


New DP-700 Test Duration - 100% Pass Quiz 2025 First-grade Microsoft DP-700: Test Implementing Data Engineering Solutions Using Microsoft Fabric Questions



BONUS!!! Download part of RealVCE DP-700 dumps for free: <https://drive.google.com/open?id=1PoLv0rbQZetr1bPEqPOh8O7EKYIj4I5L>

RealVCE also presents desktop-based Microsoft DP-700 practice test software which is usable without any internet connection after installation and only required license verification. Microsoft DP-700 practice test software is very helpful for all those who desire to practice in an actual Implementing Data Engineering Solutions Using Microsoft Fabric (DP-700) exam-like environment. Implementing Data Engineering Solutions Using Microsoft Fabric (DP-700) practice test is customizable so that you can change the timings of each session. RealVCE desktop Microsoft DP-700 practice test questions software is only compatible with windows and easy to use for everyone.

Microsoft DP-700 Exam Syllabus Topics:

Topic	Details
Topic 1	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	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	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.

>> New DP-700 Test Duration <<

Test DP-700 Questions | DP-700 Exam Cost

The pass rate is 98.95% for the DP-700 training materials, and most candidates can pass the exam just one time. We ensure you that you will refund your money if you fail to pass the exam. In addition, we offer you free update for one year, and the update

version for the DP-700 exam dumps will be sent to your email automatically, so that you can know the latest information about the DP-700 Exam Dumps. We provide you with the online chat service, and in the process of learning, if you have any questions about the DP-700 exam dumps, you can consult us.

Microsoft Implementing Data Engineering Solutions Using Microsoft Fabric Sample Questions (Q21-Q26):

NEW QUESTION # 21

You need to resolve the sales data issue. The solution must minimize the amount of data transferred. What should you do?

- A. Configure incremental refresh for the dataflow. Set Store rows from the past to 1 Month.
- B. Configure incremental refresh for the dataflow. Set Refresh rows from the past to 1 Year.
- C. Configure scheduled refresh for the dataflow.
- D. Split the dataflow into two dataflows.
- E. Configure incremental refresh for the dataflow. Set Refresh rows from the past to 1 Month.

Answer: E

Explanation:

The sales data issue can be resolved by configuring incremental refresh for the dataflow. Incremental refresh allows for only the new or changed data to be processed, minimizing the amount of data transferred and improving performance.

The solution specifies that data older than one month never changes, so setting the refresh period to 1 Month is appropriate. This ensures that only the most recent month of data will be refreshed, reducing unnecessary data transfers.

Topic 2, Contoso, LtdCase Study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview. Company Overview

Contoso, Ltd. is an online retail company that wants to modernize its analytics platform by moving to Fabric.

The company plans to begin using Fabric for marketing analytics.

Overview. IT Structure

The company's IT department has a team of data analysts and a team of data engineers that use analytics systems.

The data engineers perform the ingestion, transformation, and loading of data. They prefer to use Python or SQL to transform the data.

The data analysts query data and create semantic models and reports. They are qualified to write queries in Power Query and T-SQL.

Existing Environment. Fabric

Contoso has an F64 capacity named Cap1. All Fabric users are allowed to create items.

Contoso has two workspaces named WorkspaceA and WorkspaceB that currently use Pro license mode.

Existing Environment. Source Systems

Contoso has a point of sale (POS) system named POS1 that uses an instance of SQL Server on Azure Virtual Machines in the same Microsoft Entra tenant as Fabric. The host virtual machine is on a private virtual network that has public access blocked. POS1 contains all the sales transactions that were processed on the company's website.

The company has a software as a service (SaaS) online marketing app named MAR1. MAR1 has seven entities. The entities contain data that relates to email open rates and interaction rates, as well as website interactions. The data can be exported from MAR1 by calling REST APIs. Each entity has a different endpoint.

Contoso has been using MAR1 for one year. Data from prior years is stored in Parquet files in an Amazon Simple Storage Service (Amazon S3) bucket. There are 12 files that range in size from 300 MB to 900 MB and relate to email interactions.

Existing Environment. Product Data

POS1 contains a product list and related data. The data comes from the following three tables:

Products

ProductCategories

ProductSubcategories

In the data, products are related to product subcategories, and subcategories are related to product categories.

Existing Environment. Azure

Contoso has a Microsoft Entra tenant that has the following mail-enabled security groups:

DataAnalysts: Contains the data analysts

DataEngineers: Contains the data engineers

Contoso has an Azure subscription.

The company has an existing Azure DevOps organization and creates a new project for repositories that relate to Fabric.

Existing Environment. User Problems

The VP of marketing at Contoso requires analysis on the effectiveness of different types of email content. It typically takes a week to manually compile and analyze the data. Contoso wants to reduce the time to less than one day by using Fabric.

The data engineering team has successfully exported data from MAR1. The team experiences transient connectivity errors, which causes the data exports to fail.

Requirements. Planned Changes

Contoso plans to create the following two lakehouses:

Lakehouse1: Will store both raw and cleansed data from the sources

Lakehouse2: Will serve data in a dimensional model to users for analytical queries Additional items will be added to facilitate data ingestion and transformation.

Contoso plans to use Azure Repos for source control in Fabric.

Requirements. Technical Requirements

The new lakehouses must follow a medallion architecture by using the following three layers: bronze, silver, and gold. There will be extensive data cleansing required to populate the MAR1 data in the silver layer, including deduplication, the handling of missing values, and the standardizing of capitalization.

Each layer must be fully populated before moving on to the next layer. If any step in populating the lakehouses fails, an email must be sent to the data engineers.

Data imports must run simultaneously, when possible.

The use of email data from the Amazon S3 bucket must meet the following requirements:

Minimize egress costs associated with cross-cloud data access.

Prevent saving a copy of the raw data in the lakehouses.

Items that relate to data ingestion must meet the following requirements:

The items must be source controlled alongside other workspace items.

Ingested data must land in the bronze layer of Lakehouse1 in the Delta format.

No changes other than changes to the file formats must be implemented before the data lands in the bronze layer.

Development effort must be minimized and a built-in connection must be used to import the source data.

In the event of a connectivity error, the ingestion processes must attempt the connection again.

Lakehouses, data pipelines, and notebooks must be stored in WorkspaceA. Semantic models, reports, and dataflows must be stored in WorkspaceB.

Once a week, old files that are no longer referenced by a Delta table log must be removed.

Requirements. Data Transformation

In the POS1 product data, ProductID values are unique. The product dimension in the gold layer must include only active products from product list. Active products are identified by an IsActive value of 1.

Some product categories and subcategories are NOT assigned to any product. They are NOT analytically relevant and must be omitted from the product dimension in the gold layer.

Requirements. Data Security

Security in Fabric must meet the following requirements:

The data engineers must have read and write access to all the lakehouses, including the underlying files.

The data analysts must only have read access to the Delta tables in the gold layer.

The data analysts must NOT have access to the data in the bronze and silver layers.

The data engineers must be able to commit changes to source control in WorkspaceA.

NEW QUESTION # 22

You have a Fabric workspace that contains a lakehouse named Lakehouse1.

In an external data source, you have data files that are 500 GB each. A new file is added every day.

You need to ingest the data into Lakehouse1 without applying any transformations. The solution must meet the following requirements Trigger the process when a new file is added.

Provide the highest throughput.

Which type of item should you use to ingest the data?

- A. Dataflow Gen2
- B. Environment
- **C. Data pipeline**
- D. KQL queryset

Answer: C

Explanation:

To efficiently ingest large data files (500 GB each) into Lakehouse1 with high throughput and trigger the process when a new file is added, a Data pipeline is the most suitable solution. Data pipelines in Fabric are ideal for orchestrating data movement and can be configured to automatically trigger based on file arrivals or other events. This solution meets both requirements: ingesting the data without transformations (since you just need to copy the data) and triggering the process when new files are added.

NEW QUESTION # 23

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:

```
SELECT BikepointID, Street, Neighbourhood, No_Bikes, No_Empty_Docks, Timestamp
FROM bike_location
WHERE neighbourhood = 'Sands End'
AND no_bikes >= 15
ORDER BY no_bikes
```

Does this meet the goal?

- A. Yes
- **B. no**

Answer: B

Explanation:

This code does not meet the goal because this is an SQL-like query and cannot be executed in KQL, which is required for the database.

Correct code should look like:

```
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 # 24

HOTSPOT

You need to troubleshoot the ad-hoc query issue.

How should you complete the statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

SELECT last_run_start_time, last_run_command

FROM

queryinsights.exec_requests_history
queryinsights.exec_sessions_history
queryinsights.frequently_run_queries
queryinsights.long_running_queries

WHERE last_run_total_elapsed_time_ms > 7200000

AND

max_run_total_elapsed_time_ms > 7200000
median_total_elapsed_time_ms > 7200000
number_of_canceled_runs > 1
number_of_failed_runs > 1
number_of_runs > 1

Answer:

Explanation:

SELECT last_run_start_time, last_run_command

FROM

queryinsights.exec_requests_history
queryinsights.exec_sessions_history
queryinsights.frequently_run_queries
queryinsights.long_running_queries

WHERE last_run_total_elapsed_time_ms > 7200000

AND

max_run_total_elapsed_time_ms > 7200000
median_total_elapsed_time_ms > 7200000
number_of_canceled_runs > 1
number_of_failed_runs > 1
number_of_runs > 1

Explanation:

A screenshot of a computer Description automatically generated

Answer Area

```
SELECT last_run_start_time, last_run_command
```

```
FROM
```

queryinsights.exec_requests_history
queryinsights.exec_sessions_history
queryinsights.frequently_run_queries
queryinsights.long_running_queries

```
WHERE last_run_total_elapsed_time_ms > 7200000
```

```
AND
```

max_run_total_elapsed_time_ms > 7200000
median_total_elapsed_time_ms > 7200000
number_of_canceled_runs > 1
number_of_failed_runs > 1
number_of_runs > 1

SELECT last_run_start_time, last_run_command: These fields will help identify the execution details of the long-running queries.

FROM queryinsights.long_running_queries: The correct solution is to check the long-running queries using the queryinsights.long_running_queries view, which provides insights into queries that take longer than expected to execute.

WHERE last_run_total_elapsed_time_ms > 7200000: This condition filters queries that took more than 2 hours to complete (7200000 milliseconds), which is relevant to the issue described.

AND number_of_failed_runs > 1: This condition is key for identifying queries that have failed more than once, helping to isolate the problematic queries that cause failures and need attention.

NEW QUESTION # 25

You have a Fabric workspace that contains a lakehouse and a notebook named Notebook1. Notebook1 reads data into a DataFrame from a table named Table1 and applies transformation logic. The data from the DataFrame is then written to a new Delta table named Table2 by using a merge operation.

You need to consolidate the underlying Parquet files in Table1.

Which command should you run?

- A. CACHE
- B. VACUUM
- C. BROADCAST
- D. OPTIMIZE

Answer: D

Explanation:

To consolidate the underlying Parquet files in Table1 and improve query performance by optimizing the data layout, you should use the OPTIMIZE command in Delta Lake. The OPTIMIZE command coalesces smaller files into larger ones and reorganizes the data for more efficient reads. This is particularly useful when working with large datasets in Delta tables, as it helps reduce the number of files and improves performance for subsequent queries or operations like MERGE.

NEW QUESTION # 26

.....

we will provide you with the best Microsoft DP-700 exam dumps. You can pass the Microsoft DP-700 exam with high marks with the help of the Microsoft DP-700 exam questions. These Microsoft DP-700 exam practice questions are designed and verified by experienced and qualified DP-700 Exam Preparation trainers. They work together and put all their expertise and knowledge while verifying DP-700 exam questions all the time.

Test DP-700 Questions: https://www.realvce.com/DP-700_free-dumps.html

- Test DP-700 Dumps.zip □ DP-700 Reliable Exam Simulator □ DP-700 PDF Question □ Download ► DP-700 □ for free by simply entering ✓ www.pdf4dumps.com □ ✓ □ website □ Valid DP-700 Test Duration
- New DP-700 Test Duration Exam Pass Certify | Test DP-700 Questions □ Download □ DP-700 □ for free by simply entering □ www.pdfvce.com □ website □ Latest Test DP-700 Simulations
- DP-700 New Exam Bootcamp □ Reliable DP-700 Exam Syllabus □ DP-700 Pdf Torrent ✓ □ Search for □ DP-700 □ and download it for free immediately on ► www.prep4pass.com □ □ DP-700 Learning Mode
- DP-700 New Exam Bootcamp □ DP-700 New Exam Bootcamp □ Test DP-700 Dumps.zip □ ⇒ www.pdfvce.com ⇐ is best website to obtain □ DP-700 □ for free download □ Test DP-700 Dumps.zip
- Latest DP-700 Exam Pass4sure □ DP-700 Exam Vce □ DP-700 Valid Exam Duration □ Go to website □ www.exam4pdf.com □ open and search for □ DP-700 □ to download for free □ DP-700 Learning Mode
- Ace the Microsoft DP-700 Exam Preparation with Exams Solutions Realistic Practice Tests □ Easily obtain □ DP-700 □ for free download through ► www.pdfvce.com □ □ □ DP-700 Online Tests
- New DP-700 Test Duration Exam Pass Certify | Test DP-700 Questions □ Search for ► DP-700 □ on “ www.pass4test.com ” immediately to obtain a free download □ DP-700 Pdf Torrent
- DP-700 Online Tests □ DP-700 Online Tests □ DP-700 New Exam Bootcamp x Copy URL ► www.pdfvce.com □ □ □ open and search for ▷ DP-700 ◁ to download for free □ Reliable DP-700 Exam Syllabus
- New DP-700 Test Duration Exam Pass Certify | Test DP-700 Questions □ Enter (www.vceengine.com) and search for “ DP-700 ” to download for free □ Valid DP-700 Test Duration
- Real Microsoft DP-700 Dumps PDF Format □ □ www.pdfvce.com □ is best website to obtain { DP-700 } for free download □ DP-700 PDF Question
- DP-700 Exam Vce □ DP-700 Online Tests □ DP-700 Valid Exam Pattern □ Search for ► DP-700 ◁ and download it for free immediately on ► www.exams4collection.com □ □ □ DP-700 New Exam Bootcamp
- ncon.edu.sa, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, pct.edu.pk, ncon.edu.sa, study.stcs.edu.np, shikhaw.com, digitalrepublix.com, lms.skitbi-cuet.com, Disposable vapes

P.S. Free 2025 Microsoft DP-700 dumps are available on Google Drive shared by RealVCE: <https://drive.google.com/open?id=1PoLv0rbQZetr1bPEqPOh8O7EKYIj4I5L>