Microsoft

Exam Questions DP-600

Implementing Analytics Solutions Using Microsoft Fabric

NEW QUESTION 1
HOTSPOT - (Topic 1)
You need to create a DAX measure to calculate the average overall satisfaction score.
How should you complete the DAX code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

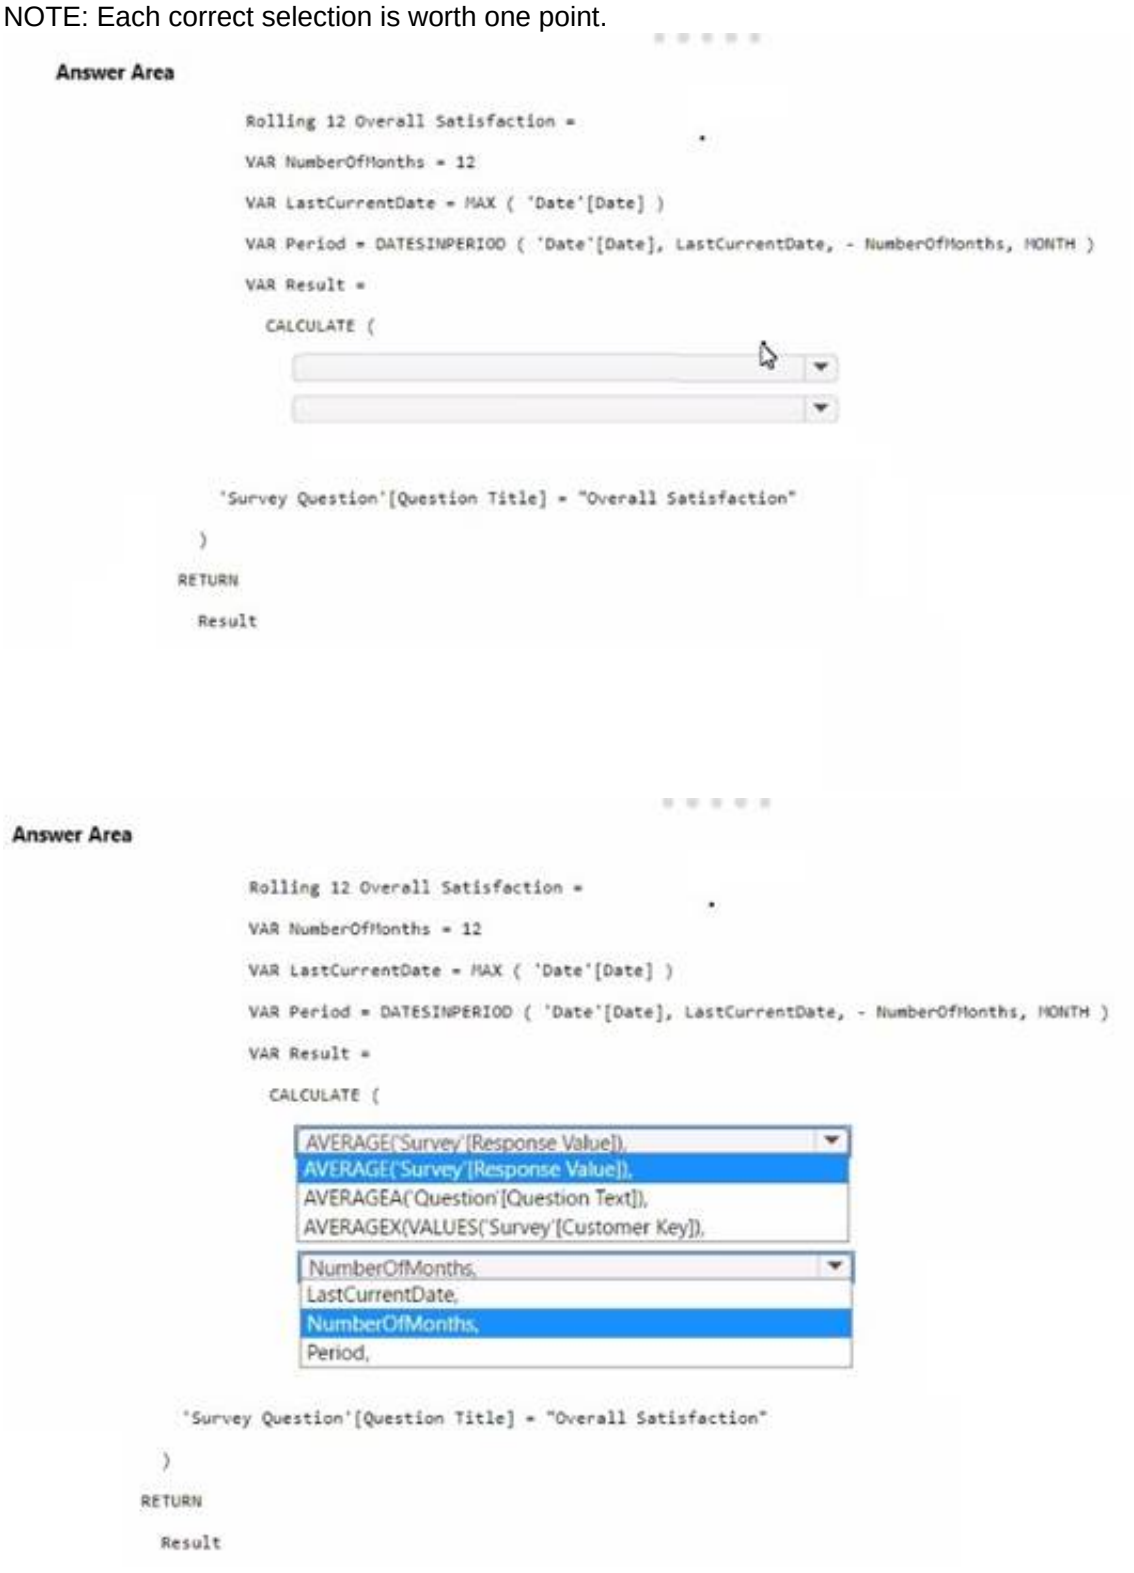

A. Mastered
B. Not Mastered
Answer: A

Explanation:
? The measure should use the AVERAGE function to calculate the average value.
? It should reference the Response Value column from the 'Survey' table.
? The 'Number of months' should be used to define the period for the average calculation.
To calculate the average overall satisfaction score using DAX, you would need to use the AVERAGE function on the response values related to satisfaction questions. The DATESINPERIOD function will help in calculating the rolling average over the last 12 months.

NEW QUESTION 2
- (Topic 1)
You need to implement the date dimension in the data store. The solution must meet the technical requirements.
What are two ways to achieve the goal? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. Populate the date dimension table by using a dataflow.
B. Populate the date dimension table by using a Stored procedure activity in a pipeline.
C. Populate the date dimension view by using T-SQL.
D. Populate the date dimension table by using a Copy activity in a pipeline.

Answer: AB

Explanation:
Both a dataflow (A) and a Stored procedure activity in a pipeline (B) are capable of creating and populating a date dimension table. A dataflow can perform the transformation needed to create the date dimension, and it aligns with the preference for using low-code tools for data ingestion when possible. A Stored procedure could be written to generate the necessary date dimension data and executed within a pipeline, which also adheres to the technical requirements for the PoC.

NEW QUESTION 3
- (Topic 2)
You have a Fabric tenant that contains 30 CSV files in OneLake. The files are updated daily.
You create a Microsoft Power Bl 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 4
- (Topic 2)
You have a Fabric tenant that uses a Microsoft tower Bl 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 5
HOTSPOT - (Topic 2)
You have a Fabric tenant that contains a warehouse named Warehouse1. Warehouse1 contains a fact table named FactSales that has one billion rows. You run the following T- SQL statement.
CREATE TABLE test.FactSales AS CLONE OF Dbo.FactSales;
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
A replica of `dbo.Sales` is created in the test schema by copying the metadata only.	○	○
Additional schema changes to `dbo.FactSales` will also apply to `test.FactSales`.	○	○
Additional data changes to `dbo.FactSales` will also apply to `test.FactSales`.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
? A replica of dbo.Sales is created in the test schema by copying the metadata only.
- No
? Additional schema changes to dbo.FactSales will also apply to test.FactSales. - No
? Additional data changes to dbo.FactSales will also apply to test.FactSales. - Yes
The CREATE TABLE AS CLONE statement creates a copy of an existing table, including its data and any associated data structures, like indexes. Therefore, the statement does not merely copy metadata; it also copies the data. However, subsequent schema changes to the original table do not automatically propagate to the cloned table. Any data changes in the original table after the clone operation will not be reflected in the clone unless explicitly updated.
References =
? CREATE TABLE AS SELECT (CTAS) in SQL Data Warehouse

NEW QUESTION 6
- (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 7
- (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 Bl 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 8
HOTSPOT - (Topic 2)
You have a data warehouse that contains a table named Stage. Customers. Stage- Customers contains all the customer record updates from a customer relationship management (CRM) system. There can be multiple updates per customer
You need to write a T-SQL query that will return the customer ID, name, postal code, and the last updated time of the most recent row for each customer ID.
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

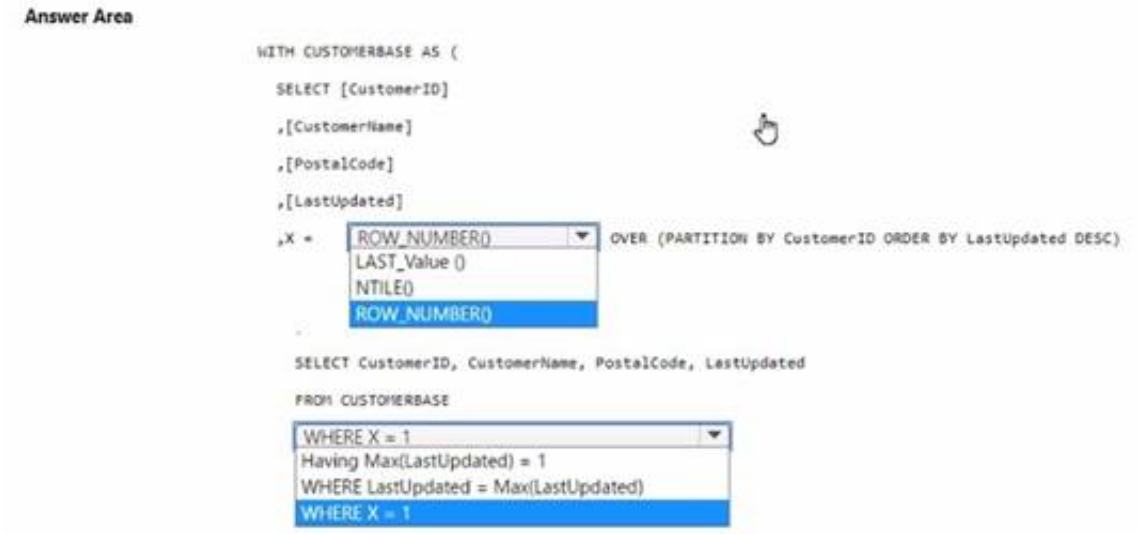

A. Mastered
B. Not Mastered

Answer: A

Explanation:
? In the ROW_NUMBER() function, choose OVER (PARTITION BY CustomerID ORDER BY LastUpdated DESC).
? In the WHERE clause, choose WHERE X = 1.
To select the most recent row for each customer ID, you use the ROW_NUMBER() window function partitioned by CustomerID and ordered by LastUpdated in descending order. This will assign a row number of 1 to the most recent update for each customer. By selecting rows where the row number (X) is 1, you get the latest update per customer. References =
? Use the OVER clause to aggregate data per partition
? Use window functions

NEW QUESTION 9
- (Topic 2)
You have a Fabric notebook that has the Python code and output shown in the following exhibit.

```
# Look at a histogram of tips by count by using Matplotlib

ax1 = sampled_taxi_pd_df['tipAmount'].plot(kind='hist', bins=25, facecolor='lightblue'
ax1.set_title('Tip amount distribution')
ax1.set_xlabel('Tip Amount ($)')
ax1.set_ylabel('Counts')
plt.suptitle('')
plt.show()
```

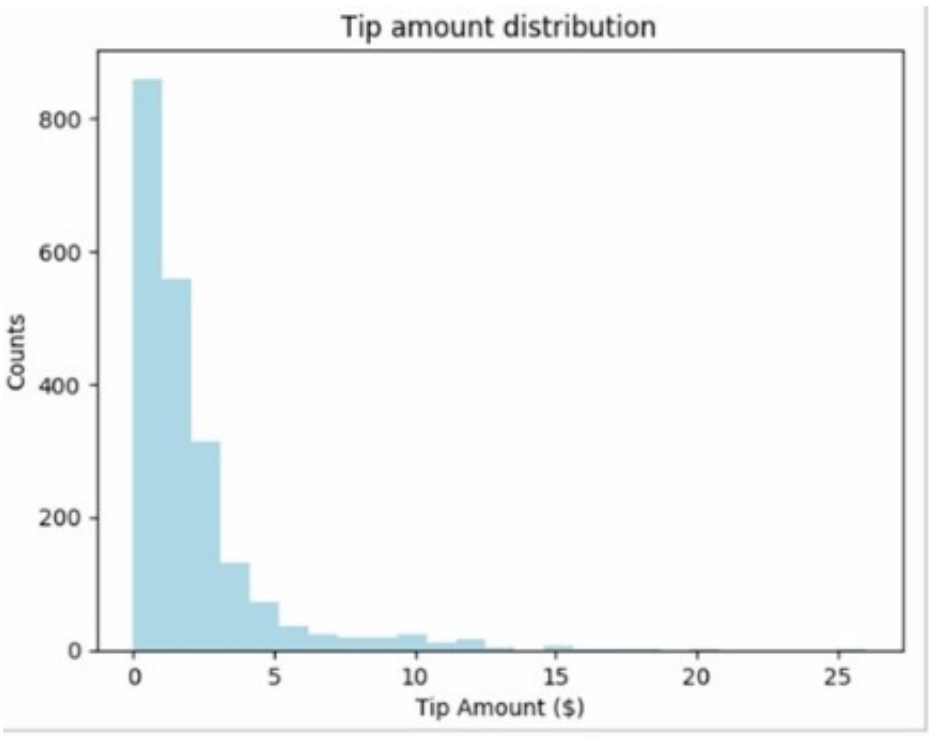

Which type of analytics are you performing?

A. predictive
B. descriptive
C. prescriptive
D. diagnostic

Answer: B

Explanation:
The Python code and output shown in the exhibit display a histogram, which is a representation of the distribution of data. This kind of analysis is descriptive analytics, which is used to describe or summarize the features of a dataset. Descriptive analytics answers the question of "what has happened" by providing insight into past data through tools such as mean, median, mode, standard deviation, and graphical representations like histograms.
References: Descriptive analytics and the use of histograms as a way to visualize data distribution are basic concepts in data analysis, often covered in introductory analytics and Python programming resources.

NEW QUESTION 10
HOTSPOT - (Topic 2)
You have the source data model shown in the following exhibit.

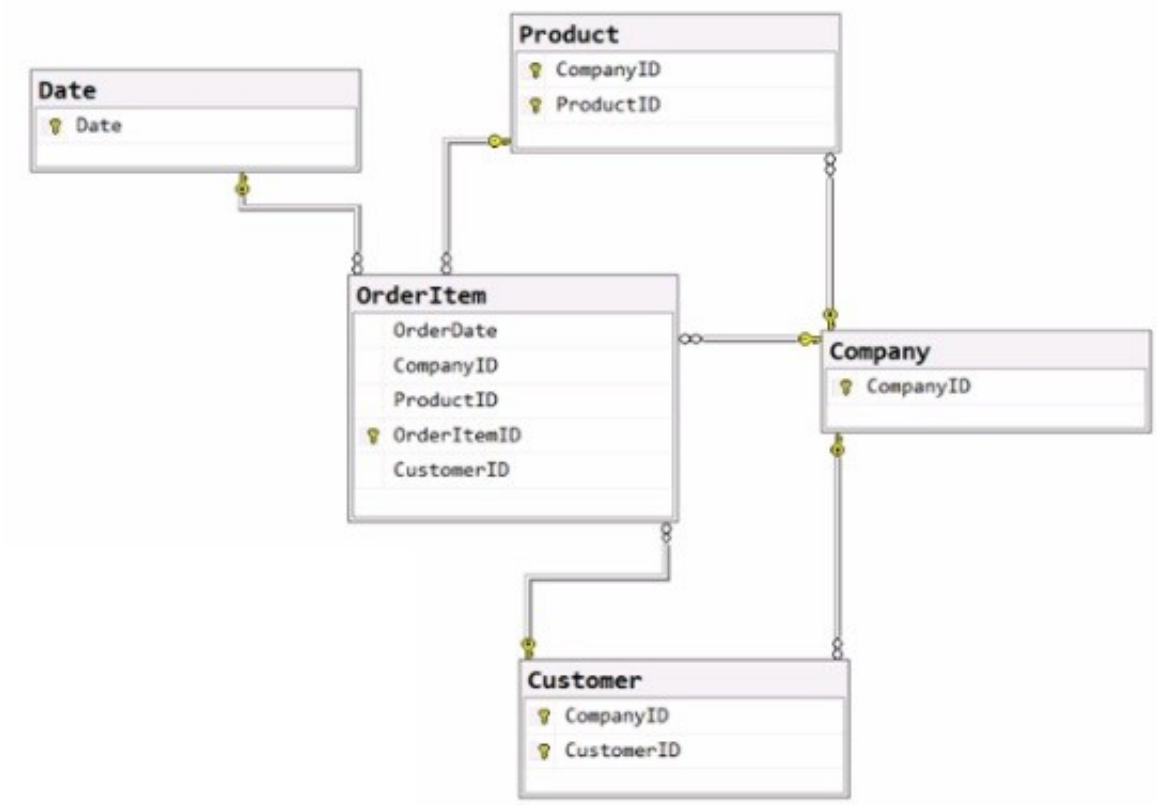

The primary keys of the tables are indicated by a key symbol beside the columns involved in each key.
You need to create a dimensional data model that will enable the analysis of order items by date, product, and customer.
What should you include in the solution? To answer, select the appropriate options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

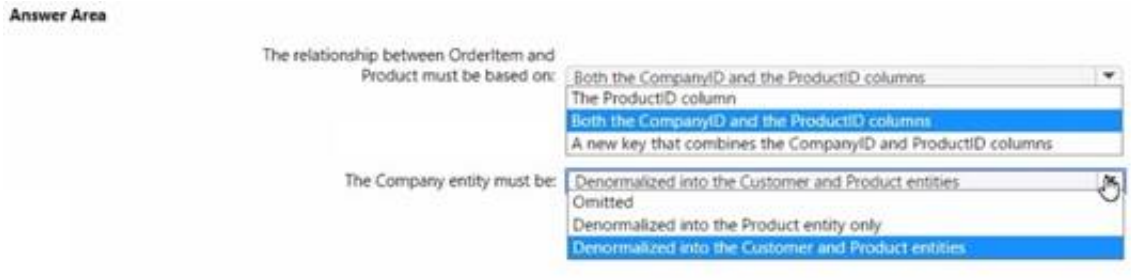

A. Mastered
B. Not Mastered

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 10
- (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 Bl 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 12
- (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 14
HOTSPOT - (Topic 2)
You have a Fabric tenant that contains a lakehouse named Lakehouse1. Lakehouse1 contains a table named Nyctaxi_raw. Nyctaxi_raw contains the following columns.

Name	Data type
pickupDateTime	Timestamp
passengerCount	Integer
fareAmount	Double
paymentType	String
tipAmount	Double

You create a Fabric notebook and attach it to lakehouse1.
You need to use PySpark code to transform the data. The solution must meet the following requirements:
• Add a column named pickupDate that will contain only the date portion of pickupDateTime.
• Filter the DataFrame to include only rows where fareAmount is a positive number that is less than 100.
How should you complete the code? 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:
? Add the pickupDate column: .withColumn("pickupDate", df["pickupDateTime"].cast("date"))
? Filter the DataFrame: .filter("fareAmount > 0 AND fareAmount < 100")
In PySpark, you can add a new column to a DataFrame using the .withColumn method, where the first argument is the new column name and the second argument is the expression to generate the content of the new column. Here, we use the .cast("date") function to extract only the date part from a timestamp. To filter the DataFrame, you use the .filter method with a condition that selects rows where fareAmount is greater than 0 and less than 100, thus ensuring only positive values less than 100 are included.

NEW QUESTION 16
- (Topic 2)
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 EC.
C. Spark SQL
D. PySpark

Answer: CD

Explanation:
The two languages you can use to perform model scoring in a Fabric notebook using the predict function are Spark SQL (option C) and PySpark (option D). These are both part of the Apache Spark ecosystem and are supported for machine learning tasks in a Fabric environment. References = You can find more information about model scoring and supported languages in the context of Fabric notebooks in the official documentation on Azure Synapse Analytics.

NEW QUESTION 17
- (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 .sumary ()
Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:
Yes, the df.summary() method does meet the goal. This method is used to compute specified statistics for numeric and string columns. By default, it provides statistics such as count, mean, stddev, min, and max. References = The PySpark API documentation details the summary() function and the statistics it provides.

NEW QUESTION 22
- (Topic 2)
You have a Fabric workspace named Workspace1 that contains a data flow named Dataflow1. Dataflow1 contains a query that returns the data shown in the following exhibit.

You need to transform the date columns into attribute-value pairs, where columns become rows.
You select the VendorID column.
Which transformation should you select from the context menu of the VendorID column?

A. Group by
B. Unpivot columns
C. Unpivot other columns
D. Split column
E. Remove other columns

Answer: B

Explanation:
The transformation you should select from the context menu of the VendorID column to transform the date columns into attribute-value pairs, where columns become rows, is Unpivot columns (B). This transformation will turn the selected columns into rows with two new columns, one for the attribute (the original column names) and one for the value (the data from the cells). References = Techniques for unpivoting columns are covered in the Power Query documentation, which explains how to use the transformation in data modeling.

NEW QUESTION 25
- (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.show()
Does this meet the goal?

A. Yes
B. No

Answer: B

Explanation:
The df.show() method also does not meet the goal. It is used to show the contents of the DataFrame, not to compute statistical functions. References = The usage of the show() function is documented in the PySpark API documentation.

NEW QUESTION 27
- (Topic 2)
You have a Fabric tenant that contains a semantic model named Model1. Model1 uses Import mode. Model1 contains a table named Orders. Orders has 100 million rows and the following fields.

Name	Data type	Description
OrderId	Integer	Column imported from the source
OrderDateTime	Date/time	Column imported from the source
Quantity	Integer	Column imported from the source
Price	Decimal	Column imported from the source
TotalSalesAmount	Decimal	Calculated column that multiplies Quantity and Price
TotalQuantity	Integer	Measure

You need to reduce the memory used by Model! and the time it takes to refresh the model. Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct answer is worth one point.

A. Split OrderDateTime into separate date and time columns.
B. Replace TotalQuantity with a calculated column.
C. Convert Quantity into the Text data type.
D. Replace TotalSalesAmount with a measure.

Answer: AD

Explanation:
To reduce memory usage and refresh time, splitting the OrderDateTime into separate date and time columns (A) can help optimize the model because date/time data types can be more memory-intensive than separate date and time columns. Moreover, replacing TotalSalesAmount with a measure (D) instead of a calculated column ensures that the calculation is performed at query time, which can reduce the size of the model as the value is not stored but calculated on the fly. References = The best practices for optimizing Power BI models are detailed in the Power BI documentation, which recommends using measures for calculations that don't need to be stored and adjusting data types to improve performance.

NEW QUESTION 31
- (Topic 2)
You have a Fabric tenant that contains a workspace named Workspace^ Workspacel is assigned to a Fabric capacity.
You need to recommend a solution to provide users with the ability to create and publish custom Direct Lake semantic models by using external tools. The solution must follow the principle of least privilege.
Which three actions in the Fabric Admin portal should you include in the recommendation? Each correct answer presents part of the solution.
NOTE: Each correct answer is worth one point.

A. From the Tenant settings, set Allow XMLA Endpoints and Analyze in Excel with on- premises datasets to Enabled
B. From the Tenant settings, set Allow Azure Active Directory guest users to access Microsoft Fabric to Enabled
C. From the Tenant settings, select Users can edit data models in the Power Bl service.
D. From the Capacity settings, set XMLA Endpoint to Read Write
E. From the Tenant settings, set Users can create Fabric items to Enabled
F. From the Tenant settings, enable Publish to Web

Answer: ACD

Explanation:
For users to create and publish custom Direct Lake semantic models using external tools, following the principle of least privilege, the actions to be included are

enabling XMLA Endpoints (A), editing data models in Power BI service (C), and setting XMLA Endpoint to Read-Write in the capacity settings (D). References = More information can be found in the Admin portal of the Power BI service documentation, detailing tenant and capacity settings.

NEW QUESTION 36
DRAG DROP - (Topic 2)
You are implementing two dimension tables named Customers and Products in a Fabric warehouse.
You need to use slowly changing dimension (SCO) to manage the versioning of data. The solution must meet the requirements shown in the following table.

Table	Change action
Customers	Create a new version of the row.
Products	Overwrite the existing value in the latest row.

Which type of SCD should you use for each table? To answer, drag the appropriate SCD types to the correct tables. Each SCD type 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:
For the Customers table, where the requirement is to create a new version of the row, you would use:
? Type 2 SCD: This type allows for the creation of a new record each time a change occurs, preserving the history of changes over time.
For the Products table, where the requirement is to overwrite the existing value in the latest row, you would use:
? Type 1 SCD: This type updates the record directly, without preserving historical data.

NEW QUESTION 38
- (Topic 2)
You have a Fabric tenant tha1 contains a takehouse 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

Explanation:
No, the REFRESH TABLE statement does not provide information on whether maintenance tasks were performed. It only updates the metadata of a table to reflect any changes on the data files. References = The use and effects of the REFRESH TABLE command are explained in the Spark SQL documentation.

NEW QUESTION 40
- (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 42
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 46
- (Topic 2)
You have a Fabric tenant.
You are creating a Fabric Data Factory pipeline.
You have a stored procedure that returns the number of active customers and their average sales for the current month.
You need to add an activity that will execute the stored procedure in a warehouse. The returned values must be available to the downstream activities of the pipeline.
Which type of activity should you add?

A. Stored procedure
B. Get metadata
C. Lookup
D. Copy data

Answer: C

Explanation:
In a Fabric Data Factory pipeline, to execute a stored procedure and make the returned values available for downstream activities, the Lookup activity is used. This activity can retrieve a dataset from a data store and pass it on for further processing. Here's how you would use the Lookup activity in this context:
? Add a Lookup activity to your pipeline.
? Configure the Lookup activity to use the stored procedure by providing the necessary SQL statement or stored procedure name.
? In the settings, specify that the activity should use the stored procedure mode.
? Once the stored procedure executes, the Lookup activity will capture the results and make them available in the pipeline's memory.
? Downstream activities can then reference the output of the Lookup activity. References: The functionality and use of Lookup activity within Azure Data Factory is documented in Microsoft's official documentation for Azure Data Factory, under the section
for pipeline activities.

NEW QUESTION 49
- (Topic 2)
You have a Fabric tenant that contains a warehouse.
A user discovers that a report that usually takes two minutes to render has been running for 45 minutes and has still not rendered.
You need to identify what is preventing the report query from completing. Which dynamic management view (DMV) should you use?

A. sys.dm-exec_requests
B. sys.dn_.exec._sessions
C. sys.dm._exec._connections
D. sys.dm_pdw_exec_requests

Answer: D

Explanation:
The correct DMV to identify what is preventing the report query from completing is sys.dm_pdw_exec_requests (D). This DMV is specific to Microsoft Analytics Platform System (previously known as SQL Data Warehouse), which is the environment assumed to be used here. It provides information about all queries and load commands currently running or that have recently run. References = You can find more about DMVs in the Microsoft documentation for Analytics Platform System.

NEW QUESTION 53
HOTSPOT - (Topic 2)
You have a Microsoft Power B1 report and a semantic model that uses Direct Lake mode. From Power Si Desktop, you open Performance analyzer as shown in the following exhibit.

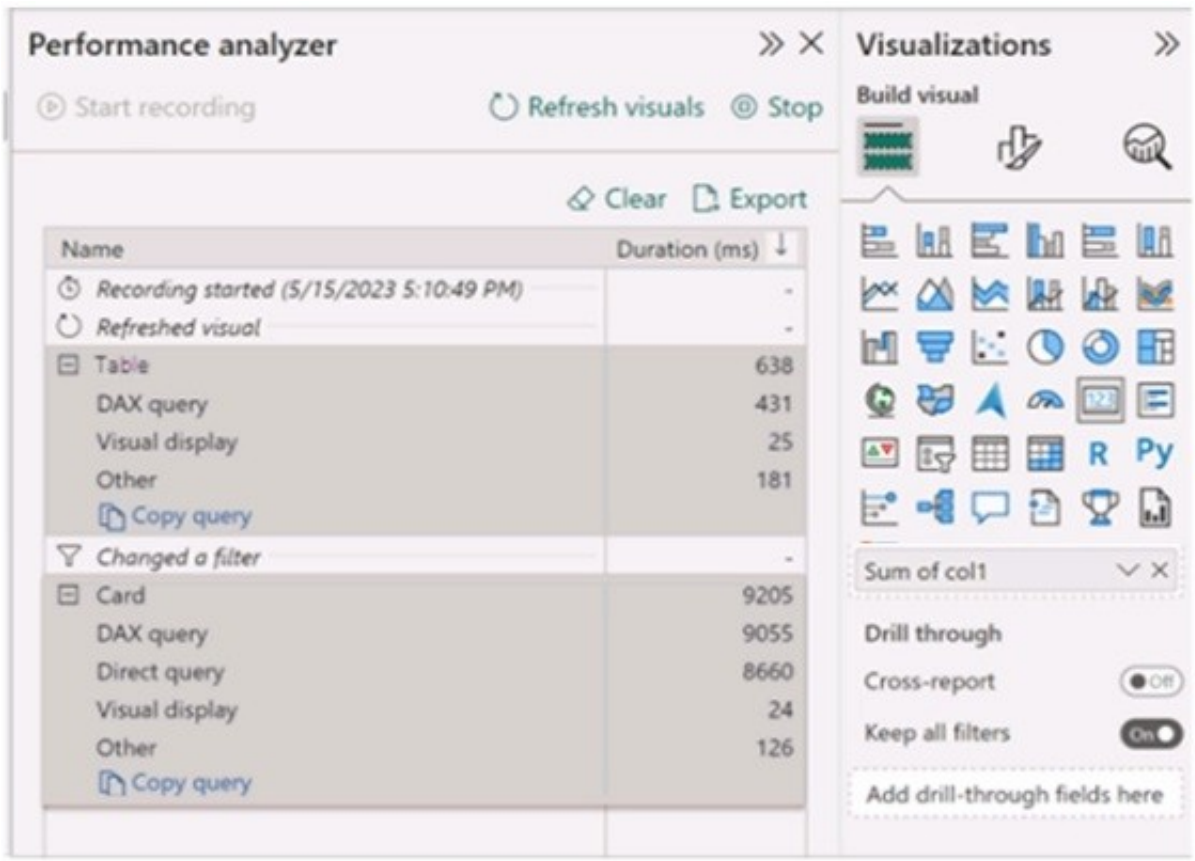

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic. NOTE: Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
? The Direct Lake fallback behavior is set to: DirectQueryOnly
? The query for the table visual is executed by using: DirectQuery
In the context of Microsoft Power BI, when using DirectQuery in Direct Lake mode, there is no caching of data and all queries are sent directly to the underlying data source. The Performance Analyzer tool shows the time taken for different operations, and from the options provided, it indicates that DirectQuery mode is being used for the visuals, which is consistent with the Direct Lake setting. DirectQueryOnly as the fallback behavior ensures that only DirectQuery will be used without reverting to import mode.

NEW QUESTION 57
- (Topic 2)
You need to provide Power Bl developers with access to the pipeline. The solution must meet the following requirements:
• Ensure that the developers can deploy items to the workspaces for Development and Test.
• Prevent the developers from deploying items to the workspace for Production.
• Follow the principle of least privilege.
Which three levels of access should you assign to the developers? Each correct answer presents part of the solution. NOTE: Each correct answer is worth one point.

A. Build permission to the production semantic models
B. Admin access to the deployment pipeline
C. Viewer access to the Development and Test workspaces
D. Viewer access to the Production workspace
E. Contributor access to the Development and Test workspaces
F. Contributor access to the Production workspace

Answer: BDE

Explanation:
To meet the requirements, developers should have Admin access to the deployment pipeline (B), Contributor access to the Development and Test workspaces (E), and Viewer access to the Production workspace (D). This setup ensures they can perform necessary actions in development and test environments without having the ability to affect production. References = The Power BI documentation on workspace access levels and deployment pipelines provides guidelines on assigning appropriate permissions.

NEW QUESTION 61
DRAG DROP - (Topic 2)
You have a Fabric tenant that contains a Microsoft Power Bl report named Report 1. Report1 is slow to render. You suspect that an inefficient DAX query is being executed.
You need to identify the slowest DAX query, and then review how long the query spends in the formula engine as compared to the storage engine.
Which five 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.

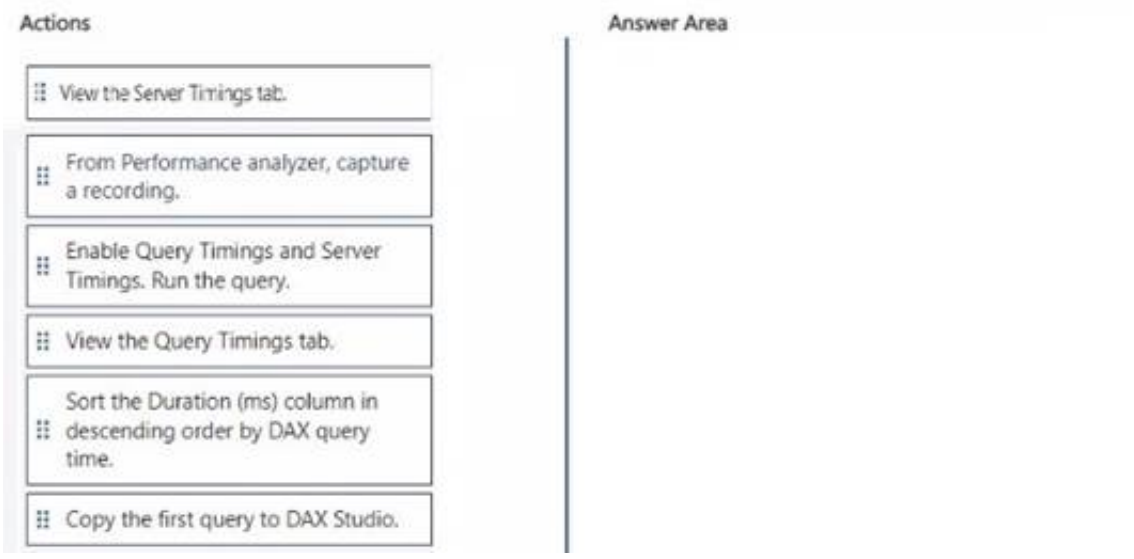

A. Mastered
B. Not Mastered

Answer: A

Explanation:
To identify the slowest DAX query and analyze the time it spends in the formula engine compared to the storage engine, you should perform the following actions in sequence:
? From Performance analyzer, capture a recording.
? View the Server Timings tab.
? Enable Query Timings and Server Timings. Run the query.
? View the Query Timings tab.
? Sort the Duration (ms) column in descending order by DAX query time.

NEW QUESTION 65
HOTSPOT - (Topic 2)
You have a Fabric warehouse that contains a table named Sales.Products. Sales.Products contains the following columns.

Name	Data type	Nullable
ProductID	Integer	No
ProductName	Varchar(30)	No
ListPrice	Decimal(18, 2)	No
WholesalePrice	Decimal(18, 2)	Yes
AgentPrice	Decimal(18, 2)	Yes

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

Name	Description
ProductID	Return the ProductID value
HighestSellingPrice	Returns the highest value from ListPrice, WholesalePrice, and AgentPrice
TradePrice	Returns the AgentPrice value if present, otherwise returns the WholesalePrice value if present, otherwise returns the ListPrice value

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

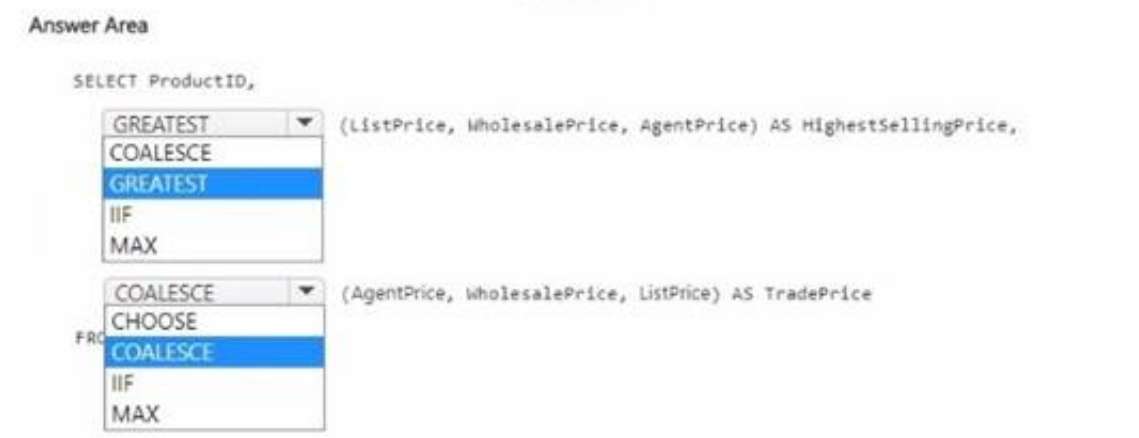

A. Mastered
B. Not Mastered

Answer: A

Explanation:
? For the HighestSellingPrice, you should use the GREATEST function to find the highest value from the given price columns. However, T-SQL does not have a GREATEST function as found in some other SQL dialects, so you would typically use a CASE statement or an IIF statement with nested MAX functions. Since neither of those are provided in the options, you should select MAX as a placeholder to indicate the function that would be used to find the highest value if

combining multiple MAX functions or a similar logic was available.
? For the TradePrice, you should use the COALESCE function, which returns the first non-null value in a list. The COALESCE function is the correct choice as it will return AgentPrice if it's not null; if AgentPrice is null, it will check WholesalePrice, and if that is also null, it will return ListPrice.
The complete code with the correct SQL functions would look like this:
SELECT ProductID,
MAX(ListPrice, WholesalePrice, AgentPrice) AS HighestSellingPrice, -- MAX is used as a placeholder
COALESCE(AgentPrice, WholesalePrice, ListPrice) AS TradePrice FROM Sales.Products
Select MAX for HighestSellingPrice and COALESCE for TradePrice in the answer area.

NEW QUESTION 70
.....

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