Question 1

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 dat

a. The solution must meet the requirements shown in the following table.

Table	Change action
Customers	Create a new version of the row.
:: Type 0	verwrite the existing value in the latest row.

:: Type 3

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.

SCD Types

:: Type 0	:: Type 1
:: Type 3	:: Type 3

Answer Area

Customers: [   ]

Products: [   ]

Answer:

Question 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	

)

Question 3

Question Type: DragDrop

You have a Fabric tenant that contains a Microsoft Power BI 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.

Actions

⠿ View the Server Timings tab.
⠿ From Performance analyzer, capture a recording.
⠿ 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.
⠿ Copy the first query to DAX Studio.

Answer Area

Question 4

Question Type: MultipleChoice

You have a Fabric tenant that contains a semantic model. The model contains 15 tables.

You need to programmatically change each column that ends in the word Key to meet the following requirements:

* Hide the column.

* Set Nullable to False.

* Set Summarize By to None

* Set Available in MDX to False.

* Mark the column as a key column.

What should you use?

Options:

A- Microsoft Power BI Desktop

B- Tabular Editor

C- ALM Toolkit

D- DAX Studio

Answer:

B

Explanation:

Tabular Editor is an advanced tool for editing Tabular models outside of Power BI Desktop that allows you to script out changes and apply them across multiple columns or tables. To accomplish the task programmatically, you would:

Open the model in Tabular Editor.

Create an Advanced Script using C# to iterate over all tables and their respective columns.

Within the script, check if the column name ends with 'Key'.

For columns that meet the condition, set the properties accordingly: IsHidden = true, IsNullable = false, SummarizeBy = None, IsAvailableInMDX = false.

Additionally, mark the column as a key column.

Save the changes and deploy them back to the Fabric tenant.

Question 5

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?

Options:

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.

Question 6

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?

Options:

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.

Question 7

Question Type: MultipleChoice

You have a Microsoft Power BI semantic model that contains measures. The measures use multiple calculate functions and a filter function.

You are evaluating the performance of the measures.

In which use case will replacing the filter function with the keepfilters function reduce execution time?

Options:

A- when the filter function uses a nested calculate function

B- when the filter function references a column from a single table that uses Import mode

C- when the filter function references columns from multiple tables

D- when the filter function references a measure

Answer:

A

Explanation:

The KEEPFILTERS function modifies the way filters are applied in calculations done through the CALCULATE function. It can be particularly beneficial to replace the FILTER function with KEEPFILTERS when the filter context is being overridden by nested CALCULATE functions, which may remove filters that are being applied on a column. This can potentially reduce execution time because KEEPFILTERS maintains the existing filter context and allows the nested CALCULATE functions to be evaluated more efficiently.

Question 8

Question Type: DragDrop

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 for SalesRegionManager. |

Answer Area

Question 9

You have a Microsoft Power BI semantic model.

You need to identify any surrogate key columns in the model that have the Summarize By property set to a value other than to None. The solution must minimize effort.

What should you use?

Options:

A- DAX Formatter in DAX Studio

B- Model view in Microsoft Power BI Desktop

C- Model explorer in Microsoft Power BI Desktop

D- Best Practice Analyzer in Tabular Editor

Answer:

D

Explanation:

To identify surrogate key columns with the 'Summarize By' property set to a value other than 'None,' the Best Practice Analyzer in Tabular Editor is the most efficient tool. The Best Practice Analyzer can analyze the entire model and provide a report on all columns that do not meet a specified best practice, such as having the 'Summarize By' property set correctly for surrogate key columns. Here's how you would proceed:

Open your Power BI model in Tabular Editor.

Go to the Advanced Scripting window.

Write or use an existing script that checks the 'Summarize By' property of each column.

Execute the script to get a report on the surrogate key columns that do not have their 'Summarize By' property set to 'None'.

You can then review and adjust the properties of the columns directly within the Tabular Editor.

Question 10

Question Type: **Hotspot**

You have a Fabric workspace named Workspace1 and an Azure Data Lake Storage Gen2 account named storage"!. Workspace1 contains a lakehouse named Lakehouse1.

You need to create a shortcut to storage! in Lakehouse1.

Which connection and endpoint should you specify? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Connection: | abfss ▼ |
| abfs |
| **abfss** |
| https |

Endpoint: | dfs ▼ |
| blob |
| **dfs** |
| file |

Question 11

You need to create a data loading pattern for a Type 1 slowly changing dimension (SCD).

Which two actions should you include in the process? Each correct answer presents part of the solution.

NOTE: Each correct answer is worth one point.

Options:

A- Update rows when the non-key attributes have changed.

B- Insert new rows when the natural key exists in the dimension table, and the non-key attribute values have changed.

C- Update the effective end date of rows when the non-key attribute values have changed.

D- Insert new records when the natural key is a new value in the table.

Answer:

A, D

Explanation:

For a Type 1 SCD, you should include actions that update rows when non-key attributes have changed (A), and insert new records when the natural key is a new value in the table (D). A Type 1 SCD does not track historical data, so you always overwrite the old data with the new data for a given key. Reference = Details on Type 1 slowly changing dimension patterns can be found in data warehousing literature and Microsoft's official documentation.

Question 12

You have a Fabric tenant that contains a lakehouse named Lakehouse1

Readings from 100 IoT 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.

Actions

Set the autoCompact table setting.

Set the optimizeWrite table setting.

Run the VACUUM command on a schedule.

Set the autoCompact SparkSession setting.

Run the OPTIMIZE command on a schedule.

Set the parallelDelete SparkSession setting.

Answer Area

Remove the files: _____

Combine the files: _____