Well DP-700 Prep & Clear DP-700 Exam

2025 Latest Prep4sureExam DP-700 PDF Dumps and DP-700 Exam Engine Free Share: https://drive.google.com/open?id=1fHGruXQfvLye8BUmQB0Hv1unpiIJr_mu

Don't be tied up in small things. Don't let your exam affect your regular work. Professionals do professionals. Only spend a little money on Microsoft DP-700 exam braindumps pdf, you will pass exam easily with only 24-36 hours preparation before the real test. Work is important, relax properly is important, Let our DP-700 Exam Braindumps pdf help you clear your exam easily so that you can achieve three things at one stroke. In fact time is money.

Microsoft DP-700 Exam Syllabus Topics:

Topic	Details
Topic 1	• 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	• 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	• 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.

>> Well DP-700 Prep <<

DP-700 new questions & DP-700 dumps VCE & DP-700 dump collection

There are totally three versions of DP-700 practice materials which are the most suitable versions for you: PDF, software and app versions. We promise ourselves and exam candidates to make these DP-700 preparation prep top notch. So if you are in a dark space, our DP-700 Study Guide can inspire you make great improvements. With the high pass rate of our DP-700 learing engine as 98% to 100%, you can be confident and ready to pass the exam easily.

Microsoft Implementing Data Engineering Solutions Using Microsoft Fabric Sample Questions (Q103-Q108):

NEW QUESTION # 103
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.

Column name	Data type
Timestamp	Datetime
GeoLocation	Dynamic
Temperature	Decimal
DeviceId	Int

Reference contains reference data in the following format.

Column name	Data type
DeviceId	Int
DeviceName	String

Both tables contain millions of rows.
You have the following KQL queryset.

```
01 Stream
02 | extend lat = todecimal(GeoLocation.Latitude), long = todecimal(GeoLocation.Longitude)
03 | join kind=inner Reference on DeviceId
04 | project Timestamp, lat, long, Temperature, DeviceName
05 | filter Temperature >= 10
06 | render scatterchart with (kind = map)
```

You need to reduce how long it takes to run the KQL queryset.
Solution: You change project to extend.
Does this meet the goal?

- A. No
- B. Yes

Answer: A

Explanation:
Using extend retains all columns in the table, potentially increasing the size of the output unnecessarily.
project is more efficient because it selects only the required columns.

NEW QUESTION # 104
You have a Fabric workspace that contains a lakehouse named Lakehouse1. Data is ingested into Lakehouse1 as one flat table. The table contains the following columns.

Name	Description
TransactionID	Contains a unique ID for each transaction
Date	Contains the date of a transaction
ProductID	Contains a unique ID for each product
ProductColor	Contains a descriptive attribute that describes the color of each product
ProductName	Contains a unique name for each product
SalesAmount	Contains the sales amount of a transaction

You plan to load the data into a dimensional model and implement a star schema. From the original flat table, you create two tables named FactSales and DimProduct. You will track changes in DimProduct.
You need to prepare the data.
Which three columns should you include in the DimProduct table? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

- A. ProductID
- B. ProductName

- C. TransactionID
- D. SalesAmount
- E. ProductColor
- F. Date

Answer: A,B,E

Explanation:
In a star schema, the DimProduct table serves as a dimension table that contains descriptive attributes about products. It will provide context for the FactSales table, which contains transactional data. The following columns should be included in the DimProduct table:
* ProductName: The ProductName is an important descriptive attribute of the product, which is needed for analysis and reporting in a dimensional model.
* ProductColor: ProductColor is another descriptive attribute of the product. In a star schema, it makes sense to include attributes like color in the dimension table to help categorize products in the analysis.
* ProductID: ProductID is the primary key for the DimProduct table, which will be used to join the FactSales table to the product dimension. It's essential for uniquely identifying each product in the model.

NEW QUESTION # 105
You are building a data loading pattern by using a Fabric data pipeline. The source is an Azure SQL database that contains 25 tables. The destination is a lakehouse.
In a warehouse, you create a control table named Control.Object as shown in the exhibit. (Click the Exhibit tab.) You need to build a data pipeline that will support the dynamic ingestion of the tables listed in the control table by using a single execution.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

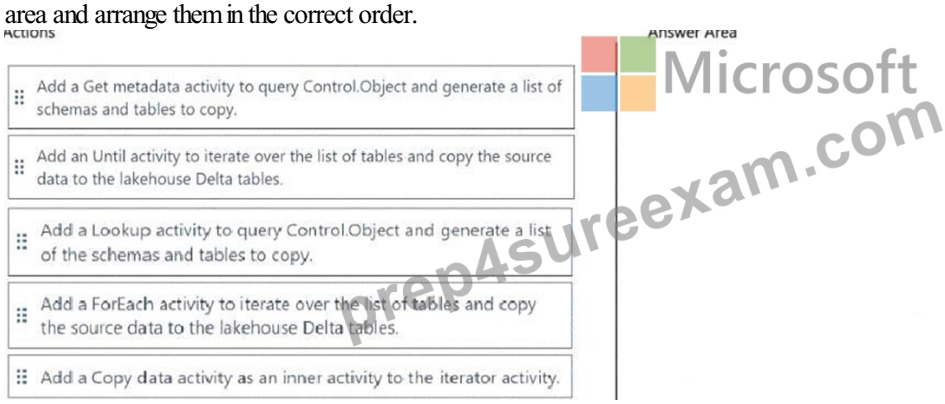

Answer:

Explanation:

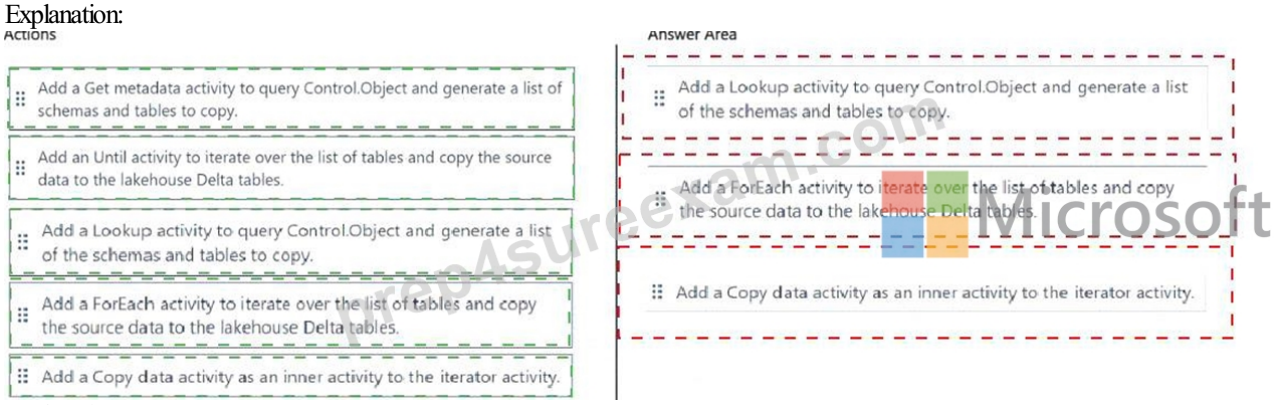

Explanation:

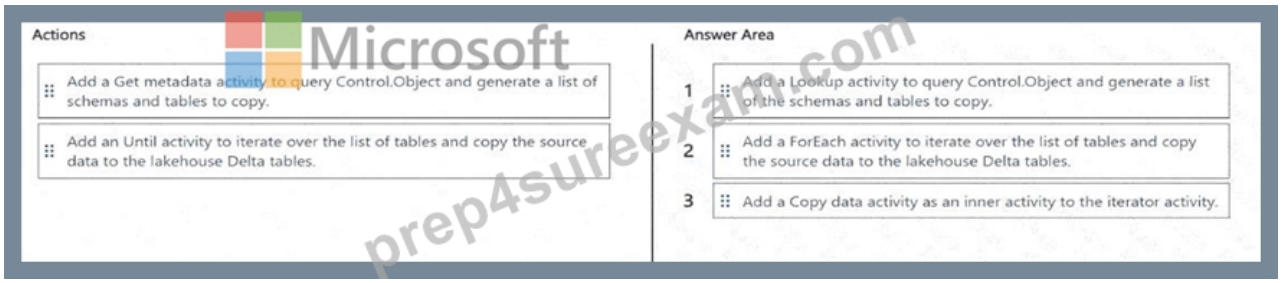

NEW QUESTION # 106
You are building a data orchestration pattern by using a Fabric data pipeline named Dynamic Data Copy as shown in the exhibit. (Click the Exhibit tab.)

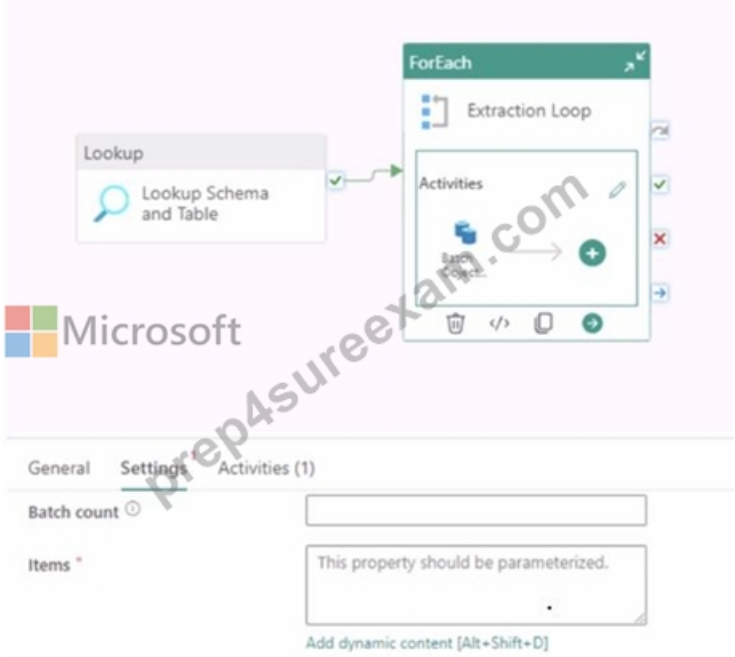

Dynamic Data Copy does NOT use parametrization.
You need to configure the ForEach activity to receive the list of tables to be copied.
How should you complete the pipeline expression? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

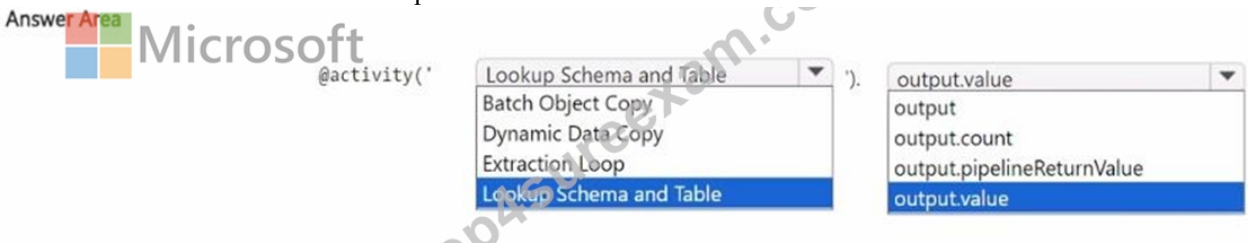

Answer:

Explanation:

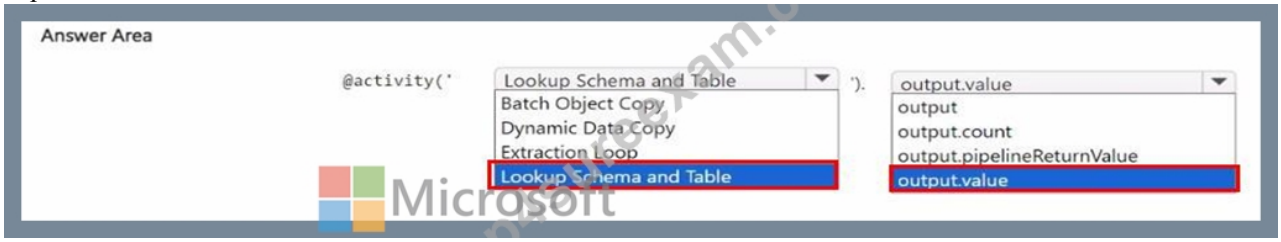

NEW QUESTION # 107

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

Answer: A

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.
Topic 1, Litware, Inc
Overview
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:
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:
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:

NEW QUESTION # 108
.....