



Microsoft

Exam Questions DP-500

Designing and Implementing Enterprise-Scale Analytics Solutions Using Microsoft Azure and Microsoft Power BI

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

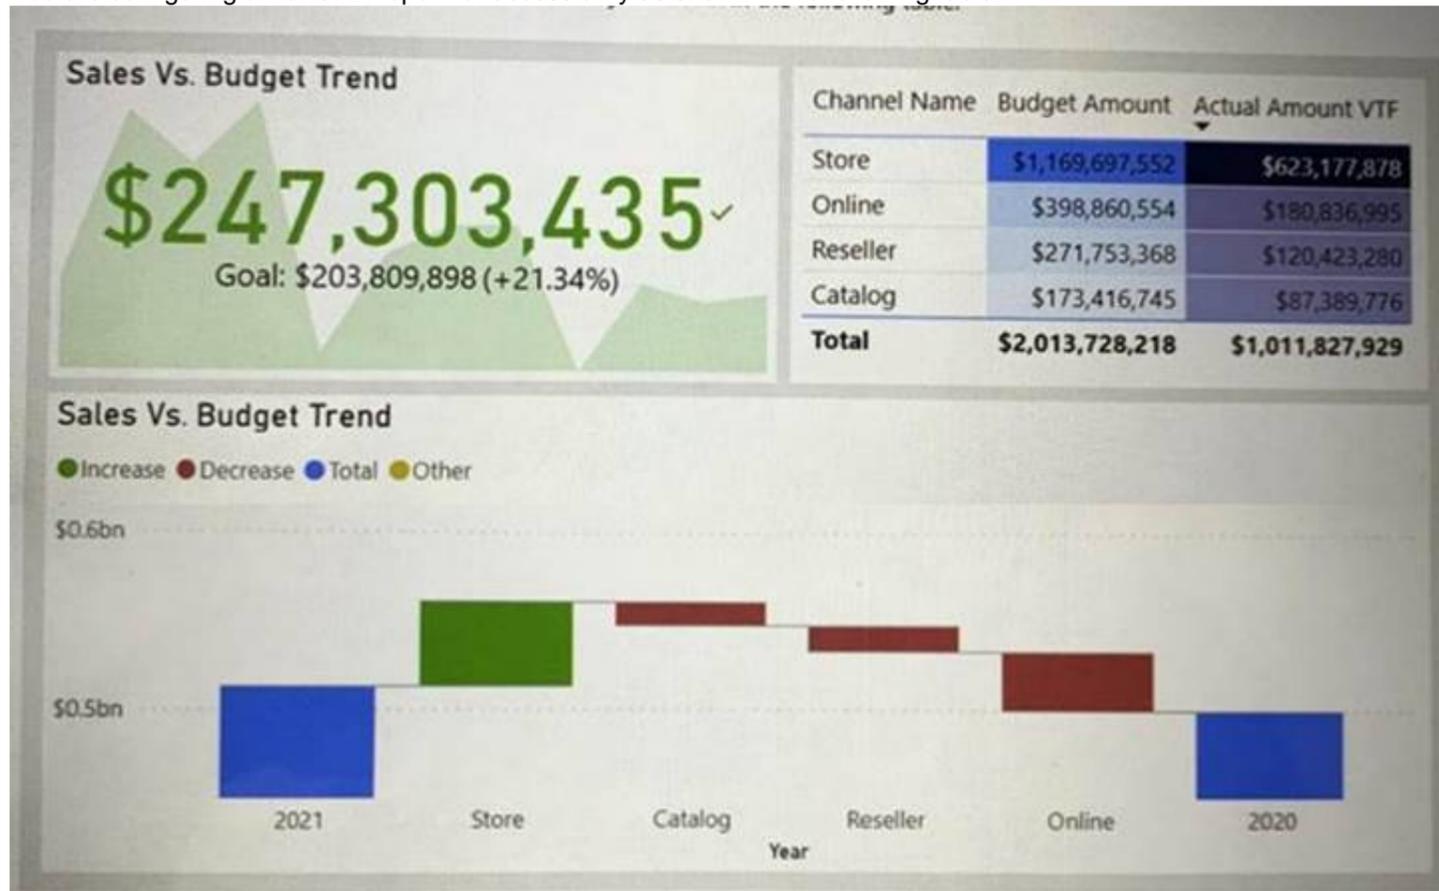
* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

- (Exam Topic 3)

You are configuring a Power BI report for accessibility as shown in the following table.



You need to change the default colors of all three visuals to make the report more accessible to users who have color vision deficiency. Which two settings should you configure in the Customize theme window? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Theme colors
- B. Sentiment colors
- C. Divergent colors
- D. First-level elements colors

Answer: AB

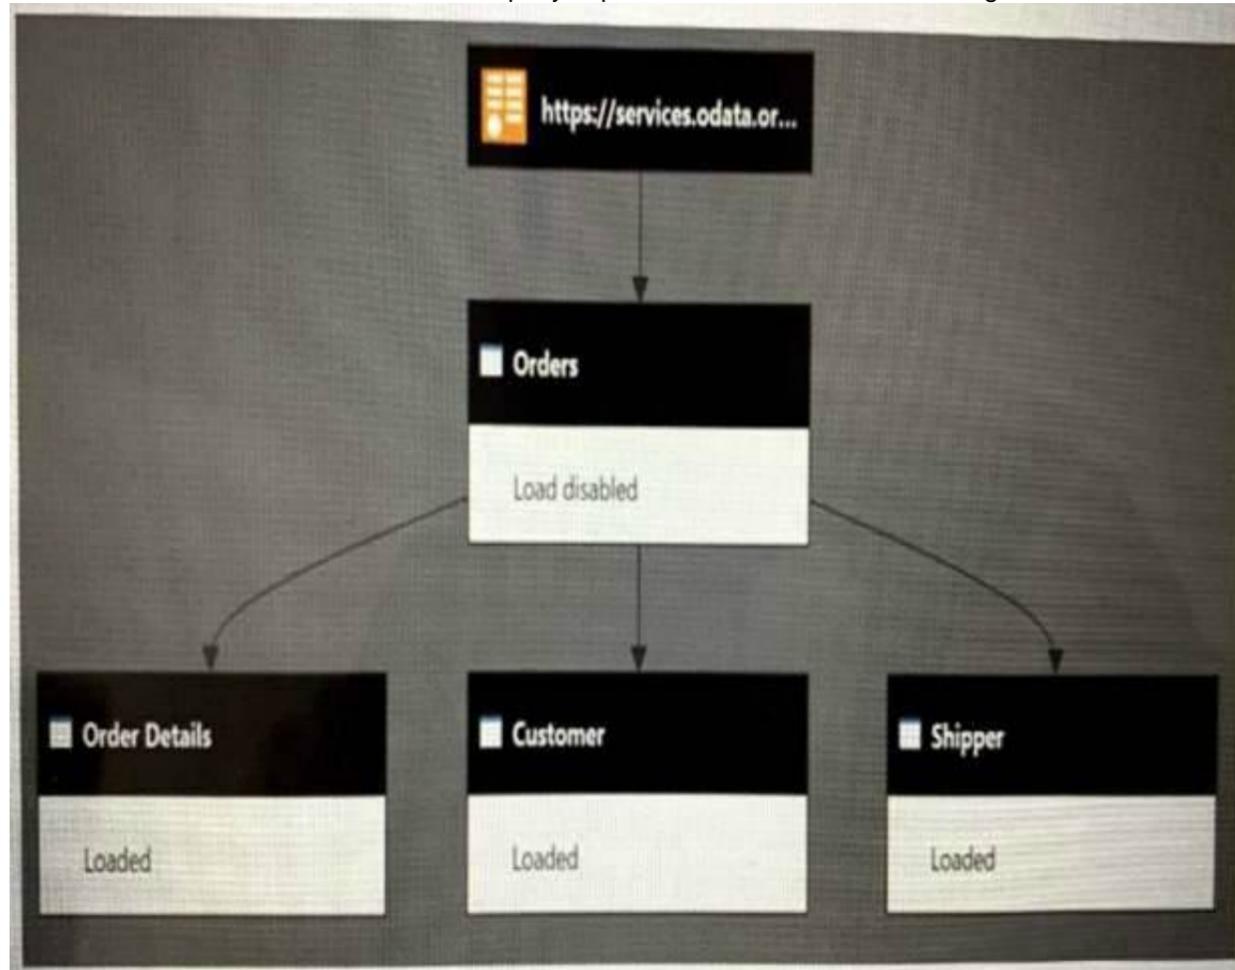
Explanation:

Reference: <https://docs.microsoft.com/en-us/power-bi/create-reports/desktop-report-themes>

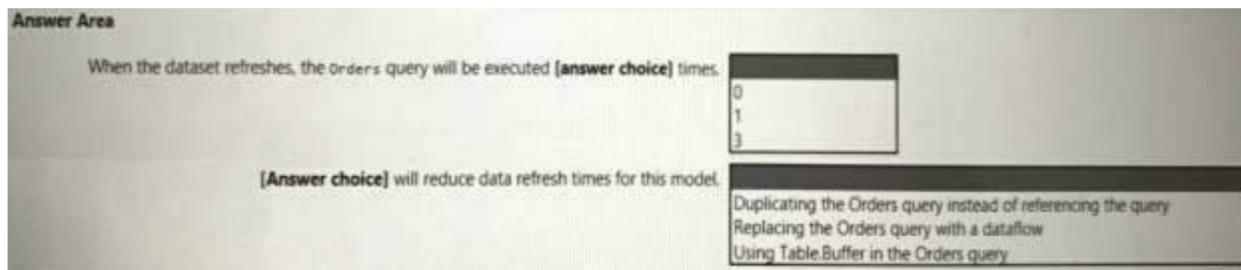
NEW QUESTION 2

- (Exam Topic 3)

You have a Power BI dataset that has the query dependencies 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:

Box 1: 3
 Power Query doesn't start at the first query and work down, it starts at the bottom (last) query and works backwards, so 3 tables from 1 will cause it to process that first source table 3 times.

Box 2: Using Table.Buffer in the Orders query

Table.Buffer buffers a table in memory, isolating it from external changes during evaluation. Buffering is shallow. It forces the evaluation of any scalar cell values, but leaves non-scalar values (records, lists, tables, and so on) as-is.

Note that using this function might or might not make your queries run faster. In some cases, it can make your queries run more slowly due to the added cost of reading all the data and storing it in memory, as well as the fact that buffering prevents downstream folding.

Example 1

Load all the rows of a SQL table into memory, so that any downstream operations will no longer be able to query the SQL server.

Usage let

Source = Sql.Database("SomeSQLServer", "MyDb"), MyTable = Source[Item="MyTable"][Data], BufferMyTable = Table.Buffer(dbo_MyTable)

in BufferMyTable Output

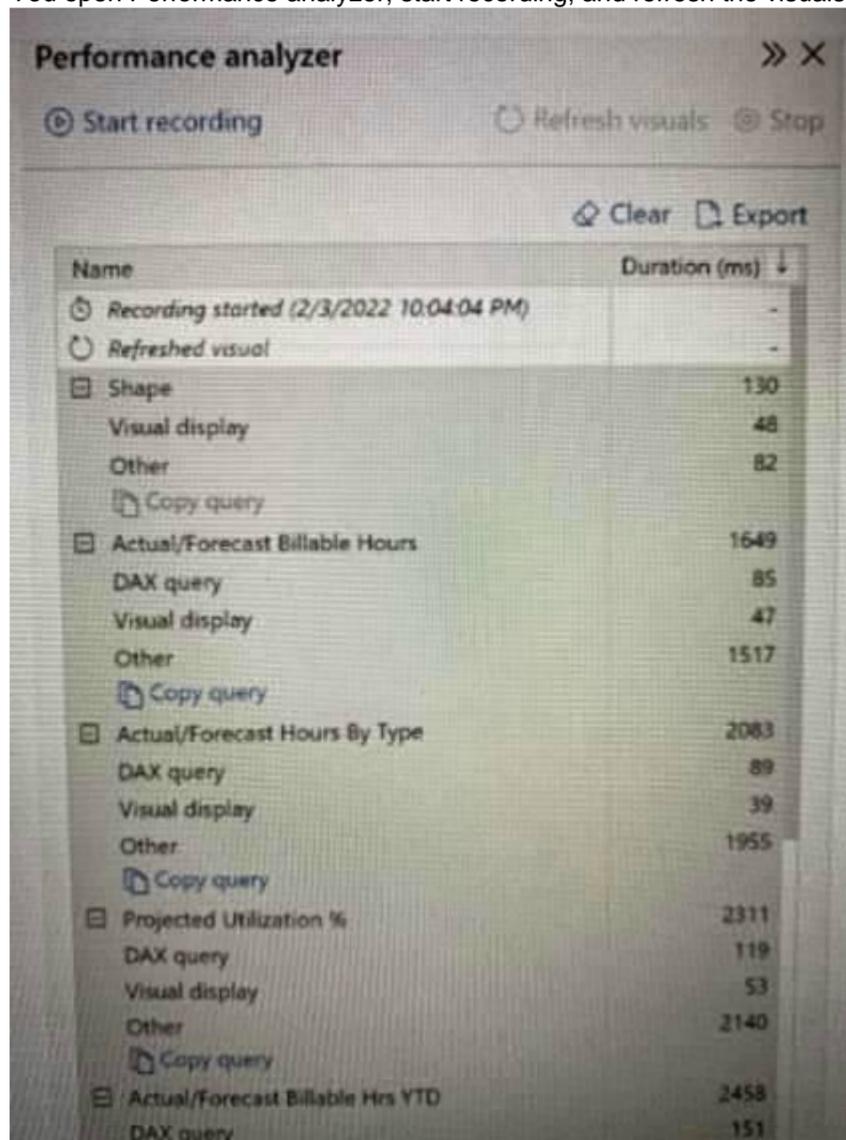
Reference: <https://radacad.com/performance-tip-for-power-bi-enable-load-sucks-memory-up> <https://docs.microsoft.com/en-us/powerquery-m/table-buffer>

NEW QUESTION 3

- (Exam Topic 3)

You open a Power BI Desktop report that contains an imported data model and a single report page.

You open Performance analyzer, start recording, and refresh the visuals on the page. The recording produces the results shown in the following exhibit



What can you identify from the results?

- A. The Actual/Forecast Hours by Type visual takes a long time to render on the report page when the data is cross-filtered.
- B. The Actual/Forecast Billable Hrs YTD visual displays the most data.
- C. Unoptimized DAX queries cause the page to load slowly.
- D. When all the visuals refresh simultaneously, the visuals spend most of the time waiting on other processes to finish.

Answer: D

Explanation:

Most time is spent in the category Other - time required by the visual for preparing queries, waiting for other visuals to complete, or performing other background processing.

Note: Each visual's log information includes the time spent (duration) to complete the following categories of tasks:
 DAX query - if a DAX query was required, this is the time between the visual sending the query, and for Analysis Services to return the results.
 Visual display - time required for the visual to draw on the screen, including time required to retrieve any web images or geocoding.
 Other - time required by the visual for preparing queries, waiting for other visuals to complete, or performing other background processing.
 Reference: <https://docs.microsoft.com/en-us/power-bi/create-reports/desktop-performance-analyzer>

NEW QUESTION 4

- (Exam Topic 3)

You have a Power Bi workspace named Workspacel in a Premium capacity. Workspacel contains a dataset. During a scheduled refresh, you receive the following error message: "Unable to save the changes since the new dataset size of 11,354 MB exceeds the limit of 10,240 MB."
 You need to ensure that you can refresh the dataset. What should you do?

- A. Turn on Large dataset storage format.
- B. Connect Workspace1 to an Azure Data Lake Storage Gen2 account
- C. Change License mode to Premium per user.
- D. Change the location of the Premium capacity.

Answer: D

Explanation:

Assigning workspaces to capacities

Workspaces can be assigned to a Premium capacity in the Power BI Admin portal or, for a workspace, in the Workspace pane.

Note: Capacity limits

Workspace storage limits, whether for My Workspace or an app workspace, depend on whether the workspace is in shared or Premium capacity.

* Shared capacity limits

For workspaces in shared capacity:

There is a per-workspace storage limit of 10 GB.

Premium Per User (PPU) tenants have a 100 TB storage limit.

When using a Pro license, the total usage can't exceed the tenant storage limit of 10 GB multiplied by the number of Pro licenses in the tenant.

* Premium capacity limits

For workspaces in Premium capacity:

There is a limit of 100 TB per Premium capacity. There is no per-user storage limit.

Workspace storage usage is shown as 0 (as shown in this screenshot) if the workspace is assigned to a Premium capacity.

Incorrect:

Not C: If your organization is using the original version of Power BI Premium, you're required to migrate to the modern Premium Gen2 platform. Microsoft began migrating all Premium capacities to Gen2.

Reference: <https://docs.microsoft.com/en-us/power-bi/enterprise/service-premium-capacity-manage-gen2> <https://docs.microsoft.com/en-us/power-bi/admin/service-admin-manage-your-data-storage-in-power-bi>

NEW QUESTION 5

- (Exam Topic 3)

You use the Vertipaq Analyzer to analyze tables in a dataset as shown in the Tables exhibit. (Click the Tables tab.)

Vertipaq Analyzer Metrics						
Tables						
Name	Cardinality	Table Size	Col Size	Data	Dictionary	Hier Size
Plan	627,876	22,823,464	21,147,552	6,697,272	10,293,184	4,157,096
Forecast Amount	101,606	22,823,464	7,400,920	1,475,640	5,112,384	812,896
Budget Amount	101,596	22,823,464	7,400,024	1,475,640	5,111,568	812,816
Row ID	627,876	22,823,464	4,185,992	1,674,344	120	2,511,528
ProductKey	628	22,823,464	842,296	818,016	19,208	5,072
Sales	858,789	20,968,092	18,674,660	12,182,384	2,587,004	3,905,272
Row ID	858,789	20,968,092	5,725,408	2,290,112	120	3,435,176
SalesAmount	36,554	20,968,092	2,960,560	1,245,904	1,422,176	292,480
TotalCost	9,711	20,968,092	1,924,272	1,238,488	608,056	77,728
Sales ID	2,000	20,968,092	1,431,192	1,374,064	41,080	16,048
Date	1,095	20,968,092	1,428,968	1,373,856	46,312	8,800

The table relationships for the dataset are shown in the Relationships exhibit. (Click the Relationships tab.)

Vertipaq Analyzer Metrics						
Relationships						
Table / Relationship	Size	Max From Cardinality	Max To Cardinality	1:M Ratio %	Missing Keys	
Plan	1,675,912	627,876	858,789	136.78%	7	
Plan[ProductKey] ∞--1 Product[ProductKey]	848	628	629	0.10%	0	
Plan[StoreKey] ∞--1 Store[Store Key]	360	306	299	0.05%	7	
Plan[GeographyKey] ∞--1 Geography[GeographyKey]	312	263	263	0.04%	0	
Plan[DateKey] ∞--1 Month & Year Distinct[Date]	32	36	36	0.01%	0	
Sales	2,293,432	858,789	1,095	0.13%	858,793	
Sales[Date] ∞--1 Calendar[Date]	1,760	1,095	1,095	0.13%	0	
Sales[GeographyKey] ∞--1 Geography[GeographyKey]	312	263	263	0.03%	0	
Sales[PromotionKey] ∞--1 Promotion[Promotion Key]	24	28	28	0.00%	0	
Sales[channelKey] ∞--1 Channel[ChannelKey]	8	4	4	0.00%	0	
Sales[Row ID] ∞--1 Plan Header Details[Row ID]	0	858,789	3	0.00%	858,786	

You need to reduce the model size by eliminating invalid relationships. Which column should you remove?

- A. Sales[Sales Amount]
- B. Sales[RowID]
- C. Sales[Sales ID]
- D. Plan[RowID]

Answer: B

Explanation:

Sales[Row ID] has 858,786 missing keys and 858,789 Max From Cardinality.

Note: The Max From Cardinality column defines the cost of the relationship which is the amount of time DAX needs to transfer the filters from the dimensions table to the fact table.

Reference: <https://blog.enterprisedna.co/vertipaq-analyzer-tutorial-relationships-referential-integrity/>

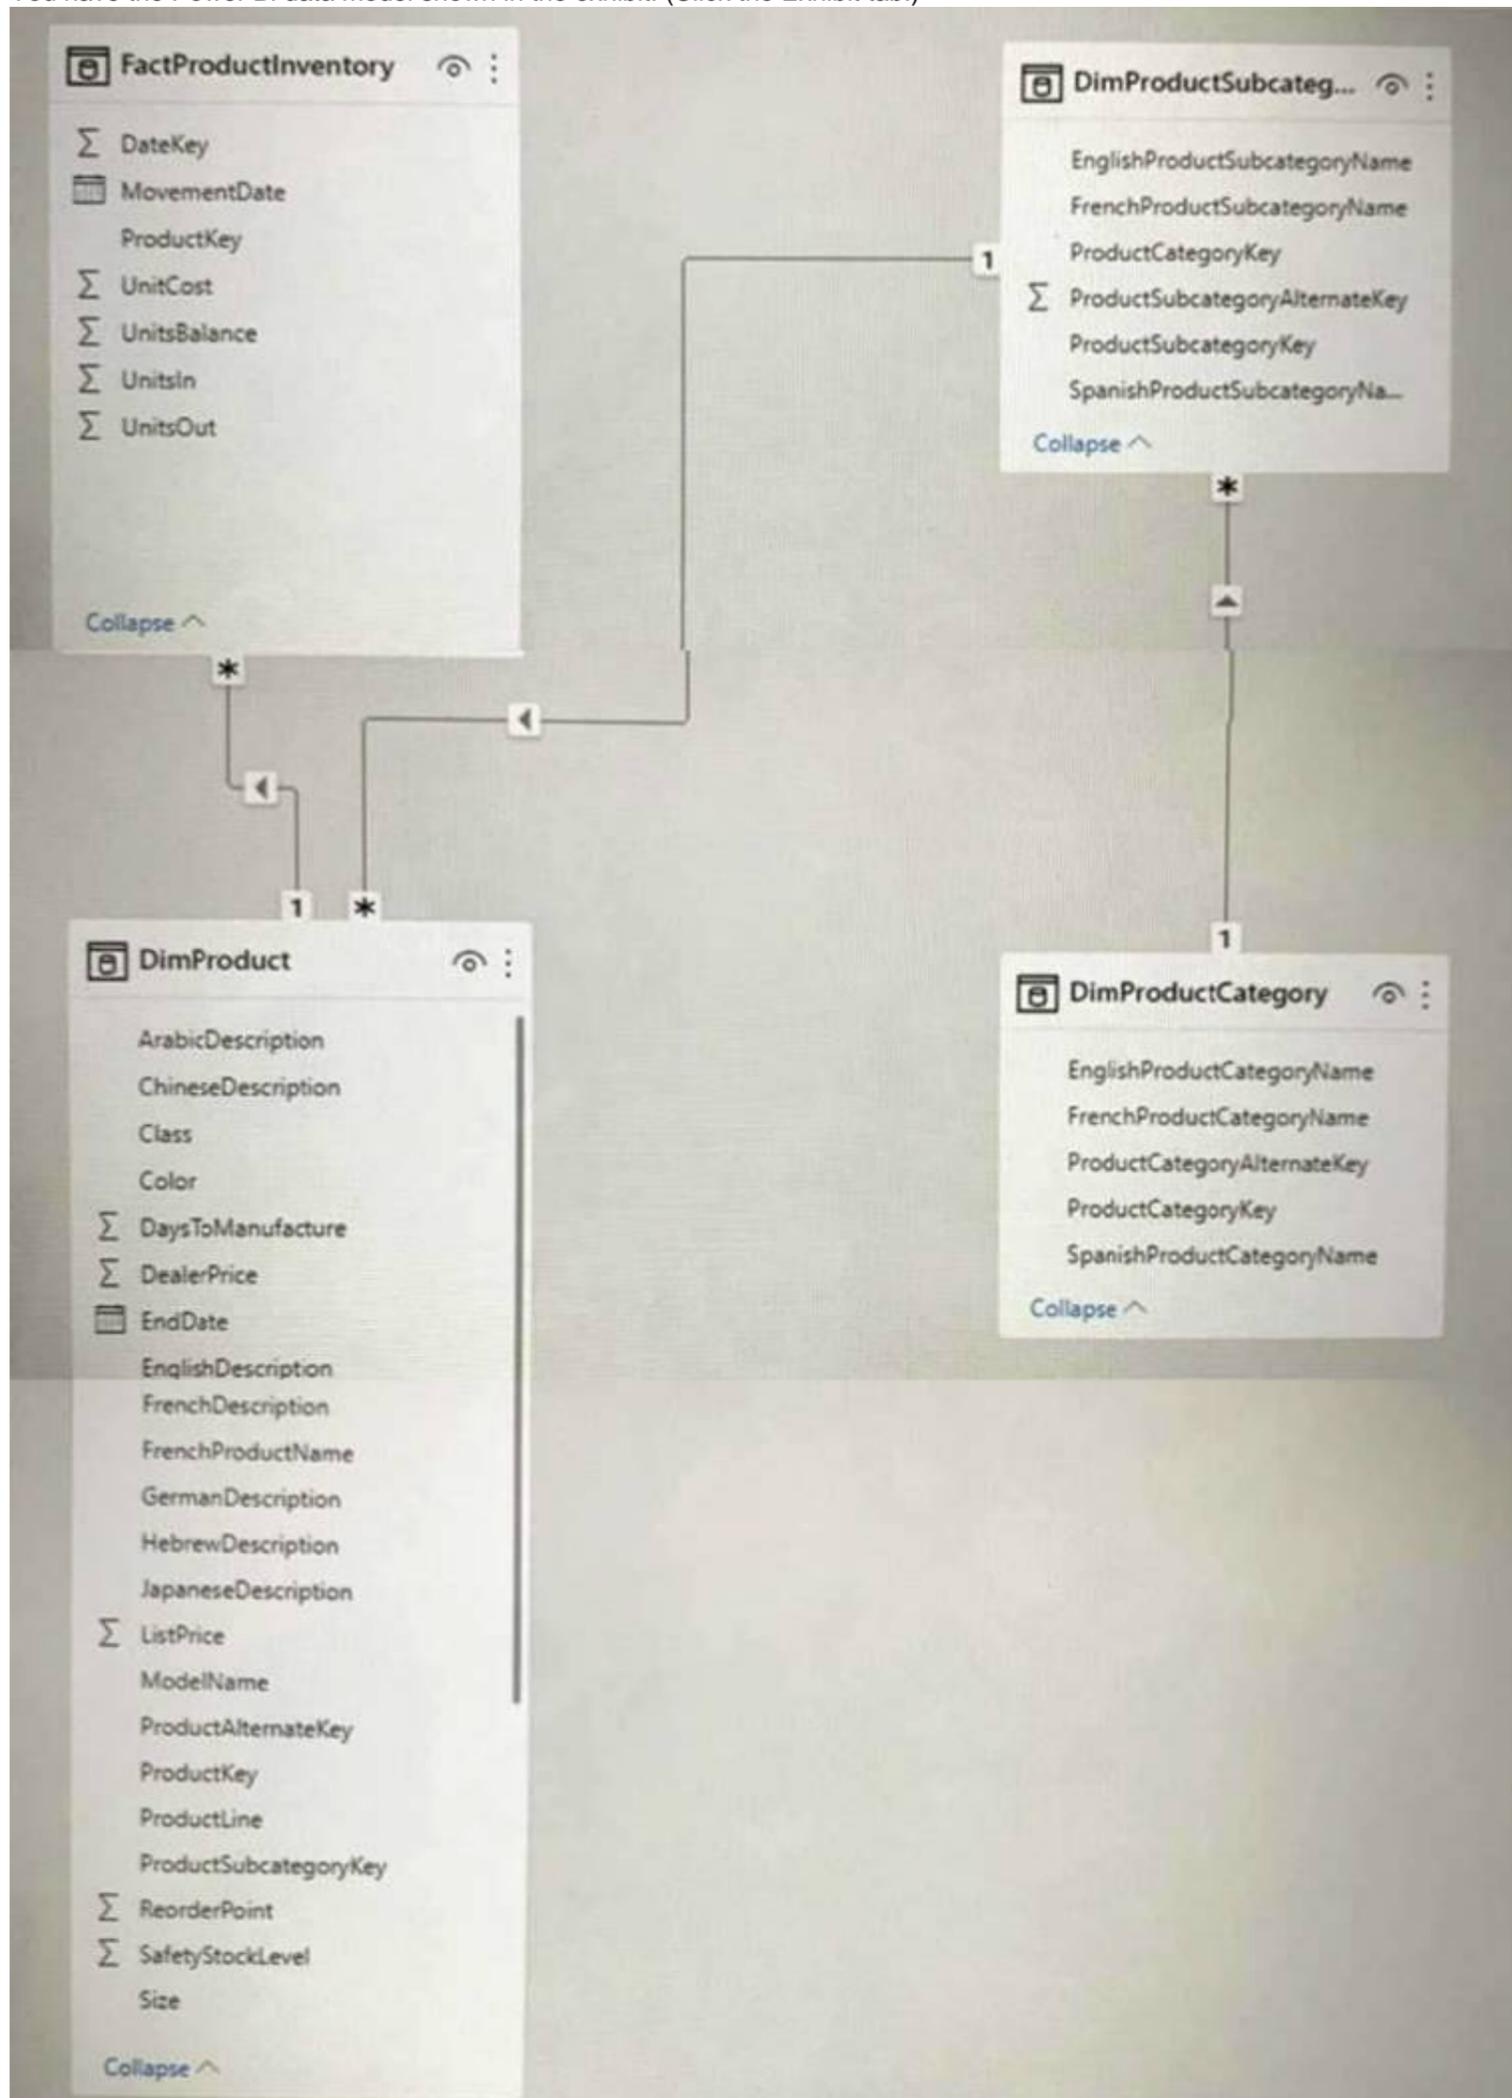
NEW QUESTION 6

- (Exam Topic 3)

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 the Power BI data model shown in the exhibit. (Click the Exhibit tab.)



Users indicate that when they build reports from the data model, the reports take a long time to load. You need to recommend a solution to reduce the load times of the reports.

Solution: You recommend normalizing the data model. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Instead denormalize For Performance.

Even though it might mean storing a bit of redundant data, schema denormalization can sometimes provide better query performance. The only question then becomes is the extra space used worth the performance benefit.

Reference: <https://www.mssqltips.com/sqlservertutorial/3211/denormalize-for-performance/>

NEW QUESTION 7

- (Exam Topic 3)

You have two Power BI reports named Report1 and Report2.

Report1 connects to a shared dataset named Dataset1.

Report2 connects to a local dataset that has the same structure as Dataset1. Report2 contains several calculated tables and parameters.

You need to prepare Report2 to use Dataset1.

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Remove the data source permissions.
- B. Delete all the Power Query Editor objects.
- C. Modify the source of each query.
- D. Update all the parameter values.
- E. Delete all the calculated tables.

Answer: CD

Explanation:

C: Power BI Desktop also comes with Power Query Editor. Use Power Query Editor to connect to one or many data sources, shape and transform the data to meet your needs, then load that model into Power BI Desktop.

D: Common uses for parameters

Here are some of the most common ways to use parameters. Control paginated report data

* Filter paginated report data at the data source by writing dataset queries that contain variables.

* Etc.

Reference: <https://docs.microsoft.com/en-us/power-bi/transform-model/desktop-query-overview> <https://docs.microsoft.com/en-us/learn/modules/dax-power-bi-add-calculated-tables/1-introduction>

NEW QUESTION 8

- (Exam Topic 3)

You have a Power BI workspace named Workspace1 that contains five dataflows.

You need to configure Workspace1 to store the dataflows in an Azure Data Lake Storage Gen2 account. What should you do first?

- A. Delete the dataflow queries.
- B. From the Power BI Admin portal, enable tenant-level storage.
- C. Disable load for all dataflow queries.
- D. Change the Data source settings in the dataflow queries.

Answer: B

Explanation:

Configuring Azure connections is an optional setting with additional properties that can optionally be set:

* Tenant Level storage, which lets you set a default, and/or

* Workspace-level storage, which lets you specify the connection per workspace

You can optionally configure tenant-level storage if you want to use a centralized data lake only, or want this to be the default option.

Reference:

<https://docs.microsoft.com/en-us/power-bi/transform-model/dataflows/dataflows-azure-data-lake-storage-integra>

NEW QUESTION 9

- (Exam Topic 3)

You are creating an external table by using an Apache Spark pool in Azure Synapse Analytics. The table will contain more than 20 million rows partitioned by date. The table will be shared with the SQL engines.

You need to minimize how long it takes for a serverless SQL pool to execute a query data against the table. In which file format should you recommend storing the table data?

- A. JSON
- B. Apache Parquet
- C. CSV
- D. Delta

Answer: B

Explanation:

Prepare files for querying

If possible, you can prepare files for better performance:

* Convert large CSV and JSON files to Parquet. Parquet is a columnar format. Because it's compressed, its file sizes are smaller than CSV or JSON files that contain the same data. Serverless SQL pool skips the columns and rows that aren't needed in a query if you're reading Parquet files. Serverless SQL pool needs less time and fewer storage requests to read it.

Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/best-practices-serverless-sql-pool> <https://stackoverflow.com/questions/65320949/parquet-vs-delta-format-in-azure-data-lake-gen-2-store>

NEW QUESTION 10

- (Exam Topic 3)

You need to save Power BI dataflows in an Azure Storage account.

Which two prerequisites are required to support the configuration? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. The storage account must be protected by using an Azure Firewall.
- B. The connection must be created by a user that is assigned the Storage Blob Data Owner role.
- C. The storage account must have hierarchical namespace enabled.
- D. Dataflows must exist already for any directly connected Power BI workspaces.
- E. The storage account must be created in a separate Azure region from the Power BI tenant and workspaces.

Answer: BC

Explanation:

Reference:

<https://docs.microsoft.com/en-us/power-bi/transform-model/dataflows/dataflows-azure-data-lake-storage-integra>

NEW QUESTION 10

- (Exam Topic 3)

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 Power BI dataset named Dataset1.

In Dataset1, you currently have 50 measures that use the same time intelligence logic. You need to reduce the number of measures, while maintaining the current functionality. Solution: From Power BI Desktop, you group the measures in a display folder.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Solution: From DAX Studio, you write a query that uses grouping sets.

A grouping is a set of discrete values that are used to group measure fields. Reference: <https://docs.microsoft.com/en-us/power-bi/developer/visuals/capabilities>

NEW QUESTION 13

- (Exam Topic 3)

You have five Power BI reports that contain R script data sources and R visuals.

You need to publish the reports to the Power BI service and configure a daily refresh of datasets. What should you include in the solution?

- A. a Power BI Embedded capacity
- B. an on-premises data gateway (standard mode)
- C. a workspace that connects to an Azure Data Lake Storage Gen2 account
- D. an on-premises data gateway (personal mode)

Answer: D

Explanation:

To schedule refresh of your R visuals or dataset, enable scheduled refresh and install an on-premises data gateway (personal mode) on the computer containing the workbook and R.

Reference: <https://docs.microsoft.com/en-us/power-bi/connect-data/desktop-r-in-query-editor>

NEW QUESTION 18

- (Exam Topic 3)

You have a deployment pipeline for a Power BI workspace. The workspace contains two datasets that use import storage mode.

A database administrator reports a drastic increase in the number of queries sent from the Power BI service to an Azure SQL database since the creation of the deployment pipeline.

An investigation into the issue identifies the following:

- One of the datasets is larger than 1 GB and has a fact table that contains more than 500 million rows.
- When publishing dataset changes to development, test, or production pipelines, a refresh is triggered against the entire dataset.

You need to recommend a solution to reduce the size of the queries sent to the database when the dataset changes are published to development, test, or production.

What should you recommend?

- A. Request the authors of the deployment pipeline datasets to reduce the number of datasets republished during development.
- B. In the dataset, delete the fact table.
- C. Configure the dataset to use a composite model that has a DirectQuery connection to the fact table.
- D. From Capacity settings in the Power BI Admin portal, reduce the Max Intermediate Row Set Count setting.

Answer: C

Explanation:

Previously in Power BI Desktop, when you used a DirectQuery in a report, no other data connections, whether DirectQuery or import, were allowed for that report. With composite models, that restriction is removed. A report can seamlessly include data connections from more than one DirectQuery or import data connection, in any combination you choose.

The composite models capability in Power BI Desktop consists of three related features:

* Composite models: Allows a report to have two or more data connections from different source groups, such as one or more DirectQuery connections and an

import connection, two or more DirectQuery connections, or any combination thereof.

* Etc.

Reference: <https://docs.microsoft.com/en-us/power-bi/transform-model/desktop-composite-models>

NEW QUESTION 22

- (Exam Topic 3)

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 are using an Azure Synapse Analytics serverless SQL pool to query a collection of Apache Parquet files by using automatic schema inference. The files contain more than 40 million rows of UTF-8-encoded business names, survey names, and participant counts. The database is configured to use the default collation.

The queries use open row set and infer the schema shown in the following table.

name	system_type_name	max_length
businessName	varchar(8000)	8000
surveyName	varchar(8000)	8000
participants	int	4

You need to recommend changes to the queries to reduce I/O reads and tempdb usage.

Solution: You recommend defining a data source and view for the Parquet files. You recommend updating the query to use the view.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Solution: You recommend using OPENROWSET WITH to explicitly specify the maximum length for businessName and surveyName.

The size of the varchar(8000) columns are too big. Better reduce their size.

A SELECT...FROM OPENROWSET(BULK...) statement queries the data in a file directly, without importing the data into a table. SELECT...FROM OPENROWSET(BULK...) statements can also list bulk-column aliases by using a format file to specify column names, and also data types.

Reference: <https://docs.microsoft.com/en-us/sql/t-sql/functions/openrowset-transact-sql>

NEW QUESTION 25

- (Exam Topic 3)

You use an Apache Spark notebook in Azure Synapse Analytics to filter and transform data. You need to review statistics for a DataFrame that includes:

The column name The column type

The number of distinct values

Whether the column has missing values Which function should you use?

- A. displayHTML()
- B. display(df, summary=true)
- C. %%configure
- D. display(df)
- E. %%!smagic

Answer: B

Explanation:

display(df) statistic details

You can use display(df, summary = true) to check the statistics summary of a given Apache Spark DataFrame that include the column name, column type, unique values, and missing values for each column. You can also select on specific column to see its minimum value, maximum value, mean value and standard deviation.

Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/spark/apache-spark-data-visualization>

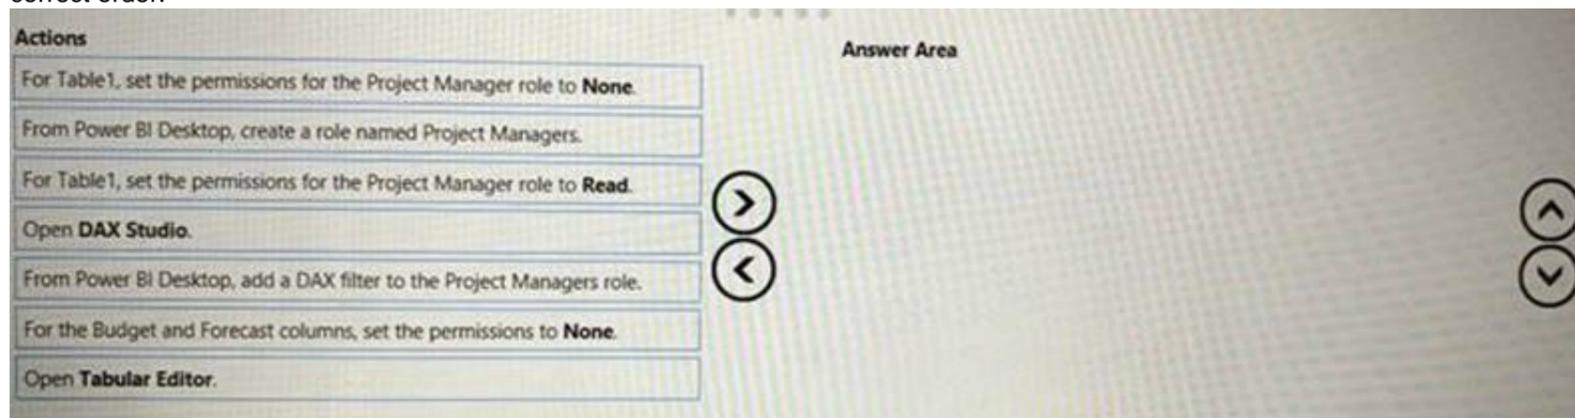
NEW QUESTION 30

- (Exam Topic 3)

You have a Power BI dataset that contains two tables named Table1 and Table2. The dataset is used by one report.

You need to prevent project managers from accessing the data in two columns in Table1 named Budget and Forecast.

Which four 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:

Step 1: From Power BI Desktop, create a role named Project Managers. Create roles

You can define roles within Power BI Desktop. Step 2: Open Tabular Editor

Under Tables, select the table to which you want to apply a DAX rule.

In the Table filter DAX expression box, enter the DAX expressions. This expression returns a value of true or false. For example: [Entity ID] = "Value".

Step 3: From Power BI Desktop, add a DAX filter to the Project Managers role. Step 4: For Table1, the Budget and Forecast columns, set the permissions to None.

Reference: <https://docs.microsoft.com/en-us/power-bi/guidance/rls-guidance>

NEW QUESTION 32

- (Exam Topic 3)

You have a Power BI dataset that has only the necessary fields visible for report development.

You need to ensure that end users see only 25 specific fields that they can use to personalize visuals. What should you do?

- A. From Tabular Editor, create a new role.
- B. Hide all the fields in the dataset.
- C. Configure object-level security (OLS).
- D. From Tabular Editor, create a new perspective.

Answer: B

NEW QUESTION 34

- (Exam Topic 3)

You are implementing a reporting solution that has the following requirements:

- Reports for external customers must support 500 concurrent requests. The data for these reports is approximately 7 GB and is stored in Azure Synapse Analytics.
 - Reports for the security team use data that must have local security rules applied at the database level to restrict access. The data being reviewed is 2 GB.
- Which storage mode provides the best response time for each group of users?

- A. DirectQuery for the external customers and import for the security team.
- B. DirectQuery for the external customers and DirectQuery for the security team.
- C. Import for the external customers and DirectQuery for the security team.
- D. Import for the external customers and import for the security team.

Answer: A

Explanation:

With DirectQuery, queries are sent back to your Azure Synapse Analytics in real time as you explore the data. Real-time queries, combined with the scale of Synapse Analytics enables users to create dynamic reports in minutes against terabytes of data.

Need import for the security team for local security rules. Reference:

<https://docs.microsoft.com/en-us/power-bi/connect-data/service-azure-sql-data-warehouse-with-direct-connect>

NEW QUESTION 36

- (Exam Topic 3)

You have a group of data scientists who must create machine learning models and run periodic experiments on a large dataset.

You need to recommend an Azure Synapse Analytics pool for the data scientists. The solution must minimize costs.

Which type of pool should you recommend?

- A. a Data Explorer pool
- B. an Apache Spark pool
- C. a dedicated SQL pool
- D. a serverless SQL pool

Answer: B

Explanation:

In Azure Synapse, training machine learning models can be performed on the Apache Spark Pools with tools like PySpark/Python, Scala, or .NET.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/machine-learning/what-is-machine-learning>

NEW QUESTION 41

- (Exam Topic 3)

You need to provide users with a reproducible method to connect to a data source and transform the data by using an AI function. The solution must meet the following requirement

- Minimize development effort.
- Avoid including data in the file. Which type of file should you create?

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

Answer: C

Explanation:

A PBIT file is a template created by Power BI Desktop, a Microsoft application used to create reports and visualizations. It contains queries, visualization settings, data models, reports, and other data added by the user.

A PBIT file acts as a Power BI template. It doesn't include any data from your source systems. Reference: <https://docs.microsoft.com/en-us/power-bi/connect-data/desktop-data-sources>

NEW QUESTION 46

- (Exam Topic 3)

You have a Power BI dataset that contains the following measure.

```
YTD Year-over-Year Var =
DIVIDE (
    (
        [Sales Amount]
        - CALCULATE (
            [Sales],
            SAMEPERIODLASTYEAR ( 'Calendar'[Date] ),
            'Calendar'[Flag] = "YTD"
        )
    ),
    CALCULATE (
        [Sales],
        SAMEPERIODLASTYEAR ( 'Calendar'[Date] ),
        'Calendar'[Flag] = "YTD"
    ),
    BLANK()
)
```

You need to improve the performance of the measure without affecting the logic or the results. What should you do?

- A. Replace both calculate functions by using a variable that contains the calculate function.
- B. Remove the alternative result of blank() from the divide function.
- C. Create a variable and replace the values for [sales Amount].
- D. Remove "calendar'[Flag] = "YTD" from the code.

Answer: A

NEW QUESTION 48

- (Exam Topic 3)

You have the following Python code in an Apache Spark notebook.

```
import matplotlib.pyplot as plt
import numpy as np
ys = 300 + np.random.randn(100)
x = [x for x in range(len(ys))]
plt.plot(x, ys, '-')
plt.fill_between(x, ys, 395, where=(ys > 395), facecolor='g', alpha=0.5)
plt.title("Chart Sample")
plt.show()
```

Which type of chart will the code produce?

- A. a stacked bar chart
- B. a pie chart
- C. a bar chart
- D. an area chart

Answer: D

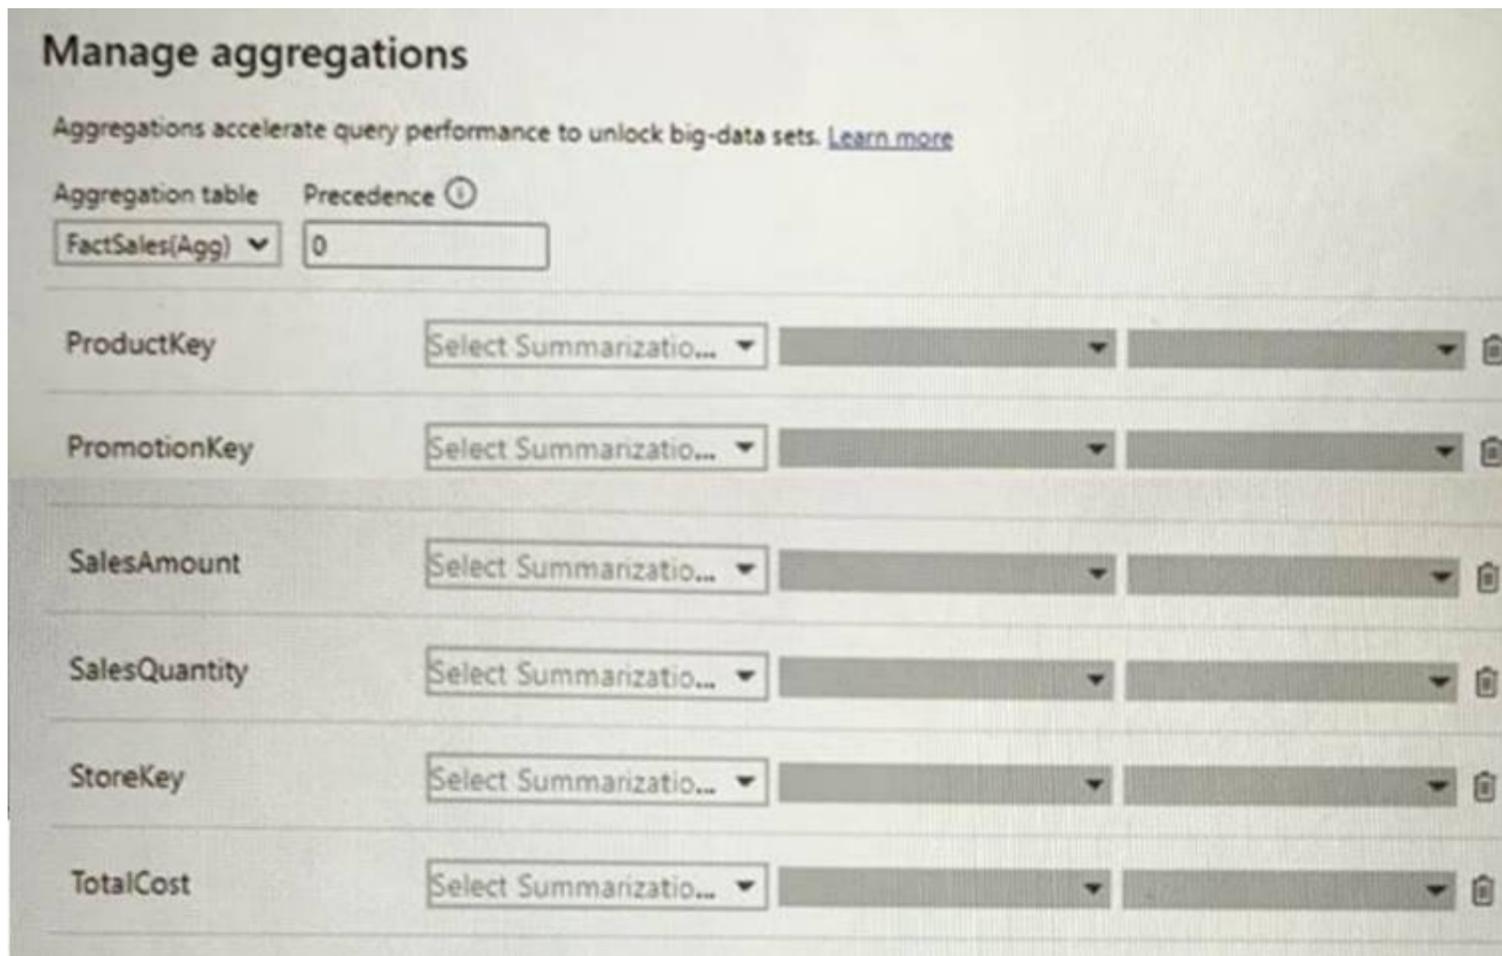
Explanation:

The matplotlib.pyplot.fill_between function fills the area between two horizontal curves. The curves are defined by the points (x, y1) and (x, y2). This creates one or multiple polygons describing the filled area. Reference: https://matplotlib.org/3.5.0/api/_as_gen/matplotlib.pyplot.fill_between.html

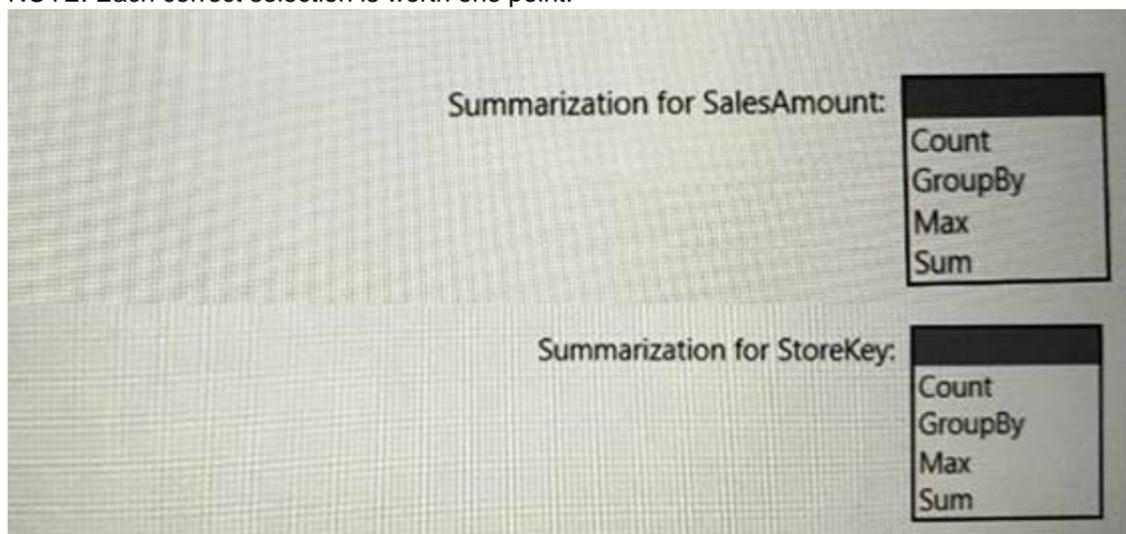
NEW QUESTION 51

- (Exam Topic 3)

You are configuring an aggregation table as shown in the following exhibit.



The detail table is named FactSales and the aggregation table is named FactSales(Agg). You need to aggregate SalesAmount for each store. Which type of summarization should you use for SalesAmount and StoreKey? To answer, select the appropriate options in the answer area, NOTE: Each correct selection is worth one point.



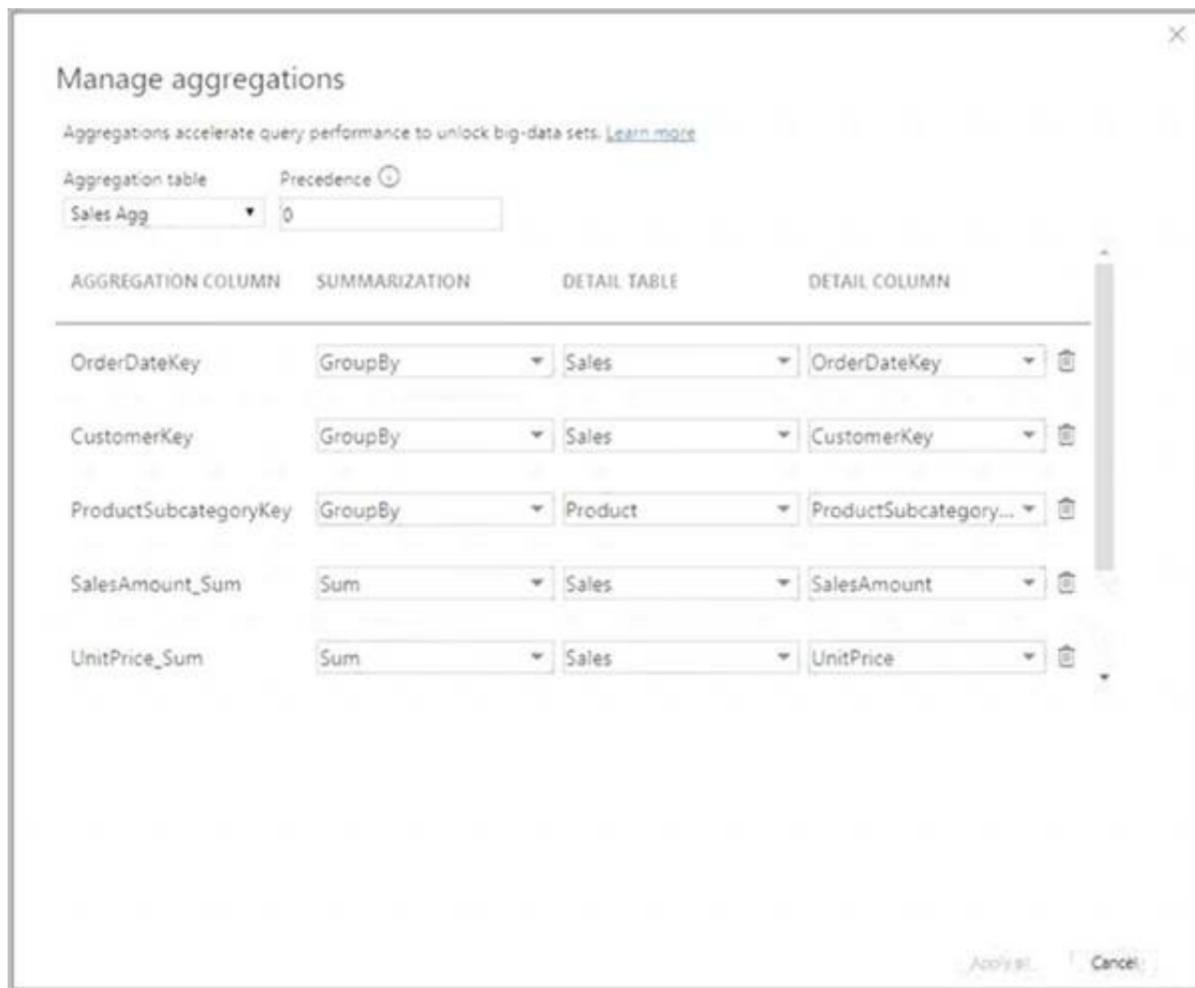
- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Sum

The Manage aggregations dialog shows a row for each column in the table, where you can specify the aggregation behavior. In the following example, queries to the Sales detail table are internally redirected to the Sales Agg aggregation table.



Box 2: GroupBy

Reference: <https://docs.microsoft.com/en-us/power-bi/transform-model/aggregations-advanced>

NEW QUESTION 54

- (Exam Topic 3)

You are optimizing a dataflow in a Power BI Premium capacity. The dataflow performs multiple joins. You need to reduce the load time of the dataflow. Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Reduce the memory assigned to the dataflows.
- B. Execute non-foldable operations before foldable operations.
- C. Execute foldable operations before non-foldable operations.
- D. Place the ingestion operations and transformation operations in a single dataflow.
- E. Place the ingestion operations and transformation operations in separate dataflows.

Answer: CE

Explanation:

Using the compute engine to improve performance

Take the following steps to enable workloads trigger the compute engine, and always improve performance: For computed and linked entities in the same workspace:

Ensure you perform the operations that fold, such as merges, joins, conversion, and others.

For ingestion focus on getting the data into the storage as fast as possible, using filters only if they reduce the overall dataset size. It's best practice to keep your transformation logic separate from this step, and allow the engine to focus on the initial gathering of ingredients. Next, separate your transformation and business logic into a separate dataflow in the same workspace, using linked or computed entities; doing so allows for the engine to activate and accelerate your computations. In our analogy, it's like food preparation in the kitchen: food preparation is typically a separate and distinct step from gathering your raw ingredients, and a pre-requisite for putting the food in the oven. Similarly, your logic needs to be prepared separately before it can take advantage of the compute engine.

Reference:

<https://docs.microsoft.com/en-us/power-bi/transform-model/dataflows/dataflows-premium-workload-configurati>

NEW QUESTION 57

- (Exam Topic 3)

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 Power BI dataset named Dataset1.

In Dataset1, you currently have 50 measures that use the same time intelligence logic. You need to reduce the number of measures, while maintaining the current functionality. Solution: From DAX Studio, you write a query that uses grouping sets.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

A grouping is a set of discrete values that are used to group measure fields. Reference: <https://docs.microsoft.com/en-us/power-bi/developer/visuals/capabilities>

NEW QUESTION 60

- (Exam Topic 3)

You have an Azure Synapse Analytics dataset that contains data about jet engine performance. You need to score the dataset to identify the likelihood of an engine failure. Which function should you use in the query?

- A. PIVOT
- B. GROUPING
- C. PREDICT
- D. CAST

Answer: A

NEW QUESTION 63

- (Exam Topic 3)

You have the following code in an Azure Synapse notebook.

```
import matplotlib.pyplot as plt
x1 = [1, 3, 4, 5, 6, 7, 9]
y1 = [4, 7, 2, 4, 7, 8, 3]
x2 = [2, 4, 6, 8, 10]
y2 = [5, 6, 2, 6, 2]
plt.bar(x1, y1, label="Blue Item", color='b')
plt.bar(x2, y2, label="Green Item", color='g')
plt.plot()
plt.xlabel("Number")
plt.ylabel("Height")
plt.title("My Chart")
plt.legend()
plt.show()
```

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

Answer Area

Running the code will create a [answer choice] in the output cell.

- clustered bar chart
- histogram
- line chart
- stacked bar chart

The legend for the resulting chart will list [answer choice] in the legend.

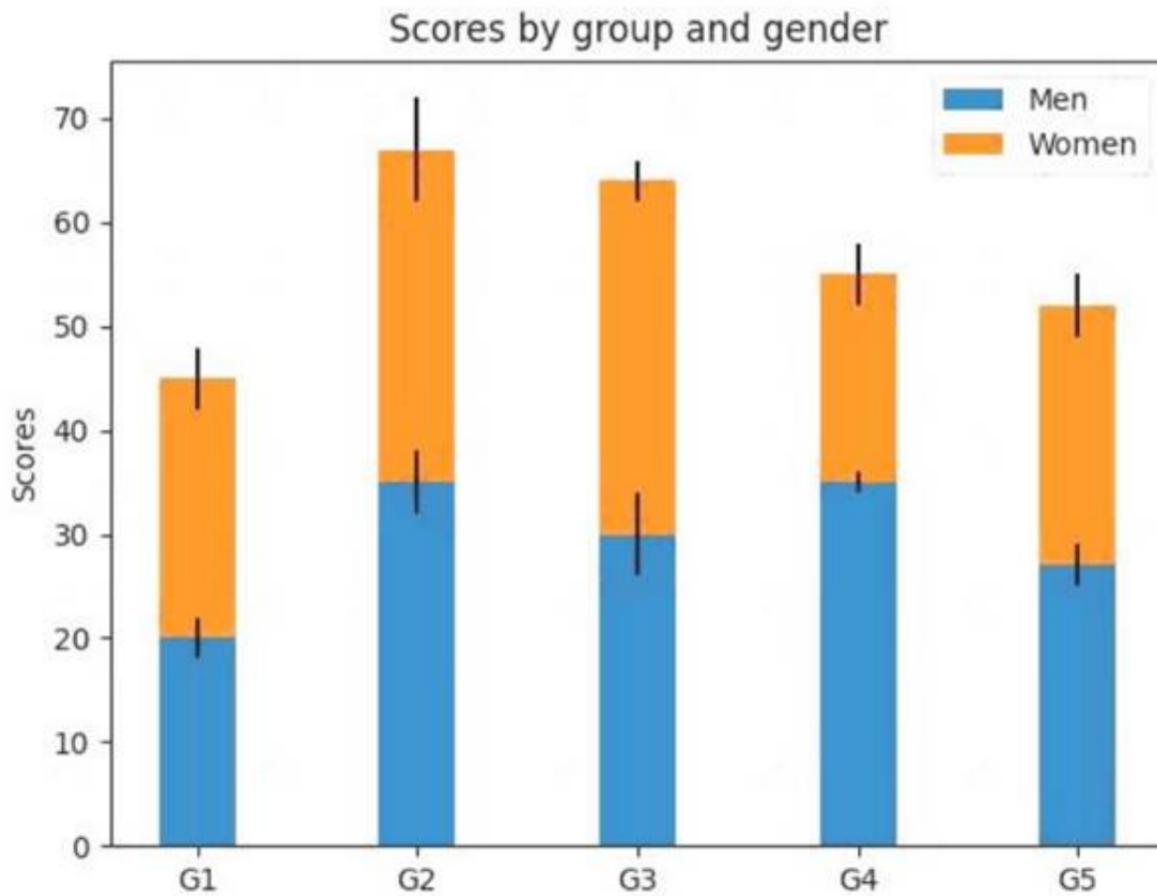
- one item
- two items
- five items
- seven items

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: stacked bar chart matplotlib.pyplot.bar makes a bar plot. The bars are positioned at x with the given alignment. Their dimensions are given by height and width. The vertical baseline is bottom (default 0). Many parameters can take either a single value applying to all bars or a sequence of values, one for each bar. Stacked bars can be achieved by passing individual bottom values per bar. Stacked bar chart This is an example of creating a stacked bar plot with error bars using bar. Note the parameters year used for error bars, and bottom to stack the women's bars on top of the men's bars.



```
import matplotlib.pyplot as plt
labels = ['G1', 'G2', 'G3', 'G4', 'G5']
men_means = [20, 35, 30, 35, 27]
women_means = [25, 32, 34, 20, 25]
men_std = [2, 3, 4, 1, 2]
women_std = [3, 5, 2, 3, 3]
width = 0.35 # the width of the bars: can also be len(x) sequence fig, ax = plt.subplots()
ax.bar(labels, men_means, width, yerr=men_std, label='Men')
ax.bar(labels, women_means, width, yerr=women_std, bottom=men_means, label='Women')
ax.set_ylabel('Scores') ax.set_title('Scores by group and gender') ax.legend()
plt.show()
```

Box 2: two items

Blue item and Green Item. matplotlib.legend

The legend module defines the Legend class, which is responsible for drawing legends associated with axes and/or figures.

Note: A Diagram Legend is an element that you can add to your diagram to provide information about the colors and/or line thicknesses and styles that have been used in the current diagram, where those colors and other styles have some particular meaning.

Reference: https://matplotlib.org/stable/api/_as_gen/matplotlib.pyplot.bar.html https://matplotlib.org/stable/gallery/lines_bars_and_markers/bar_stacked.html https://matplotlib.org/stable/api/legend_api.html

NEW QUESTION 64

- (Exam Topic 3)

You plan to create a Power BI report that will use an OData feed as the data source. You will retrieve all the entities from two different collections by using the same service root

The OData feed is still in development. The location of the feed will change once development is complete. The report will be published before the OData feed development is complete.

You need to minimize development effort to change the data source once the location changes.

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

- Create a parameter that contains the service root URI.
- From Advanced Editor, duplicate the query and change the resource path in the URL.
- Get data from an OData feed source and use the parameter to populate the first part of the URL.
- From Advanced Editor, get data from an OData feed source and use the parameter to populate the last part of the URL.
- From Advanced Editor, reference the query and change the resource path in the URL.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Step 1: Create a parameter that contains the service root URI

Step 2: Get data from OData feed source and use the parameter to populate the first part of the URL. The URI is in the first part of the query.

Example: let

Source = OData.Feed

```
("https://analytics.dev.azure.com/{organization}/{project}/_odata/v3.0-preview/WorkItemSnapshot? "
&"$apply=filter( "
&"WorkItemType eq 'Bug' "
&"and StateCategory ne 'Completed' "
&"and startswith(Area/AreaPath, '{areapath}') "
&"and DateValue ge {startdate} "
&") "
&"/groupby( "
&"(DateValue,State,WorkItemType,Priority,Severity,Area/AreaPath,Iteration/IterationPath,AreaSK), "
&"aggregate($count as Count) "
&") "
,null, [Implementation="2.0",OmitValues = ODataOmitValues.Nulls,ODataVersion = 4]) in
Source
```

Box 3: From Advanced Editor, duplicate the query and change the resource path in the URL. Choose Get Data, and then Blank Query. From the Power BI Query editor, choose Advanced Editor. The Advanced Editor window opens. Edit the query. Etc.
 Reference: <https://docs.microsoft.com/en-us/azure/devops/report/powerbi/odataquery-connect>

NEW QUESTION 66

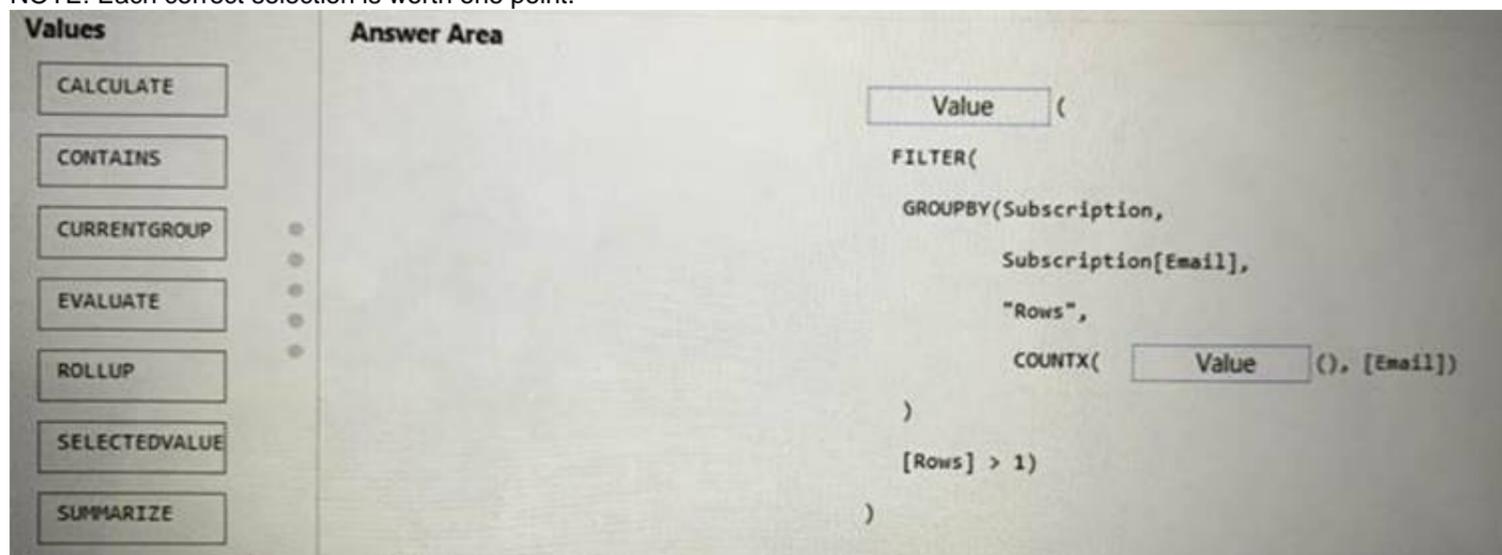
- (Exam Topic 3)

You are using DAX Studio to query an XMLA endpoint.

You need to identify the duplicate values in a column named Email in a table named Subscription.

How should you complete the DAX expression? To answer, drag the appropriate values to the targets. Each value may be used once, more than once. 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:

Box 1: CALCULATE

Box 2: CURRENTGROUP

CURRENTGROUP returns a set of rows from the table argument of a GROUPBY expression that belong to the current row of the GROUPBY result.

Remarks

This function can only be used within a GROUPBY expression.

This function takes no arguments and is only supported as the first argument to one of the following aggregation functions: AVERAGEX, COUNTAX, COUNTX, GEOMEANX, MAXX, MINX, PRODUCTX, STDEVX.S, STDEVX.P, SUMX, VARX.S, VARX.P.

Note: COUNTX counts the number of rows that contain a non-blank value or an expression that evaluates to a non-blank value, when evaluating an expression over a table.

Reference: <https://docs.microsoft.com/en-us/dax/currentgroup-function-dax>

NEW QUESTION 67

- (Exam Topic 3)

You are using an Azure Synapse Analytics serverless SQL pool to query network traffic logs in the Apache Parquet format. A sample of the data is shown in the following table.

source		destination	
name	ip	name	ip
Network01	192.168.0.1	Internet	0.0.0.0

You need to create a Transact-SQL query that will return the source IP address.

Which function should you use in the select statement to retrieve the source IP address?

- A. JSON_VALUE
- B. FOR.JSON
- C. CONVERT
- D. FIRST VALUE

Answer: A

NEW QUESTION 70

- (Exam Topic 3)

You have a Power BI dataset that uses DirectQuery against an Azure SQL database.

Multiple reports use the dataset.

A database administrator reports that too many queries are being sent from Power BI to the database. You need to reduce the number of queries sent to the database. The solution must meet the following requirements:

- DirectQuery must continue to be used.
- Visual interactions in all the reports must remain as they are configured currently.
- Consumers of the reports must only be allowed to apply filters from the Filter pane. Which two settings should you select? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Disabling cross highlighting/filtering by default
- B. Add a single Apply button to the filter pane to apply changes at once
- C. Add an Apply button to each slicer to apply changes when you're ready
- D. Add Apply buttons to all basic filters to apply changes when you're ready
- E. Ignore the Privacy Levels and potentially improve performance

Answer: BC

Explanation:

Reduce queries

Reduce the number of queries sent by Power BI using the Query reduction settings. For slicers, select the "Add an Apply button to each slicer to apply changes when you're ready" option. For filters, select "Add a single Apply button to the filter pane to apply changes at once (preview)."

Reference: <https://maqsoftware.com/insights/power-bi-best-practices>

NEW QUESTION 72

- (Exam Topic 3)

You have an Azure Data Lake Storage Gen 2 container that stores more than 300,000 files representing hourly telemetry data. The data is organized in folders by the year, month, and day according to when the telemetry was captured.

You have the following query in Power Query Editor.

```
let
    Source = AzureStorage.Blobs("https://tmppbie01.blob.core.windows.net/logs/"),
    Filtered = Table.SelectRows(Source, each Text.StartsWith([Name], "2019/12/"),
        and [Extension] = ".csv"),
    Transformed = Table.AddColumn(Filtered, "Transformed", each TransformFiles([Content])),
    Limited = Table.SelectColumns(Transformed, "Transformed"),
    Expanded = Table.ExpandTableColumn(Limited, "Transformed", {"Date", "Name", "Activity"}),
    Final = Table.TransformColumnTypes(Expanded,
        {"Date", type date}, {"Name", type text}, {"Activity", type text})
in
    Final
```

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

Statements	Yes	No
The query uses the hierarchical namespace of the storage account.	<input type="radio"/>	<input type="radio"/>
The query uses a custom function to load file data.	<input type="radio"/>	<input type="radio"/>
Changing the source to use AzureStorage.DataLake will reduce the load time of the query.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

A key mechanism that allows Azure Data Lake Storage Gen2 to provide file system performance at object storage scale and prices is the addition of a hierarchical namespace. This allows the collection of objects/files within an account to be organized into a hierarchy of directories and nested subdirectories in the same way that the file system on your computer is organized. With a hierarchical namespace enabled, a storage account becomes capable of providing the scalability and cost-effectiveness of object storage, with file system semantics that are familiar to analytics engines and frameworks.

Box 2: No

Table.SelectRows returns a table of rows from the table, that matches the selection condition. Box 3: Yes

Azure Data Lake Storage has higher throughput and IOPS.

Note: Azure Blob Storage is a general purpose, scalable object store that is designed for a wide variety of storage scenarios. Azure Data Lake Storage is a hyper-scale repository that is optimized for big data analytics workloads.

Azure Data Lake Storage use Cases: Batch, interactive, streaming analytics and machine learning data such as log files, IoT data, click streams, large datasets

Reference: <https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-namespace> <https://docs.microsoft.com/en-us/powerquery-m/table-selectrows>

<https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-comparison-with-blob-storage>

NEW QUESTION 77

- (Exam Topic 3)

You have a deployment pipeline for a Power BI workspace. The workspace contains two datasets that use import storage mode.

A database administrator reports a drastic increase in the number of queries sent from the Power BI service to an Azure SQL database since the creation of the deployment pipeline.

An investigation into the issue identifies the following:

One of the datasets is larger than 1 GB and has a fact table that contains more than 500 million rows.

When publishing dataset changes to development, test, or production pipelines, a refresh is triggered against the entire dataset.

You need to recommend a solution to reduce the size of the queries sent to the database when the dataset changes are published to development, test, or production.

What should you recommend?

- A. Turn off auto refresh when publishing the dataset changes to the Power BI service.
- B. In the datase
- C. change the fact table from an import table to a hybrid table.
- D. Enable the large dataset storage format for workspace.
- E. Create a dataset parameter to reduce the fact table row count in the development and test pipelines.

Answer: B

Explanation:

Hybrid tables

Hybrid tables are tables with incremental refresh that can have both import and direct query partitions. During a clean deployment, both the refresh policy and the hybrid table partitions are copied. When deploying to a pipeline stage that already has hybrid table partitions, only the refresh policy is copied. To update the partitions, refresh the table.

Refreshes are faster - Only the most recent data that has changed needs to be refreshed.

Reference: <https://docs.microsoft.com/en-us/power-bi/create-reports/deployment-pipelines-best-practices>

NEW QUESTION 82

- (Exam Topic 2)

You need to build a Transact-SQL query to implement the planned changes for the internal users.

How should you complete the Transact-SQL query? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```

DECLARE @model varbinary(max) = (
    SELECT native_model_object
    FROM ml_models
    WHERE model_name = 'rxLinMod'
    AND model_version = 'v1');
SELECT d.*, p.*
FROM [ ] (MODEL = @model, DATA = dbo.rx_linMod as lm)
    [ ]
    [ ]
    [ ]
    [ ]
go [ ] (model_outcome float, trade_volume float, price_Pred float) as p;
    [ ]
    [ ]
    [ ]
    [ ]
    [ ]
    
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: PREDICT

Provide internal users with the ability to incorporate machine learning models loaded to the dedicated SQL pool.

The example below shows a sample query using prediction function. An additional column with name Score and data type float is created containing the prediction results. All the input data columns as well as output prediction columns are available to display with the select statement.

-- Query for ML predictions SELECT d.*, p.Score

FROM PREDICT(MODEL = (SELECT Model FROM Models WHERE Id = 1),

DATA = dbo.mytable AS d, RUNTIME = ONNX) WITH (Score float) AS p; Box 2: WITH

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-predict>

NEW QUESTION 87

- (Exam Topic 2)

You need to identify the root cause of the data refresh issue. What should you use?

- A. the Usage Metrics Report in powerbi.com
- B. Query Diagnostics in Power Query Editor

C. Performance analyzer in Power BI Desktop

Answer: B

Explanation:

Users indicate that the data in Power BI reports is stale. You discover that the refresh process of the Power BI model occasionally times out. With Query Diagnostics, you can achieve a better understanding of what Power Query is doing at authoring and at refresh time in Power BI Desktop. While we'll be expanding on this feature in the future, including adding the ability to use it during full refreshes, at this time you can use it to understand what sort of queries you're emitting, what slowdowns you might run into during authoring refresh, and what kind of background events are happening.
 Reference: <https://docs.microsoft.com/en-us/power-query/querydiagnostics>

NEW QUESTION 88

- (Exam Topic 2)

You need to recommend a solution to add new fields to the financial data Power BI dataset with data from the Microsoft SQL Server data warehouse. What should you include in the recommendation?

- A. Azure Purview
- B. Site-to-Site VPN
- C. an XMLA endpoint
- D. the on-premises data gateway

Answer: D

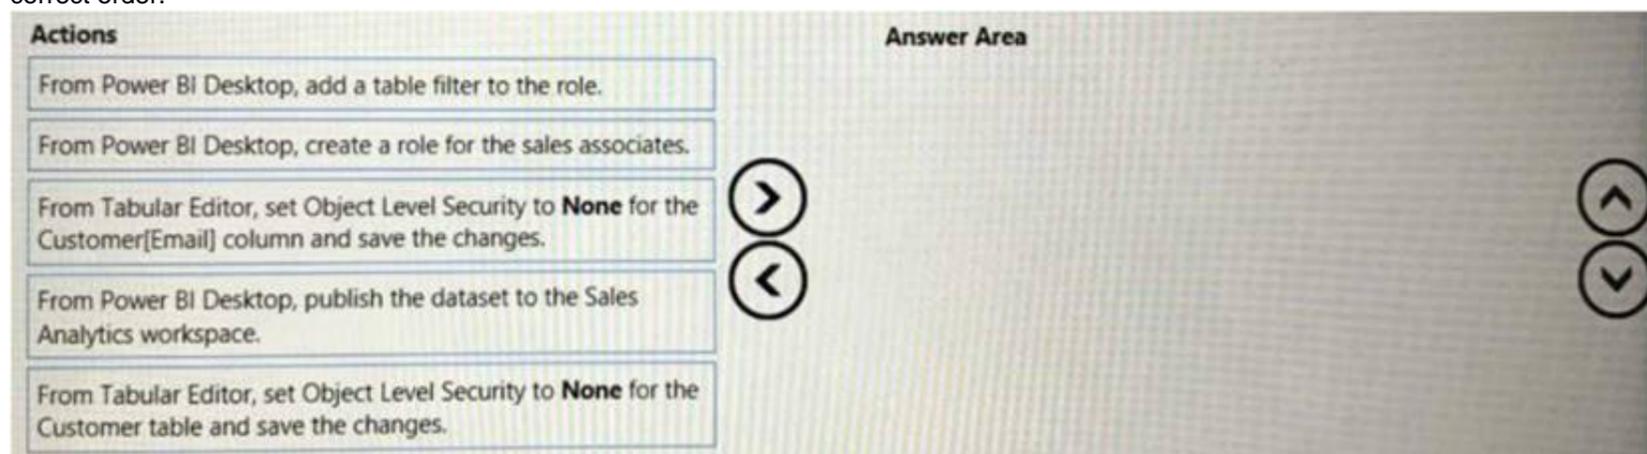
Explanation:

Refresh data from an on-premises SQL Server database
 The SQL Server database must be accessed by Power BI through an on-premises data gateway.
 You can install an on-premises data gateway on the same local computer as SQL Server (in production, it would typically be a different computer).
 Reference: <https://docs.microsoft.com/en-us/power-bi/connect-data/service-gateway-sql-tutorial>

NEW QUESTION 91

- (Exam Topic 1)

You need to implement object-level security (OLS) in the Power BI dataset for the sales associates. 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

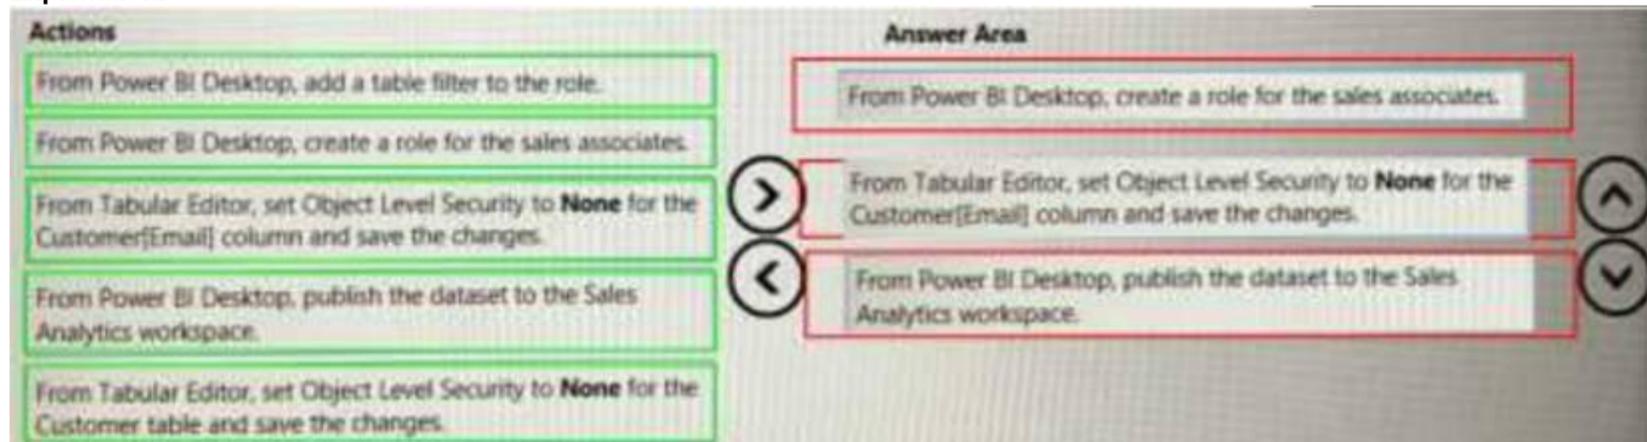
- From Power BI Desktop, add a table filter to the role.
- From Power BI Desktop, create a role for the sales associates.
- From Tabular Editor, set Object Level Security to **None** for the Customer[Email] column and save the changes.
- From Power BI Desktop, publish the dataset to the Sales Analytics workspace.
- From Tabular Editor, set Object Level Security to **None** for the Customer table and save the changes.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



Actions

- From Power BI Desktop, add a table filter to the role.
- From Power BI Desktop, create a role for the sales associates.
- From Tabular Editor, set Object Level Security to **None** for the Customer[Email] column and save the changes.
- From Power BI Desktop, publish the dataset to the Sales Analytics workspace.
- From Tabular Editor, set Object Level Security to **None** for the Customer table and save the changes.

Answer Area

- From Power BI Desktop, create a role for the sales associates.
- From Tabular Editor, set Object Level Security to **None** for the Customer[Email] column and save the changes.
- From Power BI Desktop, publish the dataset to the Sales Analytics workspace.

NEW QUESTION 93

- (Exam Topic 1)

You need to configure the Sales Analytics workspace to meet the ad hoc reporting requirements. What should you do?

- A. Grant the sales managers the Build permission for the existing Power BI datasets.
- B. Grant the sales managers admin access to the existing Power BI workspace.
- C. Create a deployment pipeline and grant the sales managers access to the pipeline.
- D. Create a PBIT file and distribute the file to the sales managers.

Answer: D

Explanation:

Allow sales managers to perform ad hoc sales reporting with minimal effort

Power BI report templates contain the following information from the report from which they were generated: Report pages, visuals, and other visual elements

The data model definition, including the schema, relationships, measures, and other model definition items All query definitions, such as queries, Query

Parameters, and other query elements

What is not included in templates is the report's data.

Report templates use the file extension .PBIT (compare to Power BI Desktop reports, which use the .PBIX extension).

Note: With Power BI Desktop, you can create compelling reports that share insights across your entire organization. With Power BI Desktop templates, you can streamline your work by creating a report template, based on an existing template, which you or other users in your organization can use as a starting point for a new report's layout, data model, and queries. Templates in Power BI Desktop help you jump-start and standardize report creation.

Reference: <https://docs.microsoft.com/en-us/power-bi/create-reports/desktop-templates>

NEW QUESTION 98

.....

Relate Links

100% Pass Your DP-500 Exam with ExamBible Prep Materials

<https://www.exambible.com/DP-500-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>