



Microsoft

Exam Questions DP-500

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

NEW QUESTION 1

- (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

Name	Duration (ms)
Recording started (2/3/2022 10:04:04 PM)	-
Refreshed visual	-
Shape	130
Visual display	48
Other	82
Copy query	
Actual/Forecast Billable Hours	1649
DAX query	85
Visual display	47
Other	1517
Copy query	
Actual/Forecast Hours By Type	2083
DAX query	89
Visual display	39
Other	1955
Copy query	
Projected Utilization %	2311
DAX query	119
Visual display	53
Other	2140
Copy query	
Actual/Forecast Billable Hrs YTD	2458
DAX query	151

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 2

- (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	Columns	Relationships	Partitions	Summary		
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.)

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.

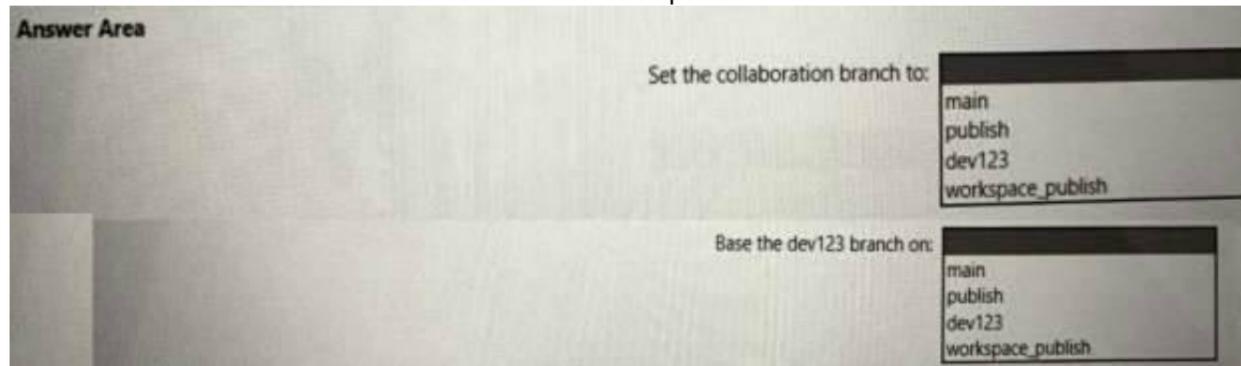
Reference: <https://blog.enterprisedna.co/vertipaq-analyzer-tutorial-relationships-referential-integrity/>

NEW QUESTION 3

- (Exam Topic 3)

You need to configure a source control solution for Azure Synapse Analytics. The solution must meet the following requirements:

- Code must always be merged to the main branch before being published, and the main branch must be used for publishing resource
- The workspace templates must be stored in the publish branch.
- A branch named dev123 will be created to support the development of a new feature. What should you do? 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: main

Code must always be merged to the main branch before being published, and the main branch must be used for publishing resources.

Collaboration branch - Your Azure Repos collaboration branch that is used for publishing. By default, its master. Change this setting in case you want to publish resources from another branch. You can select existing branches or create new.

Each Git repository that's associated with a Synapse Studio has a collaboration branch. (main or master is the default collaboration branch).

Box 2: workspace_publish

A branch named dev123 will be created to support the development of a new feature. The workspace templates must be stored in the publish branch.

Creating feature branches

Users can also create feature branches by clicking + New Branch in the branch dropdown.

By default, Synapse Studio generates the workspace templates and saves them into a branch called workspace_publish. To configure a custom publish branch, add a publish_config.json file to the root folder in the collaboration branch.

Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/cicd/source-control>

NEW QUESTION 4

- (Exam Topic 3)

You use Vertipaq Analyzer to analyze a model.

The Relationships tab contains the results shown in the following exhibit.

IsRowNumber Cardinality (Filter)	Relationship Type	Max From Cardinality	Max to Cardinality	1:M Ratio %	Missing Keys	Invalid Rows	Relationships Size	Bid. Filters	MMR
'Date' [Date] ==<-1 'LocalDateTable_39c22ddb-27f3-4e6c-8a44-a3380850fcb4' [Date]	M:1	84	2,557	3044.05%	0	0	4,056		
Fact	M:1	90	327	0.69%	22		184		
'Fact' [BU Key] ==<-1 'BU' [BU Key]	M:1	26	164	0.34%	0	0	32		
'Fact' [Customer Key] ==<-1 'Customer' [Customer]	M:1	90	327	0.69%	21	1,804	112		
'Fact' [Product Key] ==<-1 'Product' [Product Key]	M:1	7	6	0.01%	1	6,577	8		
'Fact' [Scenario Key] ==<-1 'Scenario' [Scenario Key]	M:1	2	2	0.00%	0	0	8		
'Fact' [YearPeriod] ==<-1 'Date' [YearPeriod]	M:1	16	84	0.18%	0	0	24		
Grand Total	M:1	90	2,557	3044.05%	27		4,320		

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.

Answer Area

The [answer choice] table is missing records needed by the Fact table.

▼

BU Key
Customer
Date
Scenario

There are [answer choice] blank values created by missing dimensional relationships.

▼

22
1,804
6,577
8,381

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Customer

There are 1804 invalid rows (records) in the Customer table. Box 2: 22

There are 22 missing keys.

Note: VertiPaq Analyzer in DAX Studio is useful in identifying referential integrity violations which slow down your DAX codes. It helps you determine which table or column needs to be optimized and improved. Reference: <https://blog.enterprisedna.co/vertipaq-analyzer-tutorial-relationships-referential-integrity/>

NEW QUESTION 5

- (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.

Actions

- For Table1, set the permissions for the Project Manager role to **None**.
- From Power BI Desktop, create a role named Project Managers.
- For Table1, set the permissions for the Project Manager role to **Read**.
- Open **DAX Studio**.
- From Power BI Desktop, add a DAX filter to the Project Managers role.
- For the Budget and Forecast columns, set the permissions to **None**.
- Open **Tabular Editor**.

Answer Area

>

<

↑

↓

- 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 6

- (Exam Topic 3)

You have a Power BI report that contains the visual shown in the following exhibit.

Product	Sales
Amarilla	17,747,116.06
Carretera	13,815,307.89
Montana	15,390,801.88
Paseo	33,011,143.95
Velo	18,250,059.47
VTT	20,511,921.02
Total	118,726,350.26

You need to make the visual more accessible to users who have color vision deficiency. What should you do?

- A. Change the font color of values in the Sales column to white.
- B. Change the red background color to orange.
- C. Add icons to represent the sales status of each product.
- D. Add additional measures to the table values.

Answer: A

Explanation:

Themes, contrast and colorblind-friendly colors

You should ensure that your reports have enough contrast between text and any background colors. Certain color combinations are particularly difficult for users with color vision deficiencies to distinguish.

These include the following combinations:

- **--> green and black green and red
- green and brown blue and purple green and blue
- light green and yellow blue and grey
- green and grey

Avoid using these colors together in a chart, or on the same report page.

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

NEW QUESTION 7

- (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 8

- (Exam Topic 3)

You have a Power BI dataset that contains the following measures:

- Budget
- Actuals
- Forecast

You create a report that contains 10 visuals.

You need provide users with the ability to use a slicer to switch between the measures in two visuals only. You create a dedicated measure named cg Measure switch.

How should you complete the DAX expression for the Actuals measure? To answer, drag the appropriate values to the targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: SELECTEDMEASURENAME()

SELECTEDMEASURENAME is used by expressions for calculation items to determine the measure that is in context by name.

Syntax: SELECTEDMEASURENAME()

No parameters. Example:

The following calculation item expression checks if the current measure is Expense Ratio and conditionally applies calculation logic. Since the check is based on a string comparison, it is not subject to formula fixup and will not benefit from object renaming being automatically reflected. For a similar comparison that would benefit from formula fixup, please see the ISSLECTEDMEASURE function instead.

```
IF (
SELECTEDMEASURENAME = "Expense Ratio", SELECTEDMEASURE (),
DIVIDE ( SELECTEDMEASURE (), COUNTROWS ( DimDate ) )
)
```

Box 2: SELECTEDVALUE()

SELECTEDVALUE returns the value when the context for columnName has been filtered down to one distinct value only. Otherwise returns alternateResult.

Syntax:

SELECTEDVALUE(<columnName>[, <alternateResult>]) M1, M2, ... - A list of measures.

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

NEW QUESTION 9

- (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 10

- (Exam Topic 3)

You have a sales report as shown in the following exhibit.



The sales report has the following characteristics: The measures are optimized.

The dataset uses import storage mode.

Data points, hierarchies, and fields cannot be removed or filtered from the report page. From powerbi.com, users experience slow load times when viewing the report.

You need to reduce how long it takes for the report to load without affecting the data displayed in the report. Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Change the report theme to monochromatic.
- B. Replace the single-value cards with a multi-row card.
- C. Replace the product category charts with a bar chart for sales and a hierarchy of Category and Sub Category on the axis.
- D. Replace all the filters on the Filters pane with visual slicers on the report page.

Answer: BC

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 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 11

- (Exam Topic 3)

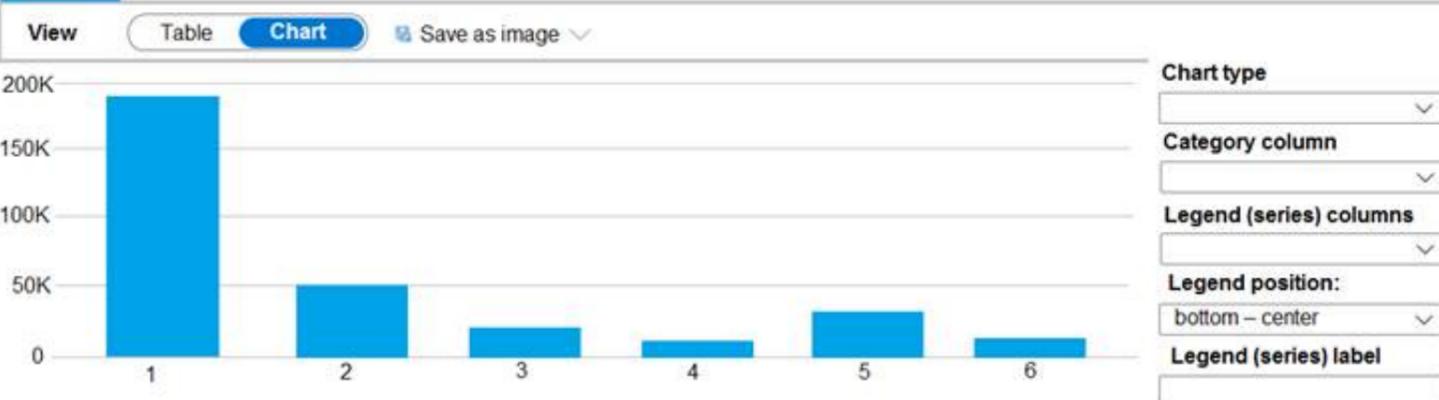
You are using Azure Synapse Studio to explore a dataset that contains data about taxi trips.

You need to create a chart that will show the total trip distance according to the number of passengers as shown in the following exhibit.

```

1 SELECT PassengerCount,
2     SUM(TripDistanceMiles) as SumTripDistance,
3     AVG(TripDistanceMiles) as AvgTripDistance,
4 FROM dbo.NYCTaxiTripSmall
5 WHERE TripDistanceMiles > 0 AND PassengerCount > 0
6 GROUP BY PassengerCount
7 ORDER BY PassengerCount;
    
```

Results Messages



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

Answer Area

Category column:

- AvgTripDistance
- PassengerCount
- SumTripDistance
- TripDistanceMiles

Legend (series) column:

- AvgTripDistance
- PassengerCount
- SumTripDistance
- TripDistanceMiles

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Category column:

- AvgTripDistance
- PassengerCount**
- SumTripDistance
- TripDistanceMiles

Legend (series) column:

- AvgTripDistance
- PassengerCount
- SumTripDistance**
- TripDistanceMiles

NEW QUESTION 16

- (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 17

- (Exam Topic 3)

You are optimizing a Power BI data model by using DAX Studio.

You need to capture the query events generated by a Power BI Desktop report. What should you use?

- A. the DMV list
- B. a Query Plan trace
- C. an All Queries trace
- D. a Server Timings trace

Answer: C

Explanation:

The All Queries trace in Dax Studio supports capturing the query events from all client tools (not just queries sent from DAX Studio like the Query Plan and Server Timings features do). The 'All Queries' trace is really useful when you wish to see the queries that are generated by a client tool like Power BI Desktop.

Reference: <https://daxstudio.org/documentation/features/all-queries-trace/>

NEW QUESTION 21

- (Exam Topic 3)

You are planning a Power BI solution for a customer.

The customer will have 200 Power BI users. The customer identifies the following requirements:

- Ensure that all the users can create paginated reports.
- Ensure that the users can create reports containing AI visuals.
- Provide autoscaling of the CPU resources during heavy usage spikes.

You need to recommend a Power BI solution for the customer. The solution must minimize costs. What should you recommend?

- A. Power BI Premium per user
- B. a Power BI Premium per capacity
- C. Power BI Pro per user
- D. Power BI Report Server

Answer: A

Explanation:

Announcing Power BI Premium Per User general availability and autoscale preview for Gen2. Power BI Premium per user features and capabilities

* Pixel perfect paginated reports are available for operational reporting capabilities based on SSRS technology. Users can create highly formatted reports in various formats such as PDF and PPT, which are embeddable in applications and are designed to be printed or shared.

* Automated machine learning (AutoML) in Power BI enables business users to build ML models to predict outcomes without having to write any code.

* Etc. Note:

Power BI empowers every business user and business analyst to get amazing insights with AI infused experiences. With Power BI Premium, we enable business analysts to not only analyze and visualize their data, but to also build an end-to-end data platform through drag and drop experiences. Everything from ingesting and transforming data at scale, to building automated machine learning models, and analyzing massive volumes of data is now possible for our millions of business analysts.

Reference:

<https://powerbi.microsoft.com/nl-be/blog/announcing-power-bi-premium-per-user-general-availability-and-auto>

NEW QUESTION 23

- (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.

- 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 26

- (Exam Topic 3)

You have the following code in an Azure Synapse notebook.

```
import matplotlib.pyplot as plt
x1 = [2, 3, 4]
y1 = [5, 5, 5]
x2 = [1, 2, 3, 4, 5]
y2 = [2, 3, 2, 3, 4]
y3 = [6, 8, 7, 8, 7]
plt.scatter(x1, y1)
plt.scatter(x2, y2, marker='v', color='r')
plt.scatter(x2, y3, marker='^', color='m')
plt.title('Scatter Plot')
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

There will be [answer choice] rendered as the output of the code.

- one scatterplot
- two scatterplots
- three scatterplots

There will be [answer choice] used in the output.

- one marker symbol
- two marker symbols
- three marker symbols

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: three scatterplots Compare Plots

Example, Draw two plots on the same figure: import matplotlib.pyplot as plt

import numpy as np

#day one, the age and speed of 13 cars:

x = np.array([5,7,8,7,2,17,2,9,4,11,12,9,6])

y = np.array([99,86,87,88,111,86,103,87,94,78,77,85,86])

plt.scatter(x, y)

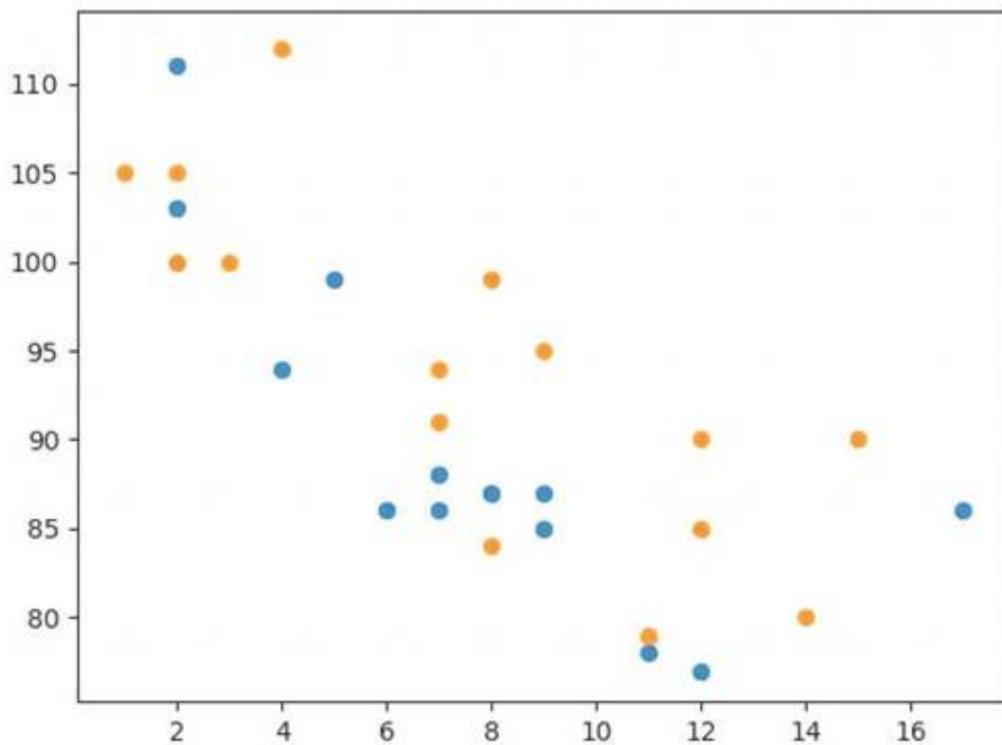
#day two, the age and speed of 15 cars:

x = np.array([2,2,8,1,15,8,12,9,7,3,11,4,7,14,12])

y = np.array([100,105,84,105,90,99,90,95,94,100,79,112,91,80,85])

plt.scatter(x, y) plt.show() Result:

Chart, scatter chart Description automatically generated



Box 2: three marker symbols

One for each scatterplot. One default, and two defined.

Default is point.

v is triangle down.

^ is triangle up.

Reference: https://www.w3schools.com/python/matplotlib_scatter.asp https://matplotlib.org/stable/api/markers_api.html

NEW QUESTION 29

- (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 34

- (Exam Topic 3)

You have a Power BI dataset named Dataset1 that uses DirectQuery against an Azure SQL database named DB1. DB1 is a transactional database in the third normal form.

You need to recommend a solution to minimize how long it takes to execute the query. The solution must maintain the current functionality. What should you include in the recommendation?

- A. Create calculated columns in Dataset1.
- B. Remove the relationships from Dataset1.
- C. Normalize the tables in DB1.
- D. Denormalize the tables in DB1.

Answer: D

Explanation:

Denormalize to improve query performance.

Note: Normalization prevents data duplications, preserves disk space, and improves the performance of the disk I/O operations. The downside of the normalization is that the queries based on these normalized tables require more table joins.

Schema denormalization (i.e. consolidation of some dimension tables) for such databases can significantly reduce costs of the analytical queries and improve the performance.

Reference:

<https://www.mssqltips.com/sqlservertip/7114/denormalization-dimensions-synapse-mapping-data-flow/>

NEW QUESTION 37

- (Exam Topic 2)

You need to recommend a solution to resolve the query issue of the serverless SQL pool. The solution must minimize impact on the users.

What should you in the recommendation?

- A. Update the statistics for the serverless SQL pool.
- B. Move the data from the serverless SQL pool to a dedicated Apache Spark pool.
- C. Execute the sp_sec_process_daca_limic stored procedure for the serverless SQL pool.
- D. Move the data from the serverless SQL pool to a dedicated SQL pool.

Answer: D

Explanation:

Users indicate that queries against the serverless SQL pool fail occasionally because the size of tempdb has been exceeded.

In the dedicated SQL pool resource, temporary tables offer a performance benefit because their results are written to local rather than remote storage.

Temporary tables in serverless SQL pool.

Temporary tables in serverless SQL pool are supported but their usage is limited. They can't be used in queries which target files.

For example, you can't join a temporary table with data from files in storage. The number of temporary tables is limited to 100, and their total size is limited to 100 MB.

Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-temporary>

NEW QUESTION 41

- (Exam Topic 1)

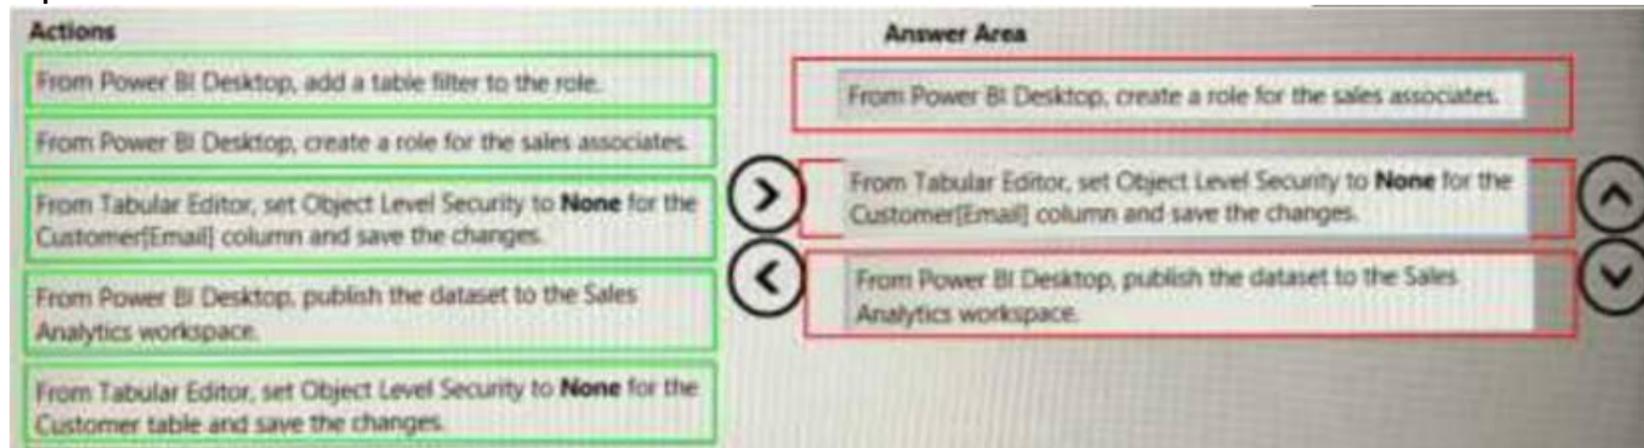
You need to implement object-level security (OLS) in the Power BI dataset for the sales associates.

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.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 45

- (Exam Topic 1)

What should you configure in the deployment pipeline?

- A. a backward deployment
- B. a selective deployment
- C. auto-binding
- D. a data source rule

Answer: D

Explanation:

Development Process Requirements

Litware identifies the following development process requirements:

SQLDW and datalake1 will act as the development environment. Once feature development is complete, all entities in synapseworkspace1 will be promoted to a test workspace, and then to a production workspace.

Power BI content must be deployed to test and production by using deployment pipelines. Create deployment rules

When working in a deployment pipeline, different stages may have different configurations. For example, each stage can have different databases or different query parameters. The development stage might query sample data from the database, while the test and production stages query the entire database.

When you deploy content between pipeline stages, configuring deployment rules enables you to allow changes to content, while keeping some settings intact. For example, if you want a dataset in a production stage to point to a production database, you can define a rule for this. The rule is defined in the production stage, under the appropriate dataset. Once the rule is defined, content deployed from test to production, will inherit the value as defined in the deployment rule, and will always apply as long as the rule is unchanged and valid.

You can configure data source rules and parameter rules.

Incorrect:

Not B: if you already have a steady production environment, you can deploy it backward (to Test or Dev, based on your need) and set up the pipeline. The feature is not limited to any sequential orders.

Reference:

<https://docs.microsoft.com/en-us/power-bi/create-reports/deployment-pipelines-get-started#step-4---create-deplo>

NEW QUESTION 47

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

DP-500 Practice Exam Features:

- * DP-500 Questions and Answers Updated Frequently
- * DP-500 Practice Questions Verified by Expert Senior Certified Staff
- * DP-500 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * DP-500 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The DP-500 Practice Test Here](#)