

Valid Microsoft DP-700 Test Sims, DP-700 New Exam Camp



P.S. Free 2026 Microsoft DP-700 dumps are available on Google Drive shared by Dumps4PDF: <https://drive.google.com/open?id=1QxS0B3wLA6K7qSJJTNDnUpL8Zx5jLyEf>

All these three Dumps4PDF Implementing Data Engineering Solutions Using Microsoft Fabric (DP-700) exam questions formats are easy to use and perfectly work with all devices, operating systems, and the latest web browsers. So rest assured that with the DP-700 Exam Dumps you will get everything that you need to learn, prepare and pass the challenging DP-700 exam with good scores.

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.

>> Valid Microsoft DP-700 Test Sims <<

DP-700 New Exam Camp & New DP-700 Test Pattern

It is known to us that our DP-700 study materials have been keeping a high pass rate all the time. There is no doubt that it must be due to the high quality of our study materials. It is a matter of common sense that pass rate is the most important standard to testify the DP-700 Study Materials. The high pass rate of our study materials means that our products are very effective and useful for all people to pass their exam and get the related certification.

Microsoft Implementing Data Engineering Solutions Using Microsoft Fabric Sample Questions (Q93-Q98):

NEW QUESTION # 93

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. Create views.
- D. Increase the capacity.C 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 # 94

HOTSPOT

You have a Fabric workspace.

You are debugging a statement and discover the following issues:

Sometimes, the statement fails to return all the expected rows.

The PurchaseDate output column is NOT in the expected format of mmm dd, yy.

You need to resolve the issues. The solution must ensure that the data types of the results are retained. The results can contain blank cells.

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

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

□ Explanation:

□ Topic 1, Litware, IncCase 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

Litware, Inc. is a publishing company that has an online bookstore and several retail bookstores worldwide.

Litware also manages an online advertising business for the authors it represents.

Existing Environment. Fabric Environment

Litware has a Fabric workspace named Workspace1. High concurrency is enabled for Workspace1.

The company has a data engineering team that uses Python for data processing.

Existing Environment. Data Processing

The retail bookstores send sales data at the end of each business day, while the online bookstore constantly provides logs and sales data to a central enterprise resource planning (ERP) system.

Litware implements a medallion architecture by using the following three layers: bronze, silver, and gold. The sales data is ingested from the ERP system as Parquet files that land in the Files folder in a lakehouse.

Notebooks are used to transform the files in a Delta table for the bronze and silver layers. The gold layer is in a warehouse that has V-Order disabled.

Litware has image files of book covers in Azure Blob Storage. The files are loaded into the Files folder.

Existing Environment. Sales Data

Month-end sales data is processed on the first calendar day of each month. Data that is older than one month never changes.

In the source system, the sales data refreshes every six hours starting at midnight each day.

The sales data is captured in a Dataflow Gen1 dataflow. When the dataflow runs, new and historical data is captured. The dataflow captures the following fields of the source:

Sales Date

Author

Price

Units

SKU

A table named AuthorSales stores the sales data that relates to each author. The table contains a column named AuthorEmail.

Authors authenticate to a guest Fabric tenant by using their email address.

Existing Environment. Security Groups

Litware has the following security groups:

Sales

Fabric Admins

Streaming Admins

Existing Environment. Performance Issues

Business users perform ad-hoc queries against the warehouse. The business users indicate that reports against the warehouse sometimes run for two hours and fail to load as expected. Upon further investigation, the data engineering team receives the following error message when the reports fail to load: "The SQL query failed while running." The data engineering team wants to debug the issue and find queries that cause more than one failure.

When the authors have new book releases, there is often an increase in sales activity. This increase slows the data ingestion process. The company's sales team reports that during the last month, the sales data has NOT been up-to-date when they arrive at work in the morning.

Requirements. Planned Changes

Litware recently signed a contract to receive book reviews. The provider of the reviews exposes the data in Amazon Simple Storage Service (Amazon S3) buckets.

Litware plans to manage Search Engine Optimization (SEO) for the authors. The SEO data will be streamed from a REST API.

Requirements. Version Control

Litware plans to implement a version control solution in Fabric that will use GitHub integration and follow the principle of least privilege.

Requirements. Governance Requirements

To control data platform costs, the data platform must use only Fabric services and items. Additional Azure resources must NOT be provisioned.

Requirements. Data Requirements

Litware identifies the following data requirements:

Process the SEO data in near-real-time (NRT).

Make the book reviews available in the lakehouse without making a copy of the data.

When a new book cover image arrives in the Files folder, process the image as soon as possible.

NEW QUESTION # 95

You have a Fabric workspace that contains a semantic model named Modell. You need to monitor the refresh history of Modell and visualize the refresh history in a chart. What should you use?

- A. a notebook
- B. a data pipeline
- C. a Dataflow Gen2 dataflow
- D. the refresh history from the settings of Modell.

Answer: D

NEW QUESTION # 96

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:

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:

Topic 1, Litware, IncCase 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

Litware, Inc. is a publishing company that has an online bookstore and several retail bookstores worldwide.

Litware also manages an online advertising business for the authors it represents.

Existing Environment. Fabric Environment

Litware has a Fabric workspace named Workspace1. High concurrency is enabled for Workspace1.

The company has a data engineering team that uses Python for data processing.

Existing Environment. Data Processing

The retail bookstores send sales data at the end of each business day, while the online bookstore constantly provides logs and sales data to a central enterprise resource planning (ERP) system.

Litware implements a medallion architecture by using the following three layers: bronze, silver, and gold. The sales data is ingested from the ERP system as Parquet files that land in the Files folder in a lakehouse.

Notebooks are used to transform the files in a Delta table for the bronze and silver layers. The gold layer is in a warehouse that has V-Order disabled.

Litware has image files of book covers in Azure Blob Storage. The files are loaded into the Files folder.

Existing Environment. Sales Data

Month-end sales data is processed on the first calendar day of each month. Data that is older than one month never changes.

In the source system, the sales data refreshes every six hours starting at midnight each day.

The sales data is captured in a Dataflow Gen1 dataflow. When the dataflow runs, new and historical data is captured. The dataflow captures the following fields of the source:

Sales Date

Author

Price

Units

SKU

A table named AuthorSales stores the sales data that relates to each author. The table contains a column named AuthorEmail.

Authors authenticate to a guest Fabric tenant by using their email address.

Existing Environment. Security Groups

Litware has the following security groups:

Sales

Fabric Admins

Streaming Admins

Existing Environment. Performance Issues

Business users perform ad-hoc queries against the warehouse. The business users indicate that reports against the warehouse sometimes run for two hours and fail to load as expected. Upon further investigation, the data engineering team receives the following error message when the reports fail to load: "The SQL query failed while running." The data engineering team wants to debug the issue and find queries that cause more than one failure.

When the authors have new book releases, there is often an increase in sales activity. This increase slows the data ingestion process.

The company's sales team reports that during the last month, the sales data has NOT been up-to-date when they arrive at work in the morning.

Requirements. Planned Changes

Litware recently signed a contract to receive book reviews. The provider of the reviews exposes the data in Amazon Simple Storage Service (Amazon S3) buckets.

Litware plans to manage Search Engine Optimization (SEO) for the authors. The SEO data will be streamed from a REST API.

Requirements. Version Control

Litware plans to implement a version control solution in Fabric that will use GitHub integration and follow the principle of least privilege.

Requirements. Governance Requirements

To control data platform costs, the data platform must use only Fabric services and items. Additional Azure resources must NOT be provisioned.

Requirements. Data Requirements

Litware identifies the following data requirements:

Process the SEO data in near-real-time (NRT).

Make the book reviews available in the lakehouse without making a copy of the data.

When a new book cover image arrives in the Files folder, process the image as soon as possible.

NEW QUESTION # 97

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 KQL database that contains two tables named Stream and Reference. Stream contains streaming data in the following format.

Reference contains reference data in the following format.

Both tables contain millions of rows.

You have the following KQL queryset.

You need to reduce how long it takes to run the KQL queryset.

Solution: You add the `make_list()` function to the output columns.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

