

Microsoft

DP-700 Exam

Implementing Data Engineering Solutions Using Microsoft Fabric

**Questions & Answers
Demo**

Version: 4.0

Question: 1

You need to ensure that the data analysts can access the gold layer lakehouse.
What should you do?

- A. Add the DataAnalyst group to the Viewer role for WorkspaceA.
- B. Share the lakehouse with the DataAnalysts group and grant the Build reports on the default semantic model permission.
- C. Share the lakehouse with the DataAnalysts group and grant the Read all SQL Endpoint data permission.
- D. Share the lakehouse with the DataAnalysts group and grant the Read all Apache Spark permission.

Answer: C

Explanation:

Data Analysts' Access Requirements must only have read access to the Delta tables in the gold layer and not have access to the bronze and silver layers.

The gold layer data is typically queried via SQL Endpoints. Granting the Read all SQL Endpoint data permission allows data analysts to query the data using familiar SQL-based tools while restricting access to the underlying files.

Question: 2

HOTSPOT

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

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

Bronze Layer: A pipeline Copy activity
 The bronze layer is used to store raw, unprocessed data

a. 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.

Question: 3

You need to ensure that usage of the data in the Amazon S3 bucket meets the technical

requirements.

What should you do?

- A. Create a workspace identity and enable high concurrency for the notebooks.
- B. Create a shortcut and ensure that caching is disabled for the workspace.
- C. Create a workspace identity and use the identity in a data pipeline.
- D. Create a shortcut and ensure that caching is enabled for the workspace.

Answer: B

Explanation:

To ensure that the usage of the data in the Amazon S3 bucket meets the technical requirements, we must address two key points:

- Minimize egress costs associated with cross-cloud data access: Using a shortcut ensures that Fabric does not replicate the data from the S3 bucket into the lakehouse but rather provides direct access to the data in its original location. This minimizes cross-cloud data transfer and avoids additional egress costs.
- Prevent saving a copy of the raw data in the lakehouses: Disabling caching ensures that the raw data is not copied or persisted in the Fabric workspace. The data is accessed on-demand directly from the Amazon S3 bucket.

Question: 4

HOTSPOT

You need to create the product dimension.

How should you complete the Apache Spark SQL code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
SELECT ProductID, ProductNumber, ProductName, ModelName, SubCategoryName, CategoryName
FROM ContosoLake.Products p
```

ContosoLake.ProductSubCategories s ON p.SubCategoryID = s.SubCategoryID
 FULL JOIN
 INNER JOIN
 LEFT ANTI JOIN
 LEFT OUTER JOIN
 OUTER JOIN

ContosoLake.ProductCategories c ON c.CategoryID = s.CategoryID
 FULL JOIN
 INNER JOIN
 LEFT ANTI JOIN
 LEFT OUTER JOIN
 OUTER JOIN

WHERE
 CategoryID = 1;
 CategoryName is not null;
 IsActive = 1;
 IsActive is not null;
 ProductNumber is not null;
 SubCategoryID = 1;
 SubCategoryName is not null;

Answer:

Explanation:

```
SELECT ProductID, ProductNumber, ProductName, ModelName, SubCategoryName, CategoryName
FROM ContosoLake.Products p
```

ContosoLake.ProductSubCategories s ON p.SubCategoryID = s.SubCategoryID
 FULL JOIN
 INNER JOIN
 LEFT ANTI JOIN
 LEFT OUTER JOIN
 OUTER JOIN

ContosoLake.ProductCategories c ON c.CategoryID = s.CategoryID
 FULL JOIN
 INNER JOIN
 LEFT ANTI JOIN
 LEFT OUTER JOIN
 OUTER JOIN

WHERE
 CategoryID = 1;
 CategoryName is not null;
 IsActive = 1;
 IsActive is not null;
 ProductNumber is not null;
 SubCategoryID = 1;
 SubCategoryName is not null;

Join between Products and ProductSubCategories:

- Use an INNER JOIN.

■The goal is to include only products that are assigned to a subcategory. An INNER JOIN ensures that only matching records (i.e., products with a valid subcategory) are included.

Join between ProductSubCategories and ProductCategories:

■Use an INNER JOIN.

■Similar to the above logic, we want to include only subcategories assigned to a valid product category. An INNER JOIN ensures this condition is met.

WHERE Clause

Condition: IsActive = 1

Only active products (where IsActive equals 1) should be included in the gold layer. This filters out inactive products.

Question: 5

You need to populate the MAR1 data in the bronze layer.

Which two types of activities should you include in the pipeline? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. ForEach
- B. Copy data
- C. WebHook
- D. Stored procedure

Answer: AB

Explanation:

MAR1 has seven entities, each accessible via a different API endpoint. A ForEach activity is required to iterate over these endpoints to fetch data from each one. It enables dynamic execution of API calls for each entity.

The Copy data activity is the primary mechanism to extract data from REST APIs and load it into the bronze layer in Delta format. It supports native connectors for REST APIs and Delta, minimizing development effort.