

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 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 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.)

Vertipaq Analyzer Metrics					
Tables	Columns	Relationships	Partitions	Summary	
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 are using a Python notebook in an Apache Spark pool in Azure Synapse Analytics. You need to present the data distribution statistics from a DataFrame in a tabular view. Which method should you invoke on the DataFrame?

- A. freqItems
- B. explain
- C. rollup
- D. describe

Answer: D

Explanation:

The aggregating statistic can be calculated for multiple columns at the same time with the describe function. Example:

```
titanic[["Age", "Fare"]].describe() Out[6]:
```

```
Age Fare
```

```
count 714.000000 891.000000
```

```
mean 29.699118 32.204208
```

```
std 14.526497 49.693429
```

```
min 0.420000 0.000000
```

```
25% 20.125000 7.910400
```

```
50% 28.000000 14.454200
```

```
75% 38.000000 31.000000
```

```
max 80.000000 512.329200
```

Reference: https://pandas.pydata.org/docs/getting_started/intro_tutorials/06_calculate_statistics.html

NEW QUESTION 4

- (Exam Topic 3)

You have a Power BI tenant.

You plan to register the tenant in an Azure Purview account.

You need to ensure that you can scan the tenant by using Azure Purview.

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

- A. From the Microsoft 365 admin center, create a Microsoft 365 group.
- B. From the Power BI Admin center, set Allow live connections to Enabled.
- C. From the Power BI Admin center, set Allow service principals to use read-only Power BI admin APIs to Enabled.
- D. From the Azure Active Directory admin center, create a security group.
- E. From the Power BI Admin center, set Share content with external users to Enabled.

Answer: CD

Explanation:

Scan same-tenant Power BI using Azure IR and Managed Identity in public network. Make sure Power BI and Microsoft Purview accounts are in the same tenant.

Make sure Power BI tenant Id is entered correctly during the registration.

From Azure portal, validate if Microsoft Purview account Network is set to public access.

From Power BI tenant Admin Portal, make sure Power BI tenant is configured to allow public network.

(D) In Azure Active Directory tenant, create a security group.

From Azure Active Directory tenant, make sure Microsoft Purview account MSI is member of the new security group.

On the Power BI Tenant Admin portal, validate if Allow service principals to use read-only Power BI admin APIs is enabled for the new security group.

Associate the security group with Power BI tenant Log into the Power BI admin portal.

Select the Tenant settings page.

(C) Select Admin API settings > Allow service principals to use read-only Power BI admin APIs (Preview). Select Specific security groups.

Select Admin API settings > Enhance admin APIs responses with detailed metadata > Enable the toggle to allow Microsoft Purview Data Map automatically

discover the detailed metadata of Power BI datasets as part of its scans.

Reference: <https://docs.microsoft.com/en-us/azure/purview/register-scan-power-bi-tenant>

NEW QUESTION 5

- (Exam Topic 3)

You have a Power BI dataset. The dataset contains data that is updated frequently. You need to improve the performance of the dataset by using incremental refreshes.

Which four actions should you perform in sequence to enable the incremental refreshes? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Answer Area

Define the incremental refresh policy for the table.

Enable query caching.

Publish the model to the Power BI service.

Create RangeStart and RangeEnd parameters.

Use the Power BI REST API to post a message to /refreshes.

Apply a custom Date/Time filter to the data.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application, chat or text message Description automatically generated

Step 1: Create RangeStart and RangeEnd parameters. Create parameters

In this task, use Power Query Editor to create RangeStart and RangeEnd parameters with default values. The default values apply only when filtering the data to be loaded into the model in Power BI Desktop. The values you enter should include only a small amount of the most recent data from your data source. When published to the service, these values are overridden by the incremental refresh policy.

Step 2: Apply a custom Date/Time filter to the data. Filter data

With RangeStart and RangeEnd parameters defined, apply a filter based on conditions in the RangeStart and RangeEnd parameters.

Before continuing with this task, verify your source table has a date column of Date/Time data type. Step 3: Define the incremental refresh policy for the table. Define policy

After you've defined RangeStart and RangeEnd parameters, and filtered data based on those parameters, you define an incremental refresh policy. The policy is applied only after the model is published to the service and a manual or scheduled refresh operation is performed.

Step 4: Publish the model to the Power BI service. Save and publish to the service

When your RangeStart and RangeEnd parameters, filtering, and refresh policy settings are complete, be sure to save your model, and then publish to the service.

Reference: <https://docs.microsoft.com/en-us/power-bi/connect-data/incremental-refresh-configure>

NEW QUESTION 6

- (Exam Topic 3)

You have a Power BI report that contains one visual.

You need to provide users with the ability to change the visual type without affecting the view for other users. What should you do?

- A. From Report setting, select Personalize visuals.
- B. From Tabular Editor, create a new perspective.
- C. From the Bookmarks pane, select Focus mode, and then select Add.
- D. From Visual options in Report settings, select Use the modern visual header with updated styling options.

Answer: A

Explanation:

Enable personalization in a report

You can enable the feature either in Power BI Desktop or the Power BI service. You can also enable it in embedded reports.

To enable the feature in Power BI Desktop, go to File > Options and settings > Options > Current file > Report settings. Make sure Personalize visuals is turned on.

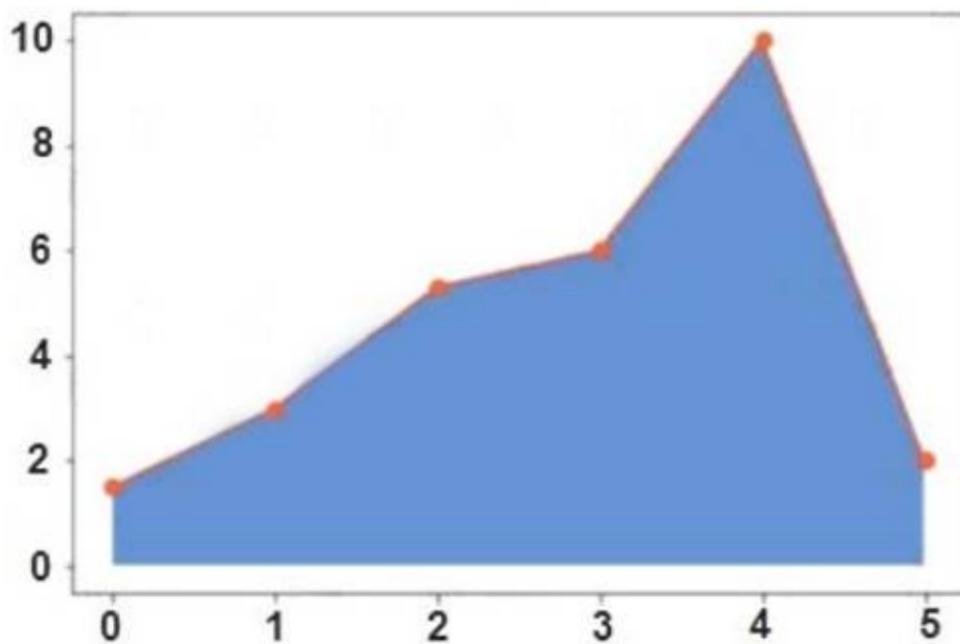
NEW QUESTION 7

- (Exam Topic 3)

You have an Azure Synapse notebook.

You need to create the visual shown in the following exhibit.

chart 1



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

```
import matplotlib.pyplot as plt
x = [0, 1, 2, 3, 4, 5]
y = [1.5, 3, 5.3, 6, 10, 2]
plt.plot(x, y, '-o', color='red')
plt.  (x, y)
plt.  ('chart 1', fontweight='bold')
plt.show()
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: fill_between
 matplotlib.pyplot.fill_between 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.
 Box 2: subtitle
 Set the title of the visual.
 subtitle adds a centred title to the figure. Reference:
https://matplotlib.org/3.1.1/api/_as_gen/matplotlib.pyplot.fill_between.html#matplotlib.pyplot.fill_between
https://matplotlib.org/3.1.1/api/_as_gen/matplotlib.pyplot.subtitle.html#matplotlib.pyplot.subtitle

NEW QUESTION 8

- (Exam Topic 3)
 You have an Azure Synapse Analytics dedicated SQL pool. You need to ensure that the SQL pool is scanned by Azure Purview. What should you do first?

- A. Register a data source.
- B. Search the data catalog.
- C. Create a data share connection.
- D. Create a data policy.

Answer: B

NEW QUESTION 9

- (Exam Topic 3)
 You have a Power BI workspace that contains one dataset and four reports that connect to the dataset. The dataset uses Import storage mode and contains the following data sources:

- A CSV file in an Azure Storage account
 - An Azure Database for PostgreSQL database
- You plan to use deployment pipelines to promote the content from development to test to production. There will be different data source locations for each stage. What should you include in the deployment pipeline to ensure that the appropriate data source locations are used during each stage?
- A. parameter rules
 - B. selective deployment
 - C. auto-binding across pipelines
 - D. data source rules

Answer: A

Explanation:

Note: 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.

NEW QUESTION 10

- (Exam Topic 3)

You are attempting to configure certification for a Power BI dataset and discover that the certification setting for the dataset is unavailable. What are two possible causes of the issue? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. The workspace is in shared capacity.
- B. You have insufficient permissions.
- C. Dataset certification is disabled for the Power BI tenant.
- D. The sensitivity level for the dataset is set to Highly Confidential.
- E. Row-level security (RLS) is missing from the dataset.

Answer: BC

Explanation:

Reference: <https://docs.microsoft.com/en-us/power-bi/admin/service-admin-setup-certification> <https://docs.microsoft.com/en-us/power-bi/collaborate-share/service-endorse-content>

NEW QUESTION 10

- (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 12

- (Exam Topic 3)

You are using DAX Studio to analyze a slow-running report query. You need to identify inefficient join operations in the query. What should you review?

- A. the query statistics
- B. the query plan
- C. the query history
- D. the server timings

Answer: B

Explanation:

Open DAX Studio.

Paste the query there, enable Query Plan display and Server Timings, run your query (with clear cache), and then study the query plan for large row counts. Once the culprit is identified you can decide how to rewrite your DAX to make that part faster.

Reference: <https://community.powerbi.com/t5/Power-Query/DAX-Query-taking-longer-time/td-p/1171961> <https://www.sqlbi.com/wp-content/uploads/DAX-Query-Plans.pdf>

NEW QUESTION 13

- (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 using openrowset with to explicitly specify the maximum length for businessName and surveyName.
 Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Instead use Solution: You recommend using OPENROWSET WITH to explicitly define the collation for businessName and surveyName as Latin1_General_100_BIN2_UTF8.

Query Parquet files using serverless SQL pool in Azure Synapse Analytics. Important

Ensure you are using a UTF-8 database collation (for example Latin1_General_100_BIN2_UTF8) because string values in PARQUET files are encoded using UTF-8 encoding. A mismatch between the text encoding in the PARQUET file and the collation may cause unexpected conversion errors. You can easily change the default collation of the current database using the following T-SQL statement: alter database current collate Latin1_General_100_BIN2_UTF8'.

Note: If you use the Latin1_General_100_BIN2_UTF8 collation you will get an additional performance boost compared to the other collations. The Latin1_General_100_BIN2_UTF8 collation is compatible with parquet string sorting rules. The SQL pool is able to eliminate some parts of the parquet files that will not contain data needed in the queries (file/column-segment pruning). If you use other collations, all data from the parquet files will be loaded into Synapse SQL and the filtering is happening within the SQL process. The Latin1_General_100_BIN2_UTF8 collation has additional performance optimization that works only for parquet and CosmosDB. The downside is that you lose fine-grained comparison rules like case insensitivity.

Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/query-parquet-files>

NEW QUESTION 17

- (Exam Topic 3)

You are creating a Power BI single-page report.

Some users will navigate the report by using a keyboard, and some users will navigate the report by using a screen reader.

You need to ensure that the users can consume content on a report page in a logical order. What should you configure on the report page?

- A. the bookmark order
- B. the X position
- C. the layer order
- D. the tab order

Answer: D

Explanation:

Tab order is the order in which users interact with the items on a page using the keyboard. Generally, we want tab order to be predictable and to closely match the visual order on the page (unless there is a good reason to deviate).

Note: If you are using the keyboard to navigate in a Power BI report, the order in which you arrive at visuals will not follow your vision unless you set the new tab order property. If you have low or no vision, this becomes an even bigger issue because you may not be able to see that you are navigating visuals out of visual order because the screen reader just reads whatever comes next.

Reference: <https://datasavvy.me/2018/12/26/tab-order-enhances-power-bi-report-accessibility/>

NEW QUESTION 19

- (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 22

- (Exam Topic 3)

You have a Power BI Premium capacity.

You need to increase the number of virtual cores associated to the capacity. Which role do you need?

- A. Power BI workspace admin

- B. capacity admin
- C. Power Platform admin
- D. Power BI admin

Answer: D

Explanation:

Change capacity size
 Power BI admins and global administrators can change Power BI Premium capacity. Capacity admins who are not a Power BI admin or global administrator don't have this option.
 Reference: <https://docs.microsoft.com/en-us/power-bi/enterprise/service-admin-premium-manage>

NEW QUESTION 26

- (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 create a hierarchy.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Instead use the 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.

Note: A hierarchy is an ordered set of values that are linked to the level above. An example of a hierarchy could be Country, State, and City. Cities are in a State, and States make up a Country. In Power BI visuals can handle hierarchy data and provide controls for the user to navigate up and down the hierarchy.

Reference: <https://docs.microsoft.com/en-us/power-bi/developer/visuals/capabilities> <https://powerbi.tips/2018/09/how-to-navigate-hierarchies/>

NEW QUESTION 28

- (Exam Topic 3)

You manage a Power BI dataset that queries a fact table named SalesDetails. SalesDetails contains three date columns named OrderDate, CreatedOnDate, and ModifiedDate.

You need to implement an incremental refresh of SalesDetails. The solution must ensure that OrderDate starts on or after the beginning of the prior year.

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.

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Actions

Answer Area

- Create RangeStart and RangeEndDateTime parameters.
- Configure an incremental refresh to archive data that starts one year before the refresh date.
- Add an applied step that filters OrderDate to the start of the prior year.
- Configure an incremental refresh to archive data that starts two years before the refresh date.
- Add an applied step that adds a custom date filter where OrderDate is between RangeStart and RangeEnd.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application, chat or text message Description automatically generated

Step 1: Create RangeStart and RangeEndDateTime parameters.

When configuring incremental refresh in Power BI Desktop, you first create two Power Query date/time parameters with the reserved, case-sensitive names RangeStart and RangeEnd. These parameters, defined in the Manage Parameters dialog in Power Query Editor are initially used to filter the data loaded into the

Power BI Desktop model table to include only those rows with a date/time within that period.

Step 2: Add an applied step that adds a custom date filter OrderDate is Between RangeStart and RangeEnd. With RangeStart and RangeEnd parameters defined, you then apply custom Date filters on your table's date

column. The filters you apply select a subset of data that will be loaded into the model when you click Apply.

Step 3: Configure an incremental refresh to archive data that starts two years before the refresh date.

After filters have been applied and a subset of data has been loaded into the model, you then define an incremental refresh policy for the table. After the model is published to the service, the policy is used by the service to create and manage table partitions and perform refresh operations. To define the policy, you will use the Incremental refresh and real-time data dialog box to specify both required settings and optional settings.

Step 4: Add an applied step that filters OrderDate to the start of the prior year.

Reference: <https://docs.microsoft.com/en-us/power-bi/connect-data/incremental-refresh-overview>

NEW QUESTION 31

- (Exam Topic 3)

You are configuring Azure Synapse Analytics pools to support the Azure Active Directory groups shown in the following table.

Name	Requirement
Group1	Analyze data to create and train machine learning models in Synapse Analytics.
Group2	Execute complex queries with multiple joins against relational data. Results will be exported by using PolyBase.
Group3	Query and load data from Apache Parquet files stored in Azure Data Lake Storage Gen2. Costs must be based on the amount of data processed.

Which type of pool should each group use? To answer, drag the appropriate pool types to the groups. Each pool type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Apache Spark pool

An Apache Spark pool provides open-source big data compute capabilities. After you've created an Apache Spark pool in your Synapse workspace, data can be loaded, modeled, processed, and distributed for faster analytic insight.

Box 2: Dedicated SQL Pool

Dedicated SQL Pool - Data is stored in relational tables

Box 3: Serverless SQL pool

Serverless SQL pool - Cost is incurred for the data processed per query

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/quickstart-create-apache-spark-pool-portal>

<https://www.royalcyber.com/blog/data-services/dedicated-sql-pool-vs-serverless-sql/>

NEW QUESTION 34

- (Exam Topic 3)

Your company is migrating its current, custom-built reporting solution to Power BI. The Power BI tenant must support the following scenarios:

- > 40 reports that will be embedded in external websites. The websites control their own security. The reports will be consumed by 50 users monthly.
- > Forty-five users that require access to the workspaces and apps in the Power BI Admin portal. Ten of the users must publish and consume datasets that are larger than 1 GB.
- > Ten developers that require Text Analytics transformations and paginated reports for datasets. An additional 15 users will consume the reports.

You need to recommend a licensing solution for the company. The solution must minimize costs.

Which two Power BI license options should you include in the recommendation? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. 70 Premium per user
- B. one Premium
- C. 70 Pro
- D. one Embedded
- E. 35 Pro
- F. 35 Premium per user

Answer: BF

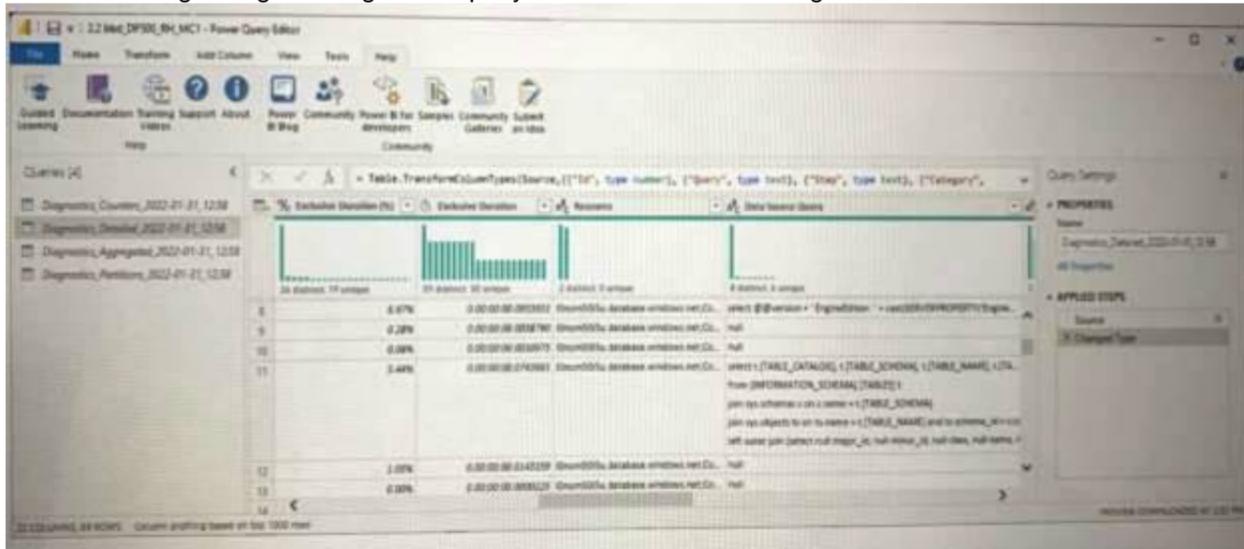
Explanation:

B:
 Free - 40 reports that will be embedded in external websites. The websites control their own security. Free - The reports will be consumed by 50 users monthly.
 Free + 1 Premium for the Worspace -Forty-five users that require access to the workspaces and apps in the Power BI Admin portal.
 F: Ten of the users must publish and consume datasets that are larger than 1 GB.
 Ten developers that require Text Analytics transformations and paginated reports for datasets. An additional 15 users will consume the reports.
 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.
 Note: There are three kinds of Power BI per-user licenses: Free, Pro, and Premium Per User. Power BI (free): Access to content in My Workspace
 Power BI (free) + Workspace is Premium: Consume content shared with them
 Power BI Pro: Publish content to other workspaces, share dashboards, subscribe to dashboards and reports, share with users who have a Pro license
 Power BI Pro + Workspace is Premium: Distribute content to users who have free licenses
 Power BI Premium Per User: Publish content to other workspaces, share dashboards, subscribe to dashboards and reports, share with users who have a Premium Per User license
 Power BI Premium Per User + Workspace is Premium: Distribute content to users who have free and Pro licenses
 Reference: <https://docs.microsoft.com/en-us/power-bi/fundamentals/service-features-license-type>

NEW QUESTION 35

- (Exam Topic 3)

You are running a diagnostic against a query as shown in the following exhibit.



What can you identify from the diagnostics query?

- A. All the query steps are folding.
- B. Elevated permissions are being used to query records.
- C. The query is timing out.
- D. Some query steps are folding.

Answer: A

Explanation:

Understanding folding with Query Diagnostics

One of the most common reasons to use Query Diagnostics is to have a better understanding of what operations were 'pushed down' by Power Query to be performed by the back-end data source, which is also known as 'folding'. If we want to see what folded, we can look at what is the 'most specific' query, or queries, that get sent to the back-end data source. We can look at this for both ODATA and SQL.

Reference: <https://docs.microsoft.com/en-us/power-query/querydiagnosticsfolding>

NEW QUESTION 38

- (Exam Topic 3)

You have a Power BI tenant that contains 10 workspaces.

You need to create dataflows in three of the workspaces. The solution must ensure that data engineers can access the resulting data by using Azure Data Factory. Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point

- A. Associate the Power BI tenant to an Azure Data Lake Storage account.
- B. Add the managed identity for Data Factory as a member of the workspaces.
- C. Create and save the dataflows to an Azure Data Lake Storage account.
- D. Create and save the dataflows to the internal storage of Power BI

Answer: AC

Explanation:

Data used with Power BI is stored in internal storage provided by Power BI by default. With the integration of dataflows and Azure Data Lake Storage Gen 2 (ADLS Gen2), you can store your dataflows in your organization's Azure Data Lake Storage Gen2 account. This essentially allows you to "bring your own storage" to Power BI dataflows, and establish a connection at the tenant or workspace level.

Reference:

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

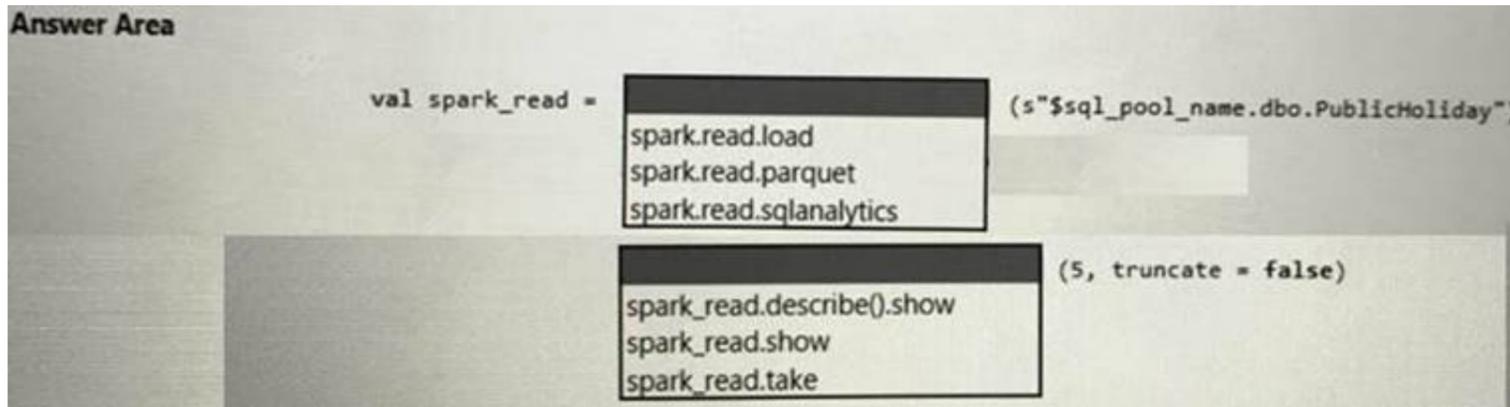
NEW QUESTION 40

- (Exam Topic 3)

You have an Azure Synapse workspace named Workspace1.

You need to use PySpark in a notebook to read data from a SQL pool as an Apache Spark DataFrame and display the top five

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



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: sqlanalytics

Read from a SQL Pool table with Spark

// Read the table we just created in the sql pool as a Spark dataframe val spark_read = spark.read. sqlanalytics(s"\$sql_pool_name.dbo.PublicHoliday")

spark_read.show(5, truncate = false)

Box 2: spark_read.show Sample output:

Graphical user interface, text, application, email Description automatically generated

```
spark_read: org.apache.spark.sql.DataFrame = [countryOrRegion: string, holidayName: string ... 3 more fields]
+-----+-----+-----+-----+-----+
|countryOrRegion|holidayName          |normalizeHolidayName |isPaidTimeOff|countryRegionCode|
+-----+-----+-----+-----+-----+
|Czech          |Den české státnosti  |Den české státnosti  |null         |CZ                |
|Norway         |Søndag               |Søndag               |null         |NO                |
|Sweden         |Søndag               |Søndag               |null         |SE                |
|India          |Gandhi Jayanti       |Gandhi Jayanti       |true         |IN                |
|Germany        |Tag der Deutschen Einheit|Tag der Deutschen Einheit|null         |DE                |
+-----+-----+-----+-----+-----+
only showing top 5 rows
```

Reference:

<https://github.com/Azure-Samples/Synapse/blob/main/Notebooks/Scala/03%20Read%20and%20write%20from>

NEW QUESTION 45

- (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 46

- (Exam Topic 3)

You have a 2-GB Power BI dataset.

You need to ensure that you can redeploy the dataset by using Tabular Editor. The solution must minimize how long it will take to apply changes to the dataset from powerbi.com.

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

- A. Enable service principal authentication for read-only admin APIs.
- B. Turn on Large dataset storage format.
- C. Connect the target workspace to an Azure Data Lake Storage Gen2 account.
- D. Enable XMLA read-write.

Answer: BD

Explanation:

Optimize datasets for write operations by enabling large models

When using the XMLA endpoint for dataset management with write operations, it's recommended you enable the dataset for large models. This reduces the overhead of write operations, which can make them considerably faster. For datasets over 1 GB in size (after compression), the difference can be significant. Tabular Editor supports Azure Analysis Services and Power BI Premium Datasets through XMLA read/write. Note: Tabular Editor - An open-source tool for creating, maintaining, and managing tabular models using an

intuitive, lightweight editor. A hierarchical view shows all objects in your tabular model. Objects are organized by display folders with support for multi-select property editing and DAX syntax highlighting. XMLA read-only is required for query operations. Read-write is required for metadata operations.

Reference: <https://docs.microsoft.com/en-us/power-bi/enterprise/service-premium-connect-tools> <https://tabulareditor.github.io/>

NEW QUESTION 49

- (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. %%ismagic

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 51

- (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 Tabular Editor, you create a calculation group.

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 54

- (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 55

- (Exam Topic 3)

You develop a solution that uses a Power BI Premium capacity. The capacity contains a dataset that is expected to consume 50 GB of memory.

Which two actions should you perform to ensure that you can publish the model successfully to the Power BI service? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Increase the Max Offline Dataset Size setting.
- B. Invoke a refresh to load historical data based on the incremental refresh policy.
- C. Restart the capacity.
- D. Publish an initial dataset that is less than 10 GB.
- E. Publish the complete dataset.

Answer: BE

Explanation:

Enable large datasets

Steps here describe enabling large datasets for a new model published to the service. For existing datasets, only step 3 is necessary.

Create a model in Power BI Desktop. If your dataset will become larger and progressively consume more memory, be sure to configure Incremental refresh.

Publish the model as a dataset to the service.

In the service > dataset > Settings, expand Large dataset storage format, set the slider to On, and then select Apply.

Enable large dataset slider

Invoke a refresh to load historical data based on the incremental refresh policy. The first refresh could take a while to load the history. Subsequent refreshes should be faster, depending on your incremental refresh policy.

Reference: <https://docs.microsoft.com/en-us/power-bi/enterprise/service-premium-large-models>

NEW QUESTION 58

- (Exam Topic 3)

You are building a Power BI dataset that will use two data sources.

The dataset has a query that uses a web data source. The web data source uses anonymous authentication. You need to ensure that the query can be used by all the other queries in the dataset.

Which privacy level should you select for the data source?

- A. Public
- B. Organizational
- C. Private
- D. None

Answer: A

Explanation:

A Public data source gives everyone visibility to the data contained in the data source. Only files, internet data sources, or workbook data can be marked Public.

Data from a Public data source may be freely folded to other sources.

Reference: <https://docs.microsoft.com/en-us/power-bi/enterprise/desktop-privacy-levels>

NEW QUESTION 61

- (Exam Topic 3)

You are using a Python notebook in an Apache Spark pool in Azure Synapse Analytics. You need to present the data distribution statistics from a DataFrame in a tabular view. Which method should you invoke on the DataFrame?

- A. rollup
- B. cov
- C. explain
- D. describe

Answer: D

Explanation:

The aggregating statistic can be calculated for multiple columns at the same time with the describe function. Example:

```
titanic[["Age", "Fare"]].describe() Out[6]:
```

```
Age Fare
```

```
count 714.000000 891.000000
```

```
mean 29.699118 32.204208
```

```
std 14.526497 49.693429
```

```
min 0.420000 0.000000
```

```
25% 20.125000 7.910400
```

```
50% 28.000000 14.454200
```

```
75% 38.000000 31.000000
```

```
max 80.000000 512.329200
```

Reference: https://pandas.pydata.org/docs/getting_started/intro_tutorials/06_calculate_statistics.html

NEW QUESTION 62

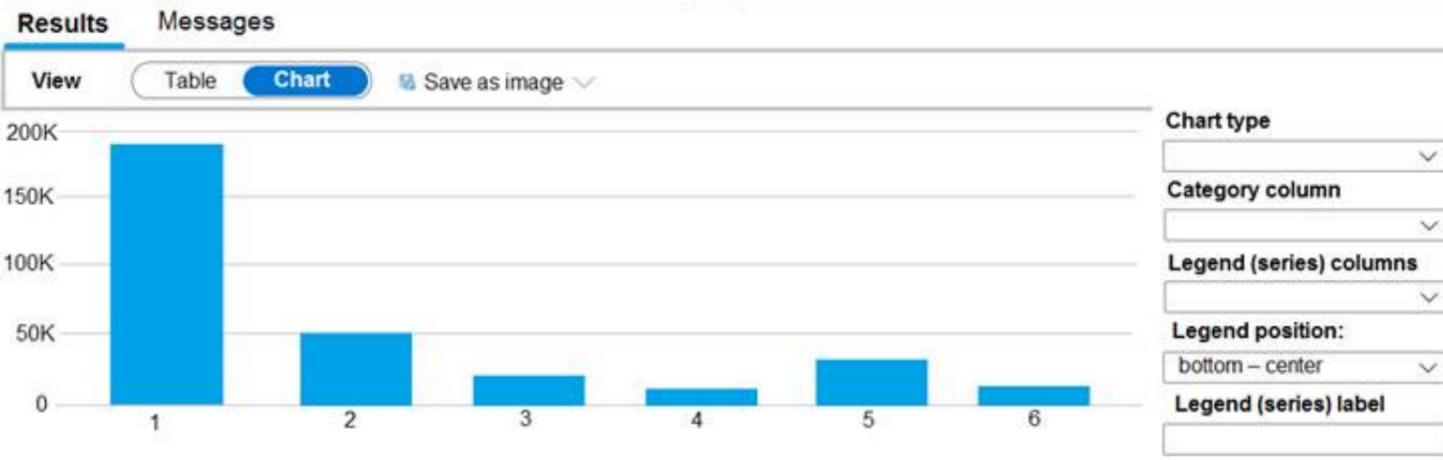
- (Exam Topic 3)

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;
    
```



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 65

- (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 68

- (Exam Topic 3)

You have new security and governance protocols for Power BI reports and datasets. The new protocols must meet the following requirements.

- New reports can be embedded only in locations that require authentication.
- Live connections are permitted only for workspaces that use Premium capacity datasets.

Which three actions should you recommend performing in the Power BI Admin portal? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. From Tenant settings, disable Allow XMLA endpoints and Analyze in Excel with on-premises datasets.
- B. From the Premium per user settings, set XMLA Endpoint to Off.
- C. From Embed Codes, delete all the codes.
- D. From Capacity settings, set XMLA Endpoint to Read Write.
- E. From Tenant settings, set Publish to web to Disable.

Answer: ADE

Explanation:

Reference: <https://docs.microsoft.com/en-us/power-bi/enterprise/service-premium-connect-tools> <https://powerbi.microsoft.com/en-us/blog/power-bi-february-service-update>

NEW QUESTION 69

- (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 72

- (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 77

- (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 82

- (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 87

- (Exam Topic 3)

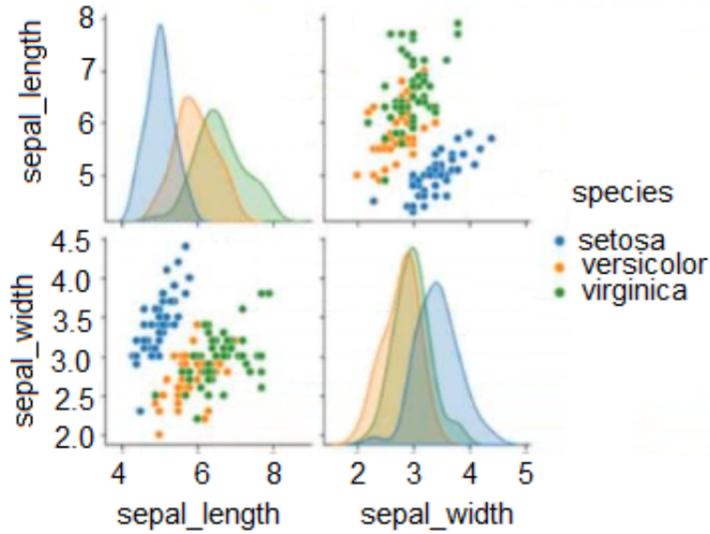
You are using an Azure Synapse notebook to create a Python visual. You run the following code cell to import a dataset named Iris.

```
iris = sns.load_dataset("iris")
iris.head()
```

A sample of the data is shown in the following table.

index	sepal_length	sepal_width	species
0	5.1	3.5	setosa
2	4.9	3	setosa
145	6.7	3	virginica
156	6.3	2.5	virginica

You need to create the visual shown in the exhibit. (Click the Exhibit tab.)



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

Answer Area

```
sns.  (iris, hue= '', height=2.5)

plt.show()
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: pairplot

A pairs plot allows us to see both distribution of single variables and relationships between two variables. Pair plots are a great method to identify trends for follow-up analysis and, fortunately, are easily implemented in Python!

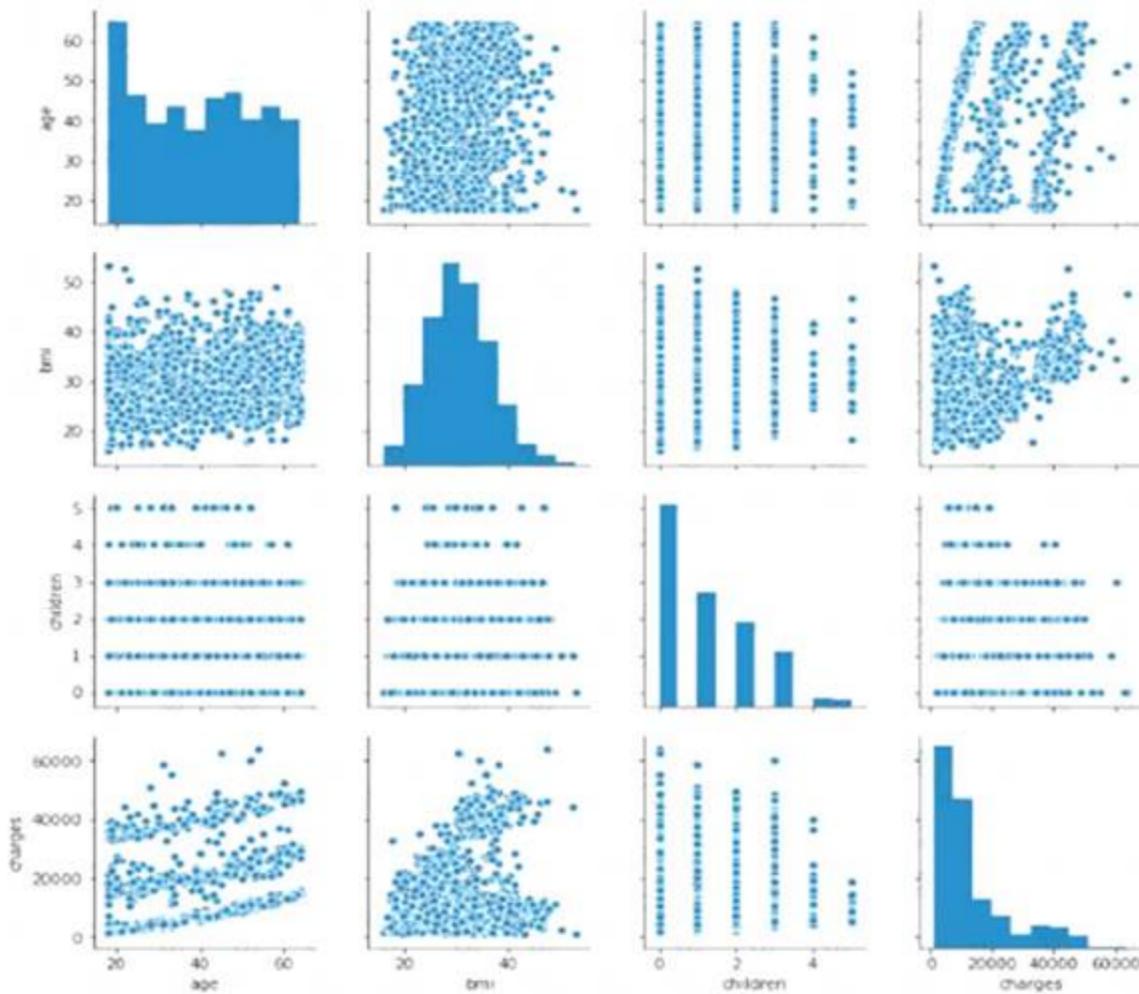
Example, let's plot data using pairplot:

From the picture below, we can observe the variations in each plot. The plots are in matrix format where the row name represents x axis and column name represents the y axis. The main-diagonal subplots are the univariate histograms (distributions) for each attribute.

A picture containing diagram Description automatically generated

```
In [4]: sns.pairplot(df)
```

```
Out[4]: <seaborn.axisgrid.PairGrid at 0x1a1fb1b64ba8>
```



Box 2: sepal_width

sepal_width is displayed with a height of 2.5 (between 2.0 and 4.5).

Reference: <https://medium.com/analytics-vidhya/pairplot-visualization-16325cd725e6>

NEW QUESTION 91

- (Exam Topic 3)

You have a dataset that contains a table named UserPermissions. UserPermissions contains the following data.

User	Region
CONTOSO\User1	1
CONTOSO\User2	2
CONTOSO\User3	1
CONTOSO\User4	3
CONTOSO\User4	5

You plan to create a security role named User Security for the dataset. You need to filter the dataset based on the current users. What should you include in the DAX expression?

- A. [UserPermissions] - USERNAME()
- B. [UserPermissions] - USERPRINCIPALNAME()
- C. [User] = USERPRINCIPALNAME()
- D. [User] = USERNAME()
- E. [User] = USEROBJECTID()

Answer: D

Explanation:

USERNAME() returns the domain name and username from the credentials given to the system at connection time.

It should be compared to column name of User, which in DAX is expressed through [User]. Reference: <https://docs.microsoft.com/en-us/dax/username-function-dax>

NEW QUESTION 96

- (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 dataset
- 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 100

- (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 105

- (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 109

- (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 113

- (Exam Topic 1)

How should you configure the Power BI dataset refresh for the dbo.SalesTransactions table?

- A. an incremental refresh of Product where the ModifiedDate value is during the last three days.
- B. an incremental refresh of dbo.SalesTransactions where the SalesDate value is during the last three days.
- C. a full refresh of all the tables
- D. an incremental refresh of dbo.SalesTransactions where the SalesDate value is during the last hour.

Answer: B

Explanation:

The sales data in SQLDW is updated every 30 minutes. Records in dbo.SalesTransactions are updated in SQLDW up to three days after being created. The records do NOT change after three days.

NEW QUESTION 114

- (Exam Topic 2)

Which two possible tools can you use to identify what causes the report to render slowly? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Synapse Studio
- B. DAX Studio
- C. Azure Data Studio
- D. Performance analyzer in Power BI Desktop

Answer: BD

Explanation:

Some users indicate that the visuals in Power BI reports are slow to render when making filter selections. B: You can investigate a slow query in a Power BI report using DAX Studio, looking at the query plan and the server timings. D: Use Power BI Desktop Performance Analyzer to optimize the report or model. Reference: <https://www.sqlbi.com/tv/analyzing-a-slow-report-query-in-dax-studio/> <https://docs.microsoft.com/en-us/power-bi/guidance/report-performance-troubleshoot>

NEW QUESTION 118

- (Exam Topic 1)

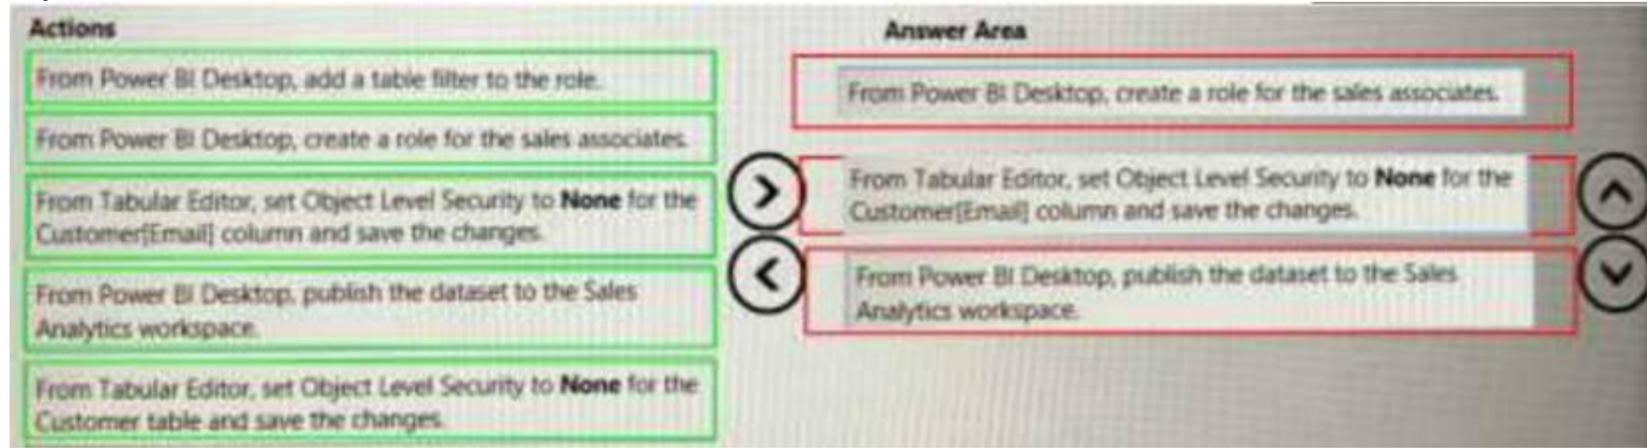
You need to implement object-level security (OLS) in the Power BI dataset for the sales associates.

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:

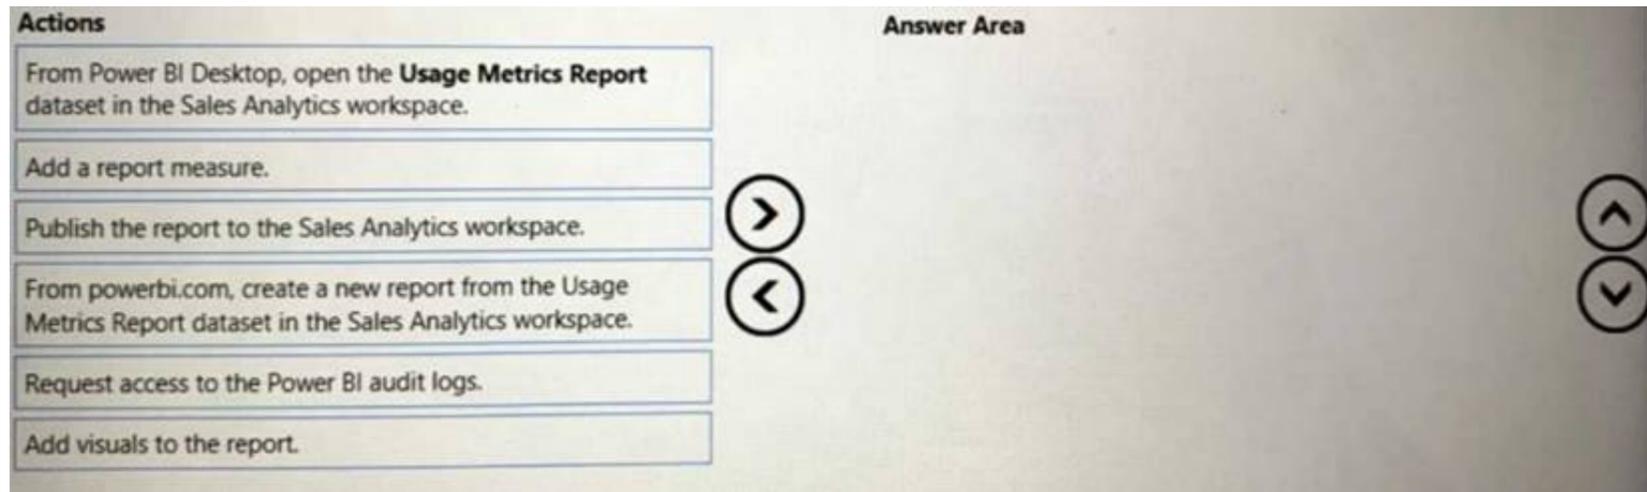


NEW QUESTION 123

- (Exam Topic 1)

You need to create the customized Power BI usage reporting. The Usage Metrics Report dataset has already been created. The solution must minimize development and administrative effort.

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 powerbi.com, create a new report..

The company wants custom Power BI usage reporting that includes the percent change of users that view reports in the Sales Analytics workspace each month.

Step 2: Add a report measure

Measures are used in some of the most common data analyses. Simple summarizations such as sums, averages, minimum, maximum and counts can be set through the Fields well. The calculated results of measures are always changing in response to your interaction with your reports, allowing for fast and dynamic ad-hoc data exploration.

Step 3: Add visuals to the report

Step 4: Publish the report to the Sales Analytics workspace

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

NEW QUESTION 127

- (Exam Topic 1)

You need to populate the CustomersWithProductScore table.

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

```

Answer Area

DECLARE @model
SELECT model
FROM MLModel
WHERE model_name = PredictPurchase
);

INSERT INTO CustomersWithProductScore (
CustomerID
, CustomerEmail
, ProductID
, ProductName
, Score
)

SELECT d.CustomerID
, d.CustomerEmail
, d.ProductID
, d.ProductName
, p.score
FROM PREDICT(MODEL = @model, DATA =
WITH (score FLOAT) AS p;

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: FLOAT

Identify which customers should receive promotional emails based on their likelihood of purchasing promoted products.

FLOT is used in the last statement of the code: WITH (score FLOAT) as p; From syntax: MODEL

The MODEL parameter is used to specify the model used for scoring or prediction. The model is specified as a variable or a literal or a scalar expression.

Box 2: dbo.CustomerWithProductScore

Identify which customers should receive promotional emails based on their likelihood of purchasing promoted products.

Only table CustomerWithProductScore has the required filed score.

From the syntax: DATA

The DATA parameter is used to specify the data used for scoring or prediction. Data is specified in the form of a table source in the query. Table source can be a table, table alias, CTE alias, view, or table-valued function.

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

NEW QUESTION 131

- (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 132

.....

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](#)