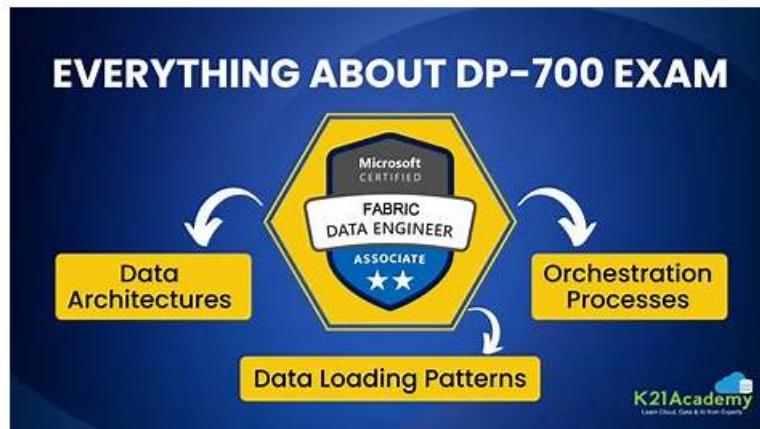


DP-700 Valid Test Cram, DP-700 Real Testing Environment



2026 Latest Dumpcollection DP-700 PDF Dumps and DP-700 Exam Engine Free Share: <https://drive.google.com/open?id=1VBUwVsx2jMwE4R8KSVZu9E9D00qQl4c2>

How can our DP-700 exam questions be the best exam materials in the field and always so popular among the candidates? There are two main reasons. First of all, we have a professional team of experts, each of whom has extensive experience on the DP-700 study guide. Secondly, before we write DP-700 Guide quiz, we collect a large amount of information and we will never miss any information points. Of course, we also fully consider the characteristics of the user. So we can make the best DP-700 learning questions.

Microsoft DP-700 Exam Syllabus Topics:

Topic	Details
Topic 1	<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.
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">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.

>> DP-700 Valid Test Cram <<

DP-700 Real Testing Environment, Test DP-700 Dumps

Our Microsoft DP-700 Exam Dumps effect in helping candidates' certification exam. Original questions are also important. These would provide a forum where certification training can be carried on. Our dumps torrent is perfect and practice test is also the latest. After you purchase our product, we offer free update service for one year.

Microsoft Implementing Data Engineering Solutions Using Microsoft Fabric

Sample Questions (Q41-Q46):

NEW QUESTION # 41

You have two Fabric workspaces named Workspace1 and Workspace2.

You have a Fabric deployment pipeline named deployPipeline1 that deploys items from Workspace1 to Workspace2.

DeployPipeline1 contains all the items in Workspace1.

You recently modified the items in Workspace1.

The workspaces currently contain the items shown in the following table.

Items in Workspace1 that have the same name as items in Workspace2 are currently paired.

You need to ensure that the items in Workspace1 overwrite the corresponding items in Workspace2. The solution must minimize effort.

What should you do?

- A. Back up the items in Workspace2, and then run deployPipeline1.
- B. **Run deployPipeline1 without modifying the items in Workspace2.**
- C. Delete all the items in Workspace2, and then run deployPipeline1.
- D. Rename each item in Workspace2 to have the same name as the items in Workspace1.

Answer: B

Explanation:

When running a deployment pipeline in Fabric, if the items in Workspace1 are paired with the corresponding items in Workspace2 (based on the same name), the deployment pipeline will automatically overwrite the existing items in Workspace2 with the modified items from Workspace1. There's no need to delete, rename, or back up items manually unless you need to keep versions. By simply running deployPipeline1, the pipeline will handle overwriting the existing items in Workspace2 based on the pairing, ensuring the latest version of the items is deployed with minimal effort.

NEW QUESTION # 42

You have two Fabric notebooks named Load_Salesperson and Load_Orders that read data from Parquet files in a lakehouse.

Load_Salesperson writes to a Delta table named dim_salesperson. Load.Orders writes to a Delta table named fact_orders and is dependent on the successful execution of Load_Salesperson.

You need to implement a pattern to dynamically execute Load_Salesperson and Load_Orders in the appropriate order by using a notebook.

How should you complete the code? To answer, drag the appropriate values the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

Explanation:

Topic 2, Contoso, Ltd Case 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 # 43

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

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 # 44

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. □
- B. □
- C. □
- D. □

Answer: A,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 # 45

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. Data pipeline**
- C. Environment
- D. KQL queryset

Answer: B

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 # 46

.....

DP-700 certification can demonstrate your mastery of certain areas of knowledge, which is internationally recognized and accepted by the general public as a certification. DP-700 certification is so high that it is not easy to obtain it. It requires you to invest time and energy. If you are not sure whether you can strictly request yourself, our DP-700 Exam Training can help you. Help is to arrange time for you and provide you with perfect service. If you use our learning materials to achieve your goals, we will be honored. DP-700 exam prep look forward to meeting you.

DP-700 Real Testing Environment: https://www.dumpcollection.com/DP-700_braindumps.html

- DP-700 Exam Cram Review Accurate DP-700 Test DP-700 Valid Test Dumps ➡ www.testkingpass.com is best website to obtain [DP-700] for free download DP-700 Reliable Braindumps Ppt
- DP-700 New Questions DP-700 Popular Exams DP-700 Reliable Braindumps Ppt Search for [DP-700] and download it for free immediately on ➡ www.pdfvce.com DP-700 Latest Guide Files
- Pass Guaranteed Microsoft - Newest DP-700 - Implementing Data Engineering Solutions Using Microsoft Fabric Valid Test Cram Enter « www.prepawayete.com » and search for ➡ DP-700 ⇌ to download for free DP-700 New Questions
- DP-700 Exam Guide - DP-700 Test Questions - DP-700 Exam Torrent Go to website www.pdfvce.com open and search for DP-700 to download for free DP-700 Official Cert Guide

P.S. Free 2026 Microsoft DP-700 dumps are available on Google Drive shared by Dumpcollection: <https://drive.google.com/open?id=1VBUwVs2jMwE4R8KSVzu9E9D00qQl4c2>