

Microsoft - DP-700 - High Pass-Rate Implementing Data Engineering Solutions Using Microsoft Fabric Reliable Dumps Pdf



2026 Latest FreePdfDump DP-700 PDF Dumps and DP-700 Exam Engine Free Share: <https://drive.google.com/open?id=1YmT5kXwq0FvThxsBskiPFp1Vy6hLTol>

The scoring system of our DP-700 exam torrent absolutely has no problem because it is intelligent and powerful. First of all, our researchers have made lots of efforts to develop the scoring system. So the scoring system of the DP-700 test answers can stand the test of practicability. Once you have submitted your practice. The scoring system will begin to count your marks of the DP-700 Exam guides quickly and correctly. At the same time, there is specific space below every question for you to make notes. So you can quickly record the important points or confusion of the DP-700 exam guides.

Microsoft DP-700 Exam Syllabus Topics:

Topic	Details
Topic 1	<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 2	<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 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.

Latest DP-700 Real Test | DP-700 Exam Certification Cost

The most important is that you just only need to spend 20 to 30 hours on practicing DP-700 exam questions before you take the exam, therefore you can arrange your time to balance learning and other things. Of course, you care more about your test pass rate. We offer you more than 99% pass guarantee if you are willing to use our DP-700 Test Guide and follow our plan of learning. And if you want to pass the DP-700 exam, you should choose our DP-700 torrent prep to help you. And We will update DP-700 learning materials to make sure you have the latest questions and answers.

Microsoft Implementing Data Engineering Solutions Using Microsoft Fabric Sample Questions (Q92-Q97):

NEW QUESTION # 92

You have a Fabric workspace.

You have semi-structured data.

You need to read the data by using T-SQL, KQL, and Apache Spark. The data will only be written by using Spark. What should you use to store the data?

- A. an eventhouse
- **B. a lakehouse**
- C. a warehouse
- D. a datamart

Answer: B

Explanation:

A lakehouse is the best option for storing semi-structured data when you need to read it using T-SQL, KQL, and Apache Spark. A lakehouse combines the flexibility of a data lake (which can handle semi-structured and unstructured data) with the performance features of a data warehouse. It allows data to be written using Apache Spark and can be queried using different technologies such as T-SQL (for SQL-based querying), KQL (Kusto Query Language for querying), and Apache Spark (for distributed processing). This solution is ideal when dealing with semi-structured data and requiring a versatile querying approach.

NEW QUESTION # 93

You have a Google Cloud Storage (GCS) container named storage1 that contains the files shown in the following table.

Name	Size
ProductFile.parquet	8 MB
StoreFile.json	500 MB
TripsFile.csv	99 MB

You have a Fabric workspace named Workspace1 that has the cache for shortcuts enabled. Workspace1 contains a lakehouse named Lakehouse1. Lakehouse1 has the shortcuts shown in the following table.

Name	Source	Last accessed
Products	ProductFile	12 hours ago
Stores	Storefile	4 hours ago
Trips	TripsFile	48 hours ago

You need to read data from all the shortcuts.

Which shortcuts will retrieve data from the cache?

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

Answer: D

Explanation:

When reading data from shortcuts in Fabric (in this case, from a lakehouse like Lakehouse1), the cache for shortcuts helps by storing the data locally for quick access. The last accessed timestamp and the cache expiration rules determine whether data is fetched from the cache or from the source (Google Cloud Storage, in this case).

Products: The ProductFile.parquet was last accessed 12 hours ago. Since the cache has data available for up to 12 hours, it is likely that this data will be retrieved from the cache, as it hasn't been too long since it was last accessed.

Stores: The StoreFile.json was last accessed 4 hours ago, which is within the cache retention period.

Therefore, this data will also be retrieved from the cache.

Trips: The TripsFile.csv was last accessed 48 hours ago. Given that it's outside the typical caching window (assuming the cache has a maximum retention period of around 24 hours), it would not be retrieved from the cache. Instead, it will likely require a fresh read from the source.

NEW QUESTION # 94

You have a Fabric F32 capacity that contains a workspace. The workspace contains a warehouse named DW1 that is modelled by using MD5 hash surrogate keys.

DW1 contains a single fact table that has grown from 200 million rows to 500 million rows during the past year.

You have Microsoft Power BI reports that are based on Direct Lake. The reports show year-over-year values.

Users report that the performance of some of the reports has degraded over time and some visuals show errors.

You need to resolve the performance issues. The solution must meet the following requirements:

Provide the best query performance.

Minimize operational costs.

Which should you do?

- A. Create views.
- B. Modify the surrogate keys to use a different data type.
- C. Change the MD5 hash to SHA256.
- D. Increase the capacity.
- E. Enable V-Order

Answer: A

Explanation:

In this case, the key issue causing performance degradation likely stems from the use of MD5 hash surrogate keys. MD5 hashes are 128-bit values, which can be inefficient for large datasets like the 500 million rows in your fact table. Using a more efficient data type for surrogate keys (such as integer or bigint) would reduce the storage and processing overhead, leading to better query performance. This approach will improve performance while minimizing operational costs because it reduces the complexity of querying and indexing, as smaller data types are generally faster and more efficient to process.

NEW QUESTION # 95

HOTSPOT

You have a Fabric workspace named Workspace1_DEV that contains the following items:

10 reports

Four notebooks

Three lakehouses

Two data pipelines

Two Dataflow Gen1 dataflows

Three Dataflow Gen2 dataflows

Five semantic models that each has a scheduled refresh policy

You create a deployment pipeline named Pipeline1 to move items from Workspace1_DEV to a new workspace named Workspace1_TEST.

You deploy all the items from Workspace1_DEV to Workspace1_TEST.

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer Area

Statements

Data from the semantic models will be deployed to the target stage.

Yes

No

The Dataflow Gen1 dataflows will be deployed to the target stage.

The scheduled refresh policies will be deployed to **the target stage**.



Answer:

Explanation:

ANSWER Area

Statements

Data from the semantic models will be deployed to the target stage.

Yes

No

The Dataflow Gen1 dataflows will be deployed to the target stage.

The scheduled refresh policies **will** be deployed to the target stage.

Explanation:

Answer Area

Statements

Data from the semantic models will be deployed to the target stage.

Yes

No

The Dataflow Gen1 dataflows will be deployed to the target stage.

The scheduled refresh policies **will** be deployed to the target stage.

NEW QUESTION # 96

You need to recommend a method to populate the POS1 data to the lakehouse medallion layers.

What should you recommend for each layer? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Bronze layer:

- A Dataflow Gen2 dataflow
- A notebook
- A pipeline Copy activity
- A pipeline stored procedure

 Microsoft
Silver layer:

- A Dataflow Gen2 dataflow
- A notebook
- A pipeline Copy activity
- A pipeline stored procedure

Answer:

Explanation:

Bronze layer:

- A Dataflow Gen2 dataflow
- A notebook
- A pipeline Copy activity
- A pipeline stored procedure

Silver layer:

- A Dataflow Gen2 dataflow
- A notebook
- A pipeline Copy activity
- A pipeline stored procedure

Explanation:

A screenshot of a computer Description automatically generated

Bronze layer:

- A Dataflow Gen2 dataflow
- A notebook
- A pipeline Copy activity
- A pipeline stored procedure

Silver layer:

- A Dataflow Gen2 dataflow
- A notebook
- A pipeline Copy activity
- A pipeline stored procedure

Bronze Layer: A pipeline Copy activity

The bronze layer is used to store raw, unprocessed data. The requirements specify that no transformations should be applied before

landing the data in this layer. Using a pipeline Copy activity ensures minimal development effort, built-in connectors, and the ability to ingest the data directly into the Delta format in the bronze layer.

Silver Layer: A notebook

The silver layer involves extensive data cleansing (deduplication, handling missing values, and standardizing capitalization). A notebook provides the flexibility to implement complex transformations and is well-suited for this task.

NEW QUESTION # 97

You don't have to worry about your problems on our DP-700 exam questions are too much or too simple. Our staff will give you a smile and then answer them carefully. All we do is just want you to concentrate on learning on our DP-700 study guide! Let other things go to us. And as long as you focus on our DP-700 Training Materials, we believe you will pass for sure for our DP-700 practice braindumps are always the latest and valid for all of our customers.

Latest DP-700 Real Test: <https://www.freepdfdump.top/DP-700-valid-torrent.html>

2026 Latest FreePdfDump DP-700 PDF Dumps and DP-700 Exam Engine Free Share: <https://drive.google.com/open?id=1YmT5kXwq0FvThxsBskjPFp1Vv6hLTol>