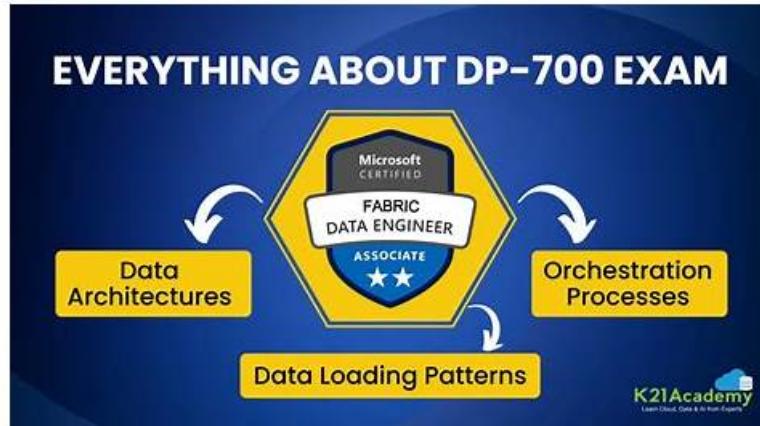


DP-700 Valid Exam Experience & DP-700 Valid Test Questions



P.S. Free & New DP-700 dumps are available on Google Drive shared by DumpStillValid: <https://drive.google.com/open?id=1xFveOm83UjyChK972sFirOIW17ypaVaU>

We are constantly updating our Microsoft DP-700 practice material to ensure that students receive the latest DP-700 questions based on the actual Implementing Data Engineering Solutions Using Microsoft Fabric exam content. Moreover, we also offer up to 1 year of free updates and free demos. DumpStillValid also offers a money-back guarantee (terms and conditions apply) for applicants who fail to pass the DP-700 test on the first try.

Microsoft DP-700 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none">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 2	<ul style="list-style-type: none">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.
Topic 3	<ul style="list-style-type: none">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.

>> DP-700 Valid Exam Experience <<

DP-700 Valid Test Questions, DP-700 Exam Training

Buying any product should choose a trustworthy company. Our DumpStillValid can give you the promise of the highest pass rate of DP-700 exam; we can give you a promise to try our DP-700 software for free, and the promise of free updates within a year after purchase. To resolve your doubts, we assure you that if you regrettably fail the DP-700 Exam, we will full refund all the cost you buy our study materials. DumpStillValid is your best partners in your preparation for DP-700 exam.

Microsoft Implementing Data Engineering Solutions Using Microsoft Fabric

Sample Questions (Q81-Q86):

NEW QUESTION # 81

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a Fabric eventstream that loads data into a table named Bike_Location in a KQL database. The table contains the following columns:

BikepointID
Street
Neighbourhood
No_Bikes
No_Empty_Docks
Timestamp

You need to apply transformation and filter logic to prepare the data for consumption. The solution must return data for a neighbourhood named Sands End when No_Bikes is at least 15. The results must be ordered by No_Bikes in ascending order. Solution: You use the following code segment:

```
bike_location
| filter Neighbourhood == "Sands End" and No_Bikes >= 15
| order by No_Bikes
| project BikepointID, Street, Neighbourhood, No_Bikes, No_Empty_Docks, Timestamp
```

Does this meet the goal?

- A. Yes
- B. no

Answer: B

Explanation:

This code does not meet the goal because it uses order by, which is not valid in KQL. The correct term in KQL is sort by. Correct code should look like:

```
Microsoft
bike_location
| filter Neighbourhood == "Sands End" and No_Bikes >= 15
| sort by No_Bikes asc
| project BikepointID, Street, Neighbourhood, No_Bikes, No_Empty_Docks, Timestamp
```

NEW QUESTION # 82

You need to develop an orchestration solution in fabric that will load each item one after the other. The solution must be scheduled to run every 15 minutes. Which type of item should you use?

- A. warehouse
- B. data pipeline
- C. Dataflow Gen2 dataflow
- D. notebook

Answer: B

NEW QUESTION # 83

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.

Name	Stored in	Size	Microsoft	Description
ProductFile	storage1	50 MB		Contains a list of products and their details
TripsFile	storage2	2 GB		Contains one month's worth of taxi trip data
StoreFile	storage2	25 MB		Contains a list of stores and their addresses

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

Answer: A

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 # 84

You have a Fabric workspace named Workspace1 that contains the items shown in the following table.

Name	Type
Notebook1	Notebook
Notebook2	Notebook
Lakehouse1	Lakehouse
Pipeline1	Data pipeline
Model1	Semantic model

For Model1, the Keep your Direct Lake data up to date option is disabled.

You need to configure the execution of the items to meet the following requirements:

How should you orchestrate each item? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Notebook1: Add Notebook1 to an Apache Spark job definition.
Add Notebook1 to Pipeline1.
From Real-Time hub, configure the execution of Notebook1

Notebook2: Add Notebook2 to an Apache Spark job definition.
Add Notebook2 to Pipeline1.
From Real-Time hub, configure the execution of Notebook2

Pipeline1: Add Pipeline1 to an Apache Spark job definition.
Configure the execution of Pipeline1 by using a schedule.
From Real-Time hub, configure the execution of Pipeline1.

Model1: Add Model1 to Pipeline1.
From Real-Time hub, configure Model1 to refresh.
Set Keep your Direct Lake data up to date to On.

Answer:

Explanation:

Microsoft

Notebook1: Add Notebook1 to an Apache Spark job definition.
Add Notebook1 to Pipeline1.
From Real-Time hub, configure the execution of Notebook1

Notebook2: Add Notebook2 to an Apache Spark job definition.
Add Notebook2 to Pipeline1.
From Real-Time hub, configure the execution of Notebook2

Pipeline1: Add Pipeline1 to an Apache Spark job definition.
Configure the execution of Pipeline1 by using a schedule.
From Real-Time hub, configure the execution of Pipeline1.

Model1: Add Model1 to Pipeline1.
From Real-Time hub, configure Model1 to refresh.
Set Keep your Direct Lake data up to date to On.

NEW QUESTION # 85

You have a Fabric notebook named Notebook1 that has been executing successfully for the last week.

During the last run, Notebook1 executed nine jobs.

You need to view the jobs in a timeline chart.

What should you use?

- A. Spark History Server
- B. Real-Time hub
- C. Monitoring hub
- D. the run series from the details of the application run**
- E. the job history from the application run

Answer: D

Explanation:

The run series from the details of the application run is the most detailed and relevant feature for visualizing job execution in a timeline format, making it the correct choice for this scenario. It provides an intuitive way to analyze job execution patterns and improve the efficiency of the notebook.

NEW QUESTION # 86

• • • • •

The Microsoft DP-700 practice exam software will provide you with feedback on your performance. The Microsoft DP-700 practice test software also includes a built-in timer and score tracker so students can monitor their progress. DP-700 Practice Exam enables applicants to practice time management, answer strategies, and all other elements of the final Microsoft DP-700 certification exam and can check their scores.

DP-700 Valid Test Questions: <https://www.dumpstillovalid.com/DP-700-prep4sure-review.html>

2025 Latest DumpStillValid DP-700 PDF Dumps and DP-700 Exam Engine Free Share: <https://drive.google.com/open?id=1xFveOm83UjyChK972sFirOIW17ypaVaU>