Microsoft

Exam Questions DP-600

Implementing Analytics Solutions Using Microsoft Fabric

NEW QUESTION 1
- (Topic 1)
You need to recommend a solution to prepare the tenant for the PoC.
Which two actions should you recommend performing from the Fabric Admin portal? Each correct answer presents part of the solution.
NOTE: Each correct answer is worth one point.

A. Enable the Users can try Microsoft Fabric paid features option for specific security groups.
B. Enable the Allow Azure Active Directory guest users to access Microsoft Fabric option for specific security groups.
C. Enable the Users can create Fabric items option and exclude specific security groups.
D. Enable the Users can try Microsoft Fabric paid features option for the entire organization.
E. Enable the Users can create Fabric items option for specific security groups.

Answer: AE

Explanation:
The PoC is planned to be completed using a Fabric trial capacity, which implies that users involved in the PoC should be able to try paid features. However, this should be limited to specific security groups involved in the PoC to prevent the entire organization from accessing these features before the trial is proven successful (A). The ability for users to create Fabric items should also be enabled for specific security groups to ensure that only the relevant team members participating in the PoC can create items in the Fabric environment (E).

NEW QUESTION 2
- (Topic 1)
Which type of data store should you recommend in the AnalyticsPOC workspace?

A. a data lake
B. a warehouse
C. a lakehouse
D. an external Hive metaStore

Answer: C

Explanation:
A lakehouse (C) should be recommended for the AnalyticsPOC workspace. It combines the capabilities of a data warehouse with the flexibility of a data lake. A lakehouse supports semi-structured and unstructured data and allows for T-SQL and Python read access, fulfilling the technical requirements outlined for Litware. References = For further understanding, Microsoft's documentation on the lakehouse architecture provides insights into how it supports various data types and analytical operations.

NEW QUESTION 3
- (Topic 1)
What should you recommend using to ingest the customer data into the data store in the AnatyticsPOC workspace?

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

Answer: D

Explanation:
For ingesting customer data into the data store in the AnalyticsPOC workspace, a dataflow (D) should be recommended. Dataflows are designed within the Power BI service to ingest, cleanse, transform, and load data into the Power BI environment. They allow for the low-code ingestion and transformation of data as needed by Litware's technical requirements. References = You can learn more about dataflows and their use in Power BI environments in Microsoft's Power BI documentation.

NEW QUESTION 4
HOTSPOT - (Topic 2)
You have a Fabric tenant that contains a warehouse named Warehouse1. Warehouse1 contains three schemas named schemaA, schemaB. and schemaC
You need to ensure that a user named User1 can truncate tables in schemaA only.
How should you complete the T-SQL statement? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
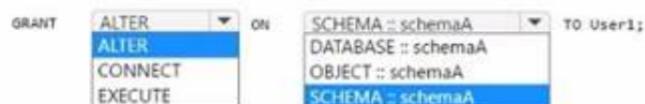
A. Mastered
B. Not Mastered

Answer: A

Explanation:
? GRANT ALTER ON SCHEMA::schemaA TO User1;
The ALTER permission allows a user to modify the schema of an object, and granting ALTER on a schema will allow the user to perform operations like

TRUNCATE TABLE on any object within that schema. It is the correct permission to grant to User1 for truncating tables in schemaA.
References =
? GRANT Schema Permissions
? Permissions That Can Be Granted on a Schema

NEW QUESTION 5
- (Topic 2)
You have a Fabric tenant that contains 30 CSV files in OneLake. The files are updated daily.
You create a Microsoft Power BI semantic model named Modell that uses the CSV files as a data source. You configure incremental refresh for Model 1 and publish the model to a Premium capacity in the Fabric tenant.
When you initiate a refresh of Model1, the refresh fails after running out of resources. What is a possible cause of the failure?

A. Query folding is occurring.
B. Only refresh complete days is selected.
C. XMLA Endpoint is set to Read Only.
D. Query folding is NOT occurring.
E. The data type of the column used to partition the data has changed.

Answer: E

Explanation:
 A possible cause for the failure is that query folding is NOT occurring (D). Query folding helps optimize refresh by pushing down the query logic to the source system, reducing the amount of data processed and transferred, hence conserving resources. References = The Power BI documentation on incremental refresh and query folding provides detailed information on this topic.

NEW QUESTION 6
- (Topic 2)
You have a Fabric tenant that uses a Microsoft tower BI Premium capacity. You need to enable scale-out for a semantic model. What should you do first?

A. At the semantic model level, set Large dataset storage format to Off.
B. At the tenant level, set Create and use Metrics to Enabled.
C. At the semantic model level, set Large dataset storage format to On.
D. At the tenant level, set Data Activator to Enabled.

Answer: C

Explanation:
 To enable scale-out for a semantic model, you should first set Large dataset storage format to On (C) at the semantic model level. This configuration is necessary to handle larger datasets effectively in a scaled-out environment. References = Guidance on configuring large dataset storage formats for scale-out is available in the Power BI documentation.

NEW QUESTION 7
- (Topic 2)
You have a Fabric tenant that contains a lakehouse. You plan to use a visual query to merge two tables.
You need to ensure that the query returns all the rows that are present in both tables. Which type of join should you use?

A. left outer
B. right anti
C. full outer
D. left anti
E. right outer
F. inner

Answer: C

Explanation:
 When you need to return all rows that are present in both tables, you use a full outer join. This type of join combines the results of both left and right outer joins and returns all rows from both tables, with matching rows from both sides where available. If there is no match, the result is NULL on the side of the join where there is no match. References: Information about joins and their use in querying data in a lakehouse can be typically found in the SQL and data processing documentation of the Fabric tenant or lakehouse solutions.

NEW QUESTION 8
- (Topic 2)
You have a Fabric tenant that contains a new semantic model in OneLake. You use a Fabric notebook to read the data into a Spark DataFrame.
You need to evaluate the data to calculate the min, max, mean, and standard deviation values for all the string and numeric columns.
Solution: You use the following PySpark expression: df.explain()
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:
 The df.explain() method does not meet the goal of evaluating data to calculate statistical functions. It is used to display the physical plan that Spark will execute.
References = The correct usage of the explain() function can be found in the PySpark documentation.

NEW QUESTION 9

- (Topic 2)
You have a semantic model named Model 1. Model 1 contains five tables that all use Import mode. Model1 contains a dynamic row-level security (RLS) role named HR. The HR role filters employee data so that HR managers only see the data of the department to which they are assigned.
You publish Model1 to a Fabric tenant and configure RLS role membership. You share the model and related reports to users.
An HR manager reports that the data they see in a report is incomplete. What should you do to validate the data seen by the HR Manager?

A. Ask the HR manager to open the report in Microsoft Power BI Desktop.
B. Select Test as role to view the data as the HR role.
C. Select Test as role to view the report as the HR manager,
D. Filter the data in the report to match the intended logic of the filter for the HR department.

Answer: B

Explanation:
To validate the data seen by the HR manager, you should use the 'Test as role' feature in Power BI service. This allows you to see the data exactly as it would appear for the HR role, considering the dynamic RLS setup. Here is how you would proceed:
? Navigate to the Power BI service and locate Model1.
? Access the dataset settings for Model1.
? Find the security/RLS settings where you configured the roles.
? Use the 'Test as role' feature to simulate the report viewing experience as the HR role.
? Review the data and the filters applied to ensure that the RLS is functioning correctly.
? If discrepancies are found, adjust the RLS expressions or the role membership as needed.
References: The 'Test as role' feature and its use for validating RLS in Power BI is covered in the Power BI documentation available on Microsoft's official documentation.

NEW QUESTION 10
DRAG DROP - (Topic 2)
You have a Fabric tenant that contains a lakehouse named Lakehouse1
Readings from 100 loT devices are appended to a Delta table in Lakehouse1. Each set of readings is approximately 25 KB. Approximately 10 GB of data is received daily.
All the table and SparkSession settings are set to the default.
You discover that queries are slow to execute. In addition, the lakehouse storage contains data and log files that are no longer used.
You need to remove the files that are no longer used and combine small files into larger files with a target size of 1 GB per file.
What should you do? To answer, drag the appropriate actions to the correct requirements. Each action 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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
? Remove the files: Run the VACUUM command on a schedule.
? Combine the files: Set the optimizeWrite table setting. or Run the OPTIMIZE command on a schedule.
To remove files that are no longer used, the VACUUM command is used in Delta Lake to clean up invalid files from a table. To combine smaller files into larger ones, you can either set the optimizeWrite setting to combine files during write operations or use the OPTIMIZE command, which is a Delta Lake operation used to compact small files into larger ones.

NEW QUESTION 10
- (Topic 2)
You have a Fabric tenant that contains a lakehouse named lakehouse1. Lakehouse1 contains an unpartitioned table named Table1.
You plan to copy data to Table1 and partition the table based on a date column in the source data.
You create a Copy activity to copy the data to Table1.
You need to specify the partition column in the Destination settings of the Copy activity. What should you do first?

A. From the Destination tab, set Mode to Append.
B. From the Destination tab, select the partition column,
C. From the Source tab, select Enable partition discovery
D. From the Destination tab, set Mode to Overwrite.

Answer: B

Explanation:
Before specifying the partition column in the Destination settings of the Copy activity, you should set Mode to Append (A). This will allow the Copy activity to add data to the table while taking the partition column into account. References = The configuration options for Copy activities and partitioning in Azure Data Factory, which are applicable to Fabric dataflows, are outlined in the official Azure Data Factory documentation.

NEW QUESTION 15
- (Topic 2)

You have a Fabric tenant that contains a lakehouse named Lakehouse1. Lakehouse1 contains a subfolder named Subfolder1 that contains CSV files. You need to convert the CSV files into the delta format that has V-Order optimization enabled. What should you do from Lakehouse explorer?

A. Use the Load to Tables feature.
B. Create a new shortcut in the Files section.
C. Create a new shortcut in the Tables section.
D. Use the Optimize feature.

Answer: D

Explanation:
To convert CSV files into the delta format with Z-Order optimization enabled, you should use the Optimize feature (D) from Lakehouse Explorer. This will allow you to optimize the file organization for the most efficient querying. References = The process for converting and optimizing file formats within a lakehouse is discussed in the lakehouse management documentation.

NEW QUESTION 18
HOTSPOT - (Topic 2)
You have a Fabric warehouse that contains a table named Sales.Orders. Sales.Orders contains the following columns.

Name	Data type	Nullable
OrderID	Integer	No
CustomerID	Integer	No
OrderDate	Date	No
Quantity	Integer	Yes
Weight	Decimal(18, 3)	Yes
ListPrice	Decimal(18, 2)	No
SalePrice	Decimal(18, 2)	Yes

You need to write a T-SQL query that will return the following columns.

Name	Description
OrderID	Returns OrderID
CustomerID	Returns CustomerID
PeriodDate	Returns a date representing the first day of the month for OrderDate
DayName	Returns the name of the day for OrderDate, such as Wednesday

How should you complete the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

```
SELECT OrderID, CustomerID,

    DATEFROMPARTS          ▼
    DATE_BUCKET
    DATEFROMPARTS
  FR DATEPART
    DATETRUNC

    DATENAME( weekday      ▼  , OrderDate) AS DayName
                day
    FROM Sales.Ord dayofyear
                weekday
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
For the PeriodDate that returns the first day of the month for OrderDate, you should use DATEFROMPARTS as it allows you to construct a date from its individual components (year, month, day).
For the DayName that returns the name of the day for OrderDate, you should use
DATENAME with the weekday date part to get the full name of the weekday. The complete SQL query should look like this:
SELECT OrderID, CustomerID,
DATEFROMPARTS(YEAR(OrderDate), MONTH(OrderDate), 1) AS PeriodDate, DATENAME(weekday, OrderDate) AS DayName
FROM Sales.Orders
Select DATEFROMPARTS for the PeriodDate and weekday for the DayName in the answer area.

NEW QUESTION 20
HOTSPOT - (Topic 2)
You have the source data model shown in the following exhibit.
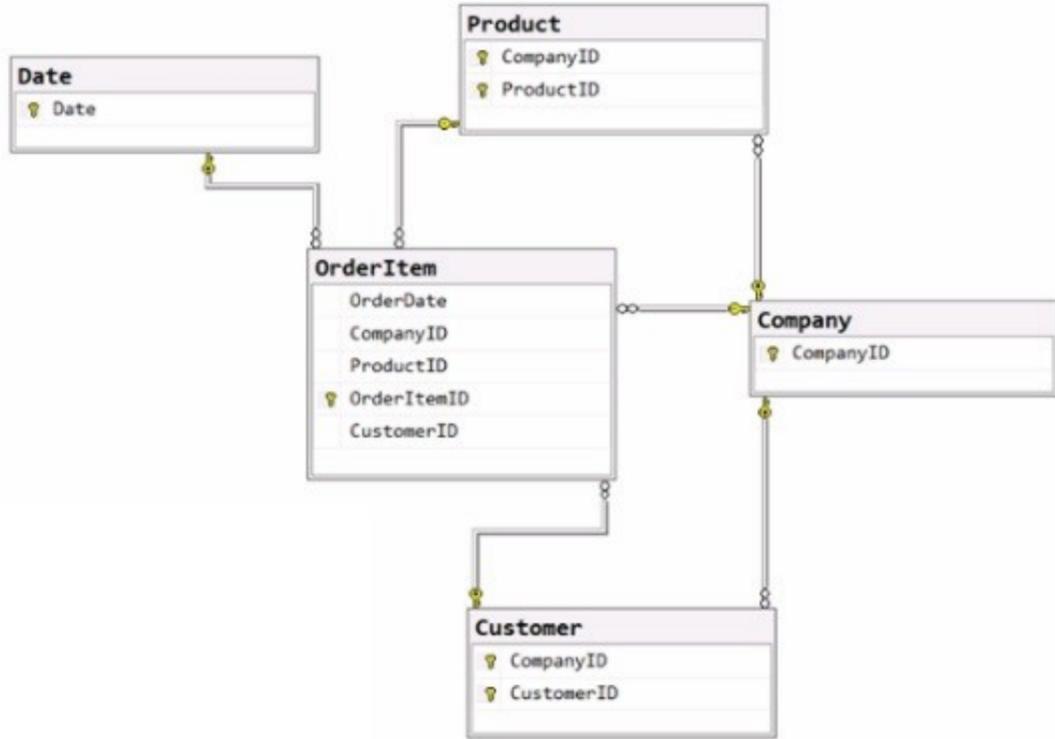
The primary keys of the tables are indicated by a key symbol beside the columns involved in each key.
You need to create a dimensional data model that will enable the analysis of order items by date, product, and customer.
What should you include in the solution? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
? The relationship between OrderItem and Product must be based on: Both the CompanyID and the ProductID columns
? The Company entity must be: Denormalized into the Customer and Product
entities
In a dimensional model, the relationships are typically based on foreign key constraints between the fact table (OrderItem) and dimension tables (Product, Customer, Date). Since CompanyID is present in both the OrderItem and Product tables, it acts as a foreign key in the relationship. Similarly, ProductID is a foreign key that relates these two tables. To enable analysis by date, product, and customer, the Company entity would need to be denormalized into the Customer and Product entities to ensure that the relevant company information is available within those dimensions for querying and reporting purposes. References =
? Dimensional modeling
? Star schema design

NEW QUESTION 23
- (Topic 2)
You have a Fabric tenant that contains a complex semantic model. The model is based on a star schema and contains many tables, including a fact table named Sales. You need to create a diagram of the model. The diagram must contain only the Sales table and related tables. What should you use from Microsoft Power BI Desktop?

A. data categories
B. Data view
C. Model view
D. DAX query view

Answer: C

Explanation:
 To create a diagram that contains only the Sales table and related tables, you should use the Model view (C) in Microsoft Power BI Desktop. This view allows you to visualize and manage the relationships between tables within your semantic model. References = Microsoft Power BI Desktop documentation outlines the functionalities available in Model view for managing semantic models.

NEW QUESTION 25
- (Topic 2)
You have a Fabric tenant named Tenant1 that contains a workspace named WS1. WS1 uses a capacity named C1 and contains a dawset named DS1. You need to ensure read- write access to DS1 is available by using the 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:
To ensure read-write access to DS1 is available by using the XMLA endpoint, the C1 settings (which refer to the capacity settings) should be modified first. XMLA endpoint configuration is a capacity feature, not specific to individual datasets or workspaces. References = The configuration of XMLA endpoints in Power BI capacities is detailed in the Power BI documentation on dataset management.

NEW QUESTION 29
- (Topic 2)
You have a Fabric tenant that contains a semantic model.
You need to prevent report creators from populating visuals by using implicit measures. What are two tools that you can use to achieve the goal? Each correct answer presents a complete solution.
NOTE: Each correct answer is worth one point.

A. Microsoft Power BI Desktop
B. Tabular Editor
C. Microsoft SQL Server Management Studio (SSMS)
D. DAX Studio

Answer: AB

Explanation:
Microsoft Power BI Desktop (A) and Tabular Editor (B) are the tools you can use to prevent report creators from using implicit measures. In Power BI Desktop, you can define explicit measures which can be used in visuals. Tabular Editor allows for advanced model editing, where you can enforce the use of explicit measures. References = Guidance on using explicit measures and preventing implicit measures in reports can be found in the Power BI and Tabular Editor official documentation.

NEW QUESTION 31
- (Topic 2)
You have a Microsoft Power BI report named Report1 that uses a Fabric semantic model. Users discover that Report1 renders slowly.
You open Performance analyzer and identify that a visual named Orders By Date is the slowest to render. The duration breakdown for Orders By Date is shown in the following table.

Name	Duration (ms)
DAX query	27
Visual display	39
Other	1047

What will provide the greatest reduction in the rendering duration of Report1?

A. Change the visual type of Orders By Dale.
B. Enable automatic page refresh.
C. Optimize the DAX query of Orders By Date by using DAX Studio.
D. Reduce the number of visuals in Report1.

Answer: C

Explanation:
Based on the duration breakdown provided, the major contributor to the rendering duration is categorized as "Other," which is significantly higher than DAX Query and Visual display times. This suggests that the issue is less likely with the DAX calculation or visual rendering times and more likely related to model performance or the complexity of the visual. However, of the options provided, optimizing the DAX query can be a crucial step, even if "Other" factors are dominant. Using DAX Studio, you can analyze and optimize the DAX queries that power your visuals for performance improvements. Here's how you might proceed:
? Open DAX Studio and connect it to your Power BI report.
? Capture the DAX query generated by the Orders By Date visual.
? Use the Performance Analyzer feature within DAX Studio to analyze the query.
? Look for inefficiencies or long-running operations.
? Optimize the DAX query by simplifying measures, removing unnecessary calculations, or improving iterator functions.
? Test the optimized query to ensure it reduces the overall duration.
References: The use of DAX Studio for query optimization is a common best practice for improving Power BI report performance as outlined in the Power BI documentation.

NEW QUESTION 35
DRAG DROP - (Topic 2)
You create a semantic model by using Microsoft Power BI Desktop. The model contains one security role named SalesRegionManager and the following tables:
• Sales
• SalesRegion
• Sales Ad dress
You need to modify the model to ensure that users assigned the SalesRegionManager role cannot see a column named Address in Sales Address.
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.

Actions

| Open the model in Power BI Desktop. |

| Set Object Level Security to **Default** for SalesRegionManager. |

| Set the Hidden property to **True**. |

| Open the model in Tabular Editor. |

| Select the **Address** column in SalesAddress. |

| Set Object Level Security to **None** for SalesRegionManager. |

Answer Area

A. Mastered
B. Not Mastered

Answer: A

Explanation:
To ensure that users assigned the SalesRegionManager role cannot see the Address column in the SalesAddress table, follow these steps in sequence:
? Open the model in Tabular Editor.
? Select the Address column in SalesAddress.
? Set Object Level Security to None for SalesRegionManager.

NEW QUESTION 37
- (Topic 2)
You have a Fabric tenant that contains a lakehouse.
You plan to query sales data files by using the SQL endpoint. The files will be in an Amazon Simple Storage Service (Amazon S3) storage bucket.
You need to recommend which file format to use and where to create a shortcut. Which two actions should you include in the recommendation? Each correct answer
presents part of the solution.
NOTE: Each correct answer is worth one point.

A. Create a shortcut in the Files section.
B. Use the Parquet format
C. Use the CSV format.
D. Create a shortcut in the Tables section.
E. Use the delta format.

Answer: BD

Explanation:
You should use the Parquet format (B) for the sales data files because it is optimized for performance with large datasets in analytical processing and create a shortcut in the Tables section (D) to facilitate SQL queries through the lakehouse's SQL endpoint. References = The best practices for working with file formats and shortcuts in a lakehouse environment are covered in the lakehouse and SQL endpoint documentation provided by the cloud data platform services.

NEW QUESTION 39
HOTSPOT - (Topic 2)
You have a Microsoft Power BI semantic model. You plan to implement calculation groups.
You need to create a calculation item that will change the context from the selected date to month-to-date (MTD).
How should you complete the DAX expression? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

CALCULATE ▼
CALCULATE
GENERATE
MEASURE

SELECTEDMEASURE ▼	(),
COMBINEVALUES	
SELECTEDMEASURE	
SELECTEDVALUE	

A. Mastered
B. Not Mastered

Answer: A

Explanation:
To create a calculation item that changes the context from the selected date to month-to- date (MTD), the appropriate DAX expression involves using the

CALCULATE function to alter the filter context and the DATESMTD function to specify the month-to-date context. The correct completion for the DAX expression would be:
? In the first dropdown, select CALCULATE.
? In the second dropdown, select SELECTEDMEASURE. This would create a DAX expression in the form:
CALCULATE(SELECTEDMEASURE(),
DATESMTD('Date'[DateColumn])
)

NEW QUESTION 44
- (Topic 2)
You are analyzing the data in a Fabric notebook.
You have a Spark DataFrame assigned to a variable named df.
You need to use the Chart view in the notebook to explore the data manually. Which function should you run to make the data available in the Chart view?

A. displayMTML
B. show
C. write
D. display

Answer: D

Explanation:
 The display function is the correct choice to make the data available in the Chart view within a Fabric notebook. This function is used to visualize Spark DataFrames in various formats including charts and graphs directly within the notebook environment. References = Further explanation of the display function can be found in the official documentation on Azure Synapse Analytics notebooks.

NEW QUESTION 47
HOTSPOT - (Topic 2)
You have a Fabric tenant that contains lakehouse named Lakehousel. Lakehousel contains a Delta table with eight columns. You receive new data that contains the same eight columns and two additional columns.
You create a Spark DataFrame and assign the DataFrame to a variable named df. The DataFrame contains the new data. You need to add the new data to the Delta table to meet the following requirements:
• Keep all the existing rows.
• Ensure that all the new data is added to the table.
How should you complete the code? To answer, select the appropriate options in the answer area.

Answer Area
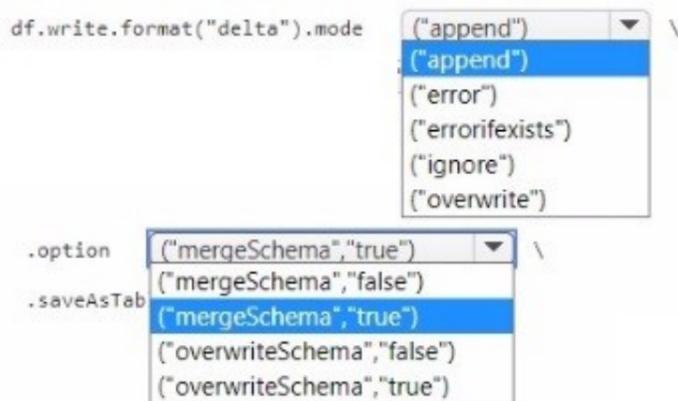
A. Mastered
B. Not Mastered

Answer: A

Explanation:
o add new data to the Delta table while meeting the specified requirements:
? You should use the append mode to ensure that all new data is added to the table without affecting the existing rows.
? You should set the mergeSchema option to true to allow the schema of the Delta table to be updated with the new columns found in the DataFrame.
The completed code would look like this:
df.write.format("delta").mode("append") option("mergeSchema", "true") saveAsTable("Lakehouse1.TableName")

NEW QUESTION 49
DRAG DROP - (Topic 2)
You are implementing a medallion architecture in a single Fabric workspace.
You have a lakehouse that contains the 8ronze and Silver layers and a warehouse that contains the Gold layer.
You create the items required to populate the layers as shown in the following table.

Layer	Data integration tool
Bronze	Pipelines with Copy activities
Silver	Dataflows
Gold	Stored procedures

You need to ensure that the layers are populated daily in sequential order such that Silver is populated only after Bronze is complete, and Gold is populated only after Silver is complete. The solution must minimize development effort and complexity.
What should you use to execute each set of items? To answer, drag the appropriate options to the correct items. Each option 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.

Execution Methods	Answer Area	
⠿ A pipeline Copy activity	Orchestration pipeline:	
⠿ A pipeline Dataflow activity	Bronze layer:	
⠿ A pipeline Stored procedure activity	Silver layer:	
⠿ A schedule	Gold layer:	
⠿ A Spark job definition		
⠿ An Invoke pipeline activity		

A. Mastered
B. Not Mastered

Answer: A

Explanation:
To execute each set of items in sequential order with minimized development effort and complexity, you should use the following options:
? Orchestration pipeline: Use a pipeline with an Invoke pipeline activity. This allows for orchestrating and scheduling the execution of other pipelines, ensuring they run in the correct sequence.
? Bronze layer: Implement a pipeline Copy activity. This aligns with the table indicating that the Bronze layer uses pipelines with Copy activities for data integration.
? Silver layer: Implement a pipeline Dataflow activity. The table specifies that Dataflows are used for the Silver layer.
? Gold layer: Implement a pipeline Stored procedure activity. Stored procedures are specified for the Gold layer according to the table.

NEW QUESTION 50
HOTSPOT - (Topic 2)
You have a Fabric tenant that contains a lakehouse.
You are using a Fabric notebook to save a large DataFrame by using the following code.

```
df.write.partitionBy("year","month","day").mode("overwrite").parquet("Files/SalesOrder")
```

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 results will form a hierarchy of folders for each partition key.	○	○
The resulting file partitions can be read in parallel across multiple nodes.	○	○
The resulting file partitions will use file compression.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
? The results will form a hierarchy of folders for each partition key. - Yes
? The resulting file partitions can be read in parallel across multiple nodes. - Yes
? The resulting file partitions will use file compression. - No
Partitioning data by columns such as year, month, and day, as shown in the DataFrame write operation, organizes the output into a directory hierarchy that reflects the partitioning structure. This organization can improve the performance of read operations, as queries that filter by the partitioned columns can scan only the relevant directories. Moreover, partitioning facilitates parallelism because each partition can be processed independently across different nodes in a distributed system like Spark. However, the code snippet provided does not explicitly specify that file compression should be used, so we cannot assume that the output will be compressed without additional context.
References =
? DataFrame write partitionBy
? Apache Spark optimization with partitioning

NEW QUESTION 54
- (Topic 2)
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 Power Query step to identify the maximum values for the numeric columns.
Which function should you include in the step?

A. Tabl
B. MaxN
C. Table.Max
D. Table.Range
E. Table.Profile

Answer: B

Explanation:
The Table.Max function should be used in a Power Query step to identify the maximum values for the numeric columns. This function is designed to calculate the maximum value across each column in a table, which suits the requirement of finding maximum values for numeric columns. References = For detailed information on Power Query functions, including Table.Max, please refer to Power Query M function reference.

NEW QUESTION 58
- (Topic 2)
You have a Fabric workspace that contains a DirectQuery semantic model. The model queries a data source that has 500 million rows.
You have a Microsoft Power BI report named Report1 that uses the model. Report! contains visuals on multiple pages.
You need to reduce the query execution time for the visuals on all the pages.
What are two features that you can use? Each correct answer presents a complete solution.
NOTE: Each correct answer is worth one point.

A. user-defined aggregations
B. automatic aggregation
C. query caching
D. OneLake integration

Answer: AB

Explanation:
 User-defined aggregations (A) and query caching (C) are two features that can help reduce query execution time. User-defined aggregations allow precalculation of large datasets, and query caching stores the results of queries temporarily to speed up future queries. References = Microsoft Power BI documentation on performance optimization offers in-depth knowledge on these features.

NEW QUESTION 63
......

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

DP-600 Practice Exam Features:

* DP-600 Questions and Answers Updated Frequently

* DP-600 Practice Questions Verified by Expert Senior Certified Staff

* DP-600 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* DP-600 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The DP-600 Practice Test Here