured is applying appropriate transformation techniques to ensure data quality.
Topic 2	• Implement and manage an analytics solution: This section of the exam measures the skills of Microsoft Data Analysts regarding configuring various workspace settings in Microsoft Fabric. It focuses on setting up Microsoft Fabric workspaces, including Spark and domain workspace configurations, as well as implementing lifecycle management and version control. One skill to be measured is creating deployment pipelines for analytics solutions.
Topic 3	• Monitor and optimize an analytics solution: This section of the exam measures the skills of Data Analysts in monitoring various components of analytics solutions in Microsoft Fabric. It focuses on tracking data ingestion, transformation processes, and semantic model refreshes while configuring alerts for error resolution. One skill to be measured is identifying performance bottlenecks in analytics workflows.

>> Microsoft DP-700 Vce Download <<

Test DP-700 Guide & DP-700 Valid Test Camp

We don't just want to make profitable deals, but also to help our users pass the exams with the least amount of time to get DP-700 certificate. Choosing our DP-700 exam practice, you only need to spend 20-30 hours to prepare for the exam. Maybe you will ask whether such a short time can finish all the content, we want to tell you that you can rest assured ,because our DP-700 Learning Materials are closely related to the exam outline and the questions of our DP-700 guide questions are related to the latest and basic knowledge. You will pass the DP-700 exam only with our DP-700 exam questions.

Microsoft Implementing Data Engineering Solutions Using Microsoft Fabric

Sample Questions (Q40-Q45):

NEW QUESTION # 40
You have a Fabric workspace that contains an eventhouse and a KQL database named Database1. Database1 has the following:
A table named Table1
A table named Table2
An update policy named Policy1
Policy1 sends data from Table1 to Table2.
The following is a sample of the data in Table2.
Recently, the following actions were performed on Table1:
An additional element named temperature was added to the StreamData column.
The data type of the Timestamp column was changed to date.
The data type of the DeviceId column was changed to string.
You plan to load additional records to Table2.
Which two records will load from Table1 to Table2? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

- A. ☐
- C. ☐
- C. ☐
- D. ☐

Answer: C,D

Explanation:
Changes to Table1 Structure:
StreamData column: An additional temperature element was added.
Timestamp column: Data type changed from datetime to date.
DeviceId column: Data type changed from guid to string.
Impact of Changes:
Only records that comply with Table2's structure will load.
Records that deviate from Table2's column data types or structure will be rejected.
Record B:
Timestamp: Matches Table2 (datetime format).
DeviceId: Matches Table2 (guid format).
StreamData: Contains only the index and eventid, which matches Table2.
Accepted because it fully matches Table2's structure and data types.
Record D:
Timestamp: Matches Table2 (datetime format).
DeviceId: Matches Table2 (guid format).
StreamData: Matches Table2's structure.
Accepted because it fully matches Table2's structure and data types.

NEW QUESTION # 41
You have an Azure Data Lake Storage Gen2 account named storage1 and an Amazon S3 bucket named storage2.
You have the Delta Parquet files shown in the following table.
You have a Fabric workspace named Workspace1 that has the cache for shortcuts enabled. Workspace1 contains a lakehouse named Lakehouse1. Lakehouse1 has the following shortcuts:
A shortcut to ProductFile aliased as Products
A shortcut to StoreFile aliased as Stores
A shortcut to TripsFile aliased as Trips
The data from which shortcuts will be retrieved from the cache?

- A. Trips and Stores only
- B. Products only
- C. Stores only
- D. Products and Store only
- E. Products. Stores, and Trips

Answer: D

Explanation:
When the cache for shortcuts is enabled in Fabric, the data retrieval is governed by the caching behavior, which generally retains data for a specific period after it was last accessed. The data from the shortcuts will be retrieved from the cache if the data is stored in locations that support caching. Here's a breakdown based on the data's location:
Products: The ProductFile is stored in Azure Data Lake Storage Gen2 (storage1). Since Azure Data Lake is a supported storage system in Fabric and the file is relatively small (50 MB), this data is most likely cached and can be retrieved from the cache.
Stores: The StoreFile is stored in Amazon S3 (storage2), and even though it is stored in a different cloud provider, Fabric can cache data from Amazon S3 if caching is enabled. This data (25 MB) is likely cached and retrievable.
Trips: The TripsFile is stored in Amazon S3 (storage2) and is significantly larger (2 GB) compared to the other files. While Fabric can cache data from Amazon S3, the larger size of the file (2 GB) may exceed typical cache sizes or retention windows, causing this file to likely be retrieved directly from the source instead of the cache.

NEW QUESTION # 42
You have an Azure Data Lake Storage Gen2 account named storage1 and an Amazon S3 bucket named storage2.
You have the Delta Parquet files shown in the following table.

You have a Fabric workspace named Workspace1 that has the cache for shortcuts enabled. Workspace1 contains a lakehouse named Lakehouse1. Lakehouse1 has the following shortcuts:
The data from which shortcuts will be retrieved from the cache?

- A. Trips and Stores only
- B. Products only
- C. Stores only
- D. Products and Store only
- E. Products. Stores, and Trips

Answer: D

Explanation:
When the cache for shortcuts is enabled in Fabric, the data retrieval is governed by the caching behavior, which generally retains data for a specific period after it was last accessed. The data from the shortcuts will be retrieved from the cache if the data is stored in locations that support caching. Here's a breakdown based on the data's location:
Products: The ProductFile is stored in Azure Data Lake Storage Gen2 (storage1). Since Azure Data Lake is a supported storage system in Fabric and the file is relatively small (50 MB), this data is most likely cached and can be retrieved from the cache.
Stores: The StoreFile is stored in Amazon S3 (storage2), and even though it is stored in a different cloud provider, Fabric can cache data from Amazon S3 if caching is enabled. This data (25 MB) is likely cached and retrievable.
Trips: The TripsFile is stored in Amazon S3 (storage2) and is significantly larger (2 GB) compared to the other files. While Fabric can cache data from Amazon S3, the larger size of the file (2 GB) may exceed typical cache sizes or retention windows, causing this file to likely be retrieved directly from the source instead of the cache.

NEW QUESTION # 43
Exhibit.
You have a Fabric workspace that contains a write-intensive warehouse named DW1. DW1 stores staging tables that are used to load a dimensional model. The tables are often read once, dropped, and then recreated to process new data.
You need to minimize the load time of DW1.
What should you do?

- A. Drop statistics.
- B. Enable V-O-der.
- C. Create statistics.
- D. Disable V-Order.

Answer: B

NEW QUESTION # 44
You have a Fabric workspace that contains a lakehouse named Lakehouse1. Lakehouse1 contains a Delta table named Table1.

You analyze Table1 and discover that Table1 contains 2,000 Parquet files of 1 MB each.
You need to minimize how long it takes to query Table1.
What should you do?

- A. Disable V-Order and run the OPTIMIZE command.
- B. Run the OPTIMIZE and VACUUM commands.
- C. Disable V-Order and run the VACUUM command.

Answer: B

Explanation:
Problem Overview:
Solution:
Commands and Their Roles:
- Compacts small Parquet files into larger files to improve query performance.
- It supports optional features like V-Order, which organizes data for efficient scanning.
- Removes old, unreferenced data files and metadata from the Delta table.
- Running VACUUM after OPTIMIZE ensures unnecessary files are cleaned up, reducing storage overhead and improving performance.

NEW QUESTION # 45
......

Another version of Implementing Data Engineering Solutions Using Microsoft Fabric (DP-700) practice exams is also available at Getcertkey and that is web-based. It has all specifications we have discussed above in the section of the Microsoft DP-700 desktop practice test software. But the only difference is that this web-based DP-700 practice exam software works online and needs no software installation. Furthermore, this DP-700 Practice Exam is supported by both Windows and iOS, Android, Mac, and Linux. Since it is the web-based DP-700 practice exam, you can take it from Opera, Chrome, Safari, Firefox, or any other popular browser.

Test DP-700 Guide: https://www.getcertkey.com/DP-700_braindumps.html

id=1UJIwWblC3sQMh8V8xqaTntk3gQztOTpX