



Vendor: Microsoft

Exam Code: DP-600

Exam Name: Implementing Analytics Solutions Using
Microsoft Fabric

Version: DEMO

QUESTION 1

Case Study 1 - Contoso

Overview

Contoso, Ltd. is a US-based health supplements company. Contoso has two divisions named Sales and Research. The Sales division contains two departments named Online Sales and Retail Sales. The Research division assigns internally developed product lines to individual teams of researchers and analysts.

Existing Environment

Identity Environment

Contoso has a Microsoft Entra tenant named contoso.com. The tenant contains two groups named ResearchReviewersGroup1 and ResearchReviewersGroup2.

Data Environment

Contoso has the following data environment:

- The Sales division uses a Microsoft Power BI Premium capacity.
- The semantic model of the Online Sales department includes a fact table named Orders that uses Import mode. In the system of origin, the OrderID value represents the sequence in which orders are created.
- The Research department uses an on-premises, third-party data warehousing product.
- Fabric is enabled for contoso.com.
- An Azure Data Lake Storage Gen2 storage account named storage1 contains Research division data for a product line named Productline1. - The data is in the delta format.
- A Data Lake Storage Gen2 storage account named storage2 contains Research division data for a product line named Productline2. The data is in the CSV format.

Requirements

Planned Changes

Contoso plans to make the following changes:

- Enable support for Fabric in the Power BI Premium capacity used by the Sales division.
- Make all the data for the Sales division and the Research division available in Fabric.
- For the Research division, create two Fabric workspaces named Productline1ws and Productline2ws.
- In Productline1ws, create a lakehouse named Lakehouse1.
- In Lakehouse1, create a shortcut to storage1 named ResearchProduct.

You need to ensure that Contoso can use version control to meet the data analytics requirements and the general requirements.

What should you do?

- A. Store all the semantic models and reports in Data Lake Gen2 storage.
- B. Modify the settings of the Research workspaces to use a GitHub repository.
- C. Modify the settings of the Research division workspaces to use an Azure Repos repository.
- D. Store all the semantic models and reports in Microsoft OneDrive.

Answer: C

Explanation:

Currently, only Git in Azure Repos is supported.

<https://learn.microsoft.com/en-us/fabric/cicd/git-integration/intro-to-git-integration#considerations-and-limitations>

QUESTION 2

Case Study 2 - Litware, Inc

Overview

Litware, Inc. is a manufacturing company that has offices throughout North America. The analytics team at Litware contains data engineers, analytics engineers, data analysts, and data scientists.

Existing Environment

Fabric Environment

Litware has been using a Microsoft Power BI tenant for three years. Litware has NOT enabled any Fabric capacities and features.

Available Data

Litware has data that must be analyzed as shown in the following table.

Description	Original source	Total size
Customer data	Customer relationship management (CRM) system	50 MB
Product data	Customer relationship management (CRM) system	200 MB
Customer satisfaction surveys	SurveyMonkey	500 GB

The Product data contains a single table and the following columns.

Name	Data type
ProductID	Integer
ProductName	String
ProductCategory	String
ListPrice	Decimal

The customer satisfaction data contains the following tables:

- Survey
- Question
- Response

For each survey submitted, the following occurs:

- One row is added to the Survey table.
- One row is added to the Response table for each question in the survey.
- The Question table contains the text of each survey question. The third question in each survey response is an overall satisfaction score. Customers can submit a survey after each purchase.

What should you recommend using to ingest the customer data into the data store in the AnalyticsPOC workspace?

- A. a stored procedure
- B. a pipeline that contains a KQL activity
- C. a Spark notebook
- D. a dataflow

Answer: D

Explanation:

Even though the text reads "Data will be loaded without transformation in one area of the AnalyticsPOC data store": in general, dataflows are used when data transformations are involved after ingestion. As suggested by user BHARAT, the Copy Activity should be the optimal solution.

QUESTION 3

You have a Fabric tenant that contains a warehouse.

You use a dataflow to load a new dataset from OneLake to the warehouse.

You need to add a PowerQuery step to identify the maximum values for the numeric columns.

Which function should you include in the step?

- A. Table.MaxN
- B. Table.Max
- C. Table.Range
- D. Table.Profile

Answer: D

Explanation:

<https://learn.microsoft.com/en-us/powerquery-m/table-profile>

QUESTION 4

You have a Fabric tenant that contains a machine learning model registered in a Fabric workspace.

You need to use the model to generate predictions by using the PREDICT function in a Fabric notebook.

Which two languages can you use to perform model scoring? Each correct answer presents a complete solution.

NOTE: Each correct answer is worth one point.

- A. T-SQL
- B. DAX
- C. Spark SQL
- D. PySpark

Answer: CD

Explanation:

<https://learn.microsoft.com/en-us/azure/synapse-analytics/machine-learning/tutorial-score-model-predict-spark-pool>

QUESTION 5

You have a Fabric tenant that contains a Microsoft Power BI report named Report1. Report1 includes a Python visual.

Data displayed by the visual is grouped automatically and duplicate rows are NOT displayed.

You need all rows to appear in the visual.

What should you do?

- A. Reference the columns in the Python code by index.
- B. Modify the Sort Column By property for all columns.
- C. Add a unique field to each row.
- D. Modify the Summarize By property for all columns.

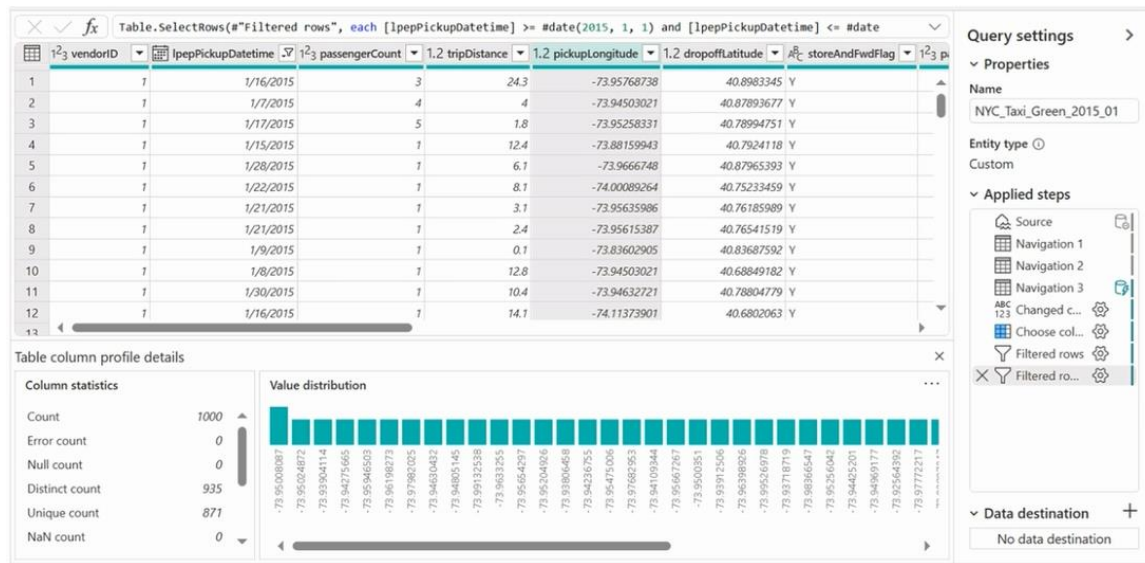
Answer: D

Explanation:

By setting the "Summarize By" property to "None" for all columns, you disable automatic aggregation and ensure all rows, including duplicates, are displayed in the Python visual.

QUESTION 6

You have a Fabric workspace named Workspace1 that contains a dataflow named Dataflow1. Dataflow1 has a query that returns 2,000 rows. You view the query in Power Query as shown in the following exhibit.



What can you identify about the pickupLongitude column?

- A. The column has duplicate values.
- B. All the table rows are profiled.
- C. The column has missing values.
- D. There are 935 values that occur only once.

Answer: A

QUESTION 7

You have a Fabric tenant named Tenant1 that contains a workspace named WS1. WS1 uses a capacity named C1 and contains a dataset named DS1. You need to ensure read-write access to DS1 is available by using XMLA endpoint. What should be modified first?

- A. the DS1 settings
- B. the WS1 settings
- C. the C1 settings
- D. the Tenant1 settings

Answer: C

Explanation:

<https://learn.microsoft.com/en-us/power-bi/enterprise/service-premium-connect-tools>

QUESTION 8

You are creating a semantic model in Microsoft Power BI Desktop.
You plan to make bulk changes to the model by using the Tabular Model Definition Language (TMDL) extension for Microsoft Visual Studio Code.
You need to save the semantic model to a file.
Which file format should you use?

- A. PBIP
- B. PBIX
- C. PBIT
- D. PBIDS

Answer: A

Explanation:

The PBIP will create one file and two folders, PBIP.Dataset contains definition folder that is use to host the .tmdl files.

QUESTION 9

Hotspot Question

You have a Fabric tenant.

You plan to create a Fabric notebook that will use Spark DataFrames to generate Microsoft Power BI visuals.

You run the following code.

```
from powerbiclient import QuickVisualize, get_dataset_config, Report

PBI_visualize = QuickVisualize(get_dataset_config(df))
PBI_visualize
```

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
The code embeds an existing Power BI report.	<input type="radio"/>	<input type="radio"/>
The code creates a Power BI report.	<input type="radio"/>	<input type="radio"/>
The code displays a summary of the DataFrame.	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statements	Yes	No
The code embeds an existing Power BI report.	<input type="radio"/>	<input checked="" type="radio"/>
The code creates a Power BI report.	<input checked="" type="radio"/>	<input type="radio"/>
The code displays a summary of the DataFrame.	<input type="radio"/>	<input checked="" type="radio"/>

QUESTION 10

Drag and Drop Question

You have a Fabric tenant that contains a semantic model. The model contains data about retail stores.

You need to write a DAX query that will be executed by using the XMLA endpoint. The query must return a table of stores that have opened since December 1, 2023.

How should you complete the DAX expression? To answer, drag the appropriate values to 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.

Values	Answer Area
DEFINE	<input type="text"/>
EVALUATE	VAR_SalesSince = DATE (2023, 12, 01)
FILTER	<input type="text"/>
SUMMARIZE	FILTER (<input type="text"/> (Store, Store[Name], Store[OpenDate]), Store[OpenDate] >= _SalesSince)
TABLE	

Answer:

Values	Answer Area
	<div>DEFINE</div> <div>VAR_SalesSince = DATE (2023, 12, 01)</div>
<div>FILTER</div>	<div>EVALUATE</div> <div>FILTER (</div>
<div>TABLE</div>	<div>SUMMARIZE (Store, Store[Name], Store[OpenDate]), Store[OpenDate] >= _SalesSince)</div>

QUESTION 11

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 tenant that contains a lakehouse named Lakehouse1. Lakehouse1 contains a Delta table named Customer.

When you query Customer, you discover that the query is slow to execute. You suspect that maintenance was NOT performed on the table.

You need to identify whether maintenance tasks were performed on Customer.

Solution: You run the following Spark SQL statement: REFRESH TABLE customer

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Thank You for Trying Our Product

Braindump2go Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.braindump2go.com/all-products.html>



Microsoft



ORACLE



CITRIX



JUNIPER
NETWORKS



EMC²
where information lives

10% Discount Coupon Code: ASTR14