VCE & PDF
Pass4itSure.com

DP-500^{Q&As}

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

Pass Microsoft DP-500 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/dp-500.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft Official Exam Center

Instant Download After Purchase

100% Money Back Guarantee

365 Days Free Update

800,000+ Satisfied Customers

QUESTION 1

You have a Power BI workspace named workspace1 that contains three reports and two dataflows.

You have an Azure Data Lake Storage account named storage1.

You need to integrate workspace1 and storage1.

What should you do first?

A. In storage1, create a folder named dataflows.

B. Delete the dataflows from workspace1.

C. In storage1, create a folder named reports.

D. Delete the reports from workspace1.

Correct Answer: B

QUESTION 2

HOTSPOT

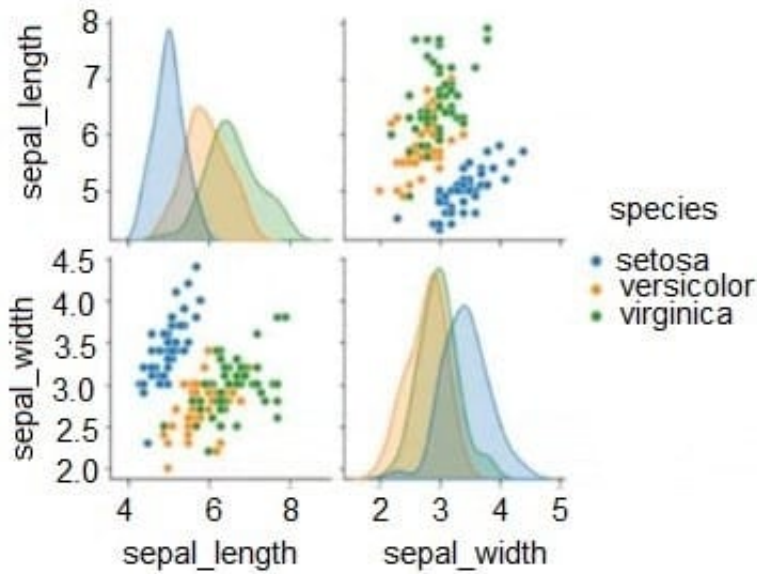You are using an Azure Synapse notebook to create a Python visual.

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.

Hot Area:

Correct Answer:

QUESTION 3

DRAG DROP

You have an Azure Synapse Analytics serverless SQL pool.

You need to return a list of files and the number of rows in each file.

How should you complete the Transact-SQL statement? 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.

Select and Place:

Values

| APPROX_COUNT_DISTINCT |
| COUNT_BIG |
| OPENDATASOURCE |
| OPENJSON |
| OPENQUERY |
| OPENROWSET |

Answer Area

```
SELECT
        asa.filename() AS [filename]
      ,  [                    ] (*) AS [rows]
FROM
        [                    ]
        BULK 'parquet/production/year=2017/month=9/*.parquet',
        DATA_SOURCE = 'DataLake1',
        FORMAT= 'PARQUET'
      ) asa
GROUP BY [filename]
```

Correct Answer:

Values

| |
| COUNT_BIG |
| OPENDATASOURCE |
| OPENJSON |
| OPENQUERY |
| |

Answer Area

```
SELECT
        asa.filename() AS [filename]
      ,  APPROX_COUNT_DISTINCT (*) AS [rows]
FROM
            OPENROWSET
        BULK 'parquet/production/year=2017/month=9/*.parquet',
        DATA_SOURCE = 'DataLake1',
        FORMAT= 'PARQUET'
      ) asa
GROUP BY [filename]
```

Box 1: APPROX_COUNT_DISTINCT

The APPROX_COUNT_DISTINCT function returns the approximate number of unique non- null values in a group.

Box 2: OPENROWSET

OPENROWSET function in Synapse SQL reads the content of the file(s) from a data source. The data source is an Azure storage account and it can be explicitly referenced in the OPENROWSET function or can be dynamically inferred from

URL of the files that you want to read. The OPENROWSET function can optionally contain a DATA_SOURCE parameter to specify the data source that contains files.

The OPENROWSET function can be referenced in the FROM clause of a query as if it were a table name OPENROWSET. It supports bulk operations through a built-in BULK provider that enables data from a file to be read and returned as a

rowset.

QUESTION 4

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

A. the Usage Metrics Report in powerbi.com

B. Query Diagnostics in Power Query Editor

C. Performance analyzer in Power BI Desktop

Correct Answer: B

QUESTION 5

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

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

NOTE: Each correct selection is worth one point.

A. The storage account must be protected by using an Azure Firewall.

B. The connection must be created by a user that is assigned the Storage Blob Data Owner role.

C. The storage account must have hierarchical namespace enabled.

D. Dataflows must exist already for any directly connected Power BI workspaces.

E. The storage account must be created in a separate Azure region from the Power BI tenant and workspaces.

Correct Answer: BC

QUESTION 6

You have an Azure Synapse Analytics serverless SQL pool.

You need to catalog the serverless SQL pool by using Azure Purview.

Which three actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Create a managed identity in Azure Active Directory (Azure AD).

B. Assign the Storage Blob Data Reader role to the Azure Purview managed service identity (MSI) for the storage account associated to the Synapse Analytics workspace.

C. Assign the Owner role to the Azure Purview managed service identity (MSI) for the Azure Purview resource group.

D. Register a data source.

E. Assign the Reader role to the Azure Purview managed service identity (MSI) for the Synapse Analytics workspace.

Correct Answer: BDE

QUESTION 7

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.

Correct Answer: BD

QUESTION 8

You have a Power BI data model.

You need to refresh the data from the source every 15 minutes.

What should you do first?

A. Enable the XMLA endpoint.

B. Define an incremental refresh policy.

C. Change the storage mode of the dataset.

D. Configure a scheduled refresh.

Correct Answer: A

QUESTION 9

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. a Power BI Premium per capacity

B. Power BI Report Server

C. Power BI Premium per user

D. Power BI Pro per user

Correct Answer: A

QUESTION 10

DRAG DROP

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.

Select and Place:

Correct Answer:

QUESTION 11

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.

Correct Answer: A

QUESTION 12

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

Correct Answer: C

QUESTION 13

You have an Azure SQL database named DB1.

You have a Power BI tenant that contains 20 workspaces. Multiple workspaces contain datasets that use DB1 as a data source.

You need to identity how many datasets use DB1 as a data source across the tenant.

What should you use?

A. an XLMA endpoint

B. the Power BI REST API

C. Power BI data source files

D. lineage view

Correct Answer: D

QUESTION 14

You have an Azure subscription that contains an Azure Synapse Analytics workspace. You create an Azure Data Lake Storage Gen2 account and upload a CSV file named Filel.csv. You need to use Synapse Studio to query the data in Filel.csv by using a serverless SQL pool. Which Transact-SQL operator should you include in the query?

A. STRIMO_SPLIT

B. OPENOUERY

C. OPCNROWSET

D. OPEMDATASOURCE

Correct Answer: C

QUESTION 15

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

Correct Answer: BE

Since we need 45 users to have access to the workspaces and apps, 35 Premium Per user would not work because for the Premium Per User workspace content to be accessed we need that all users have Premium Per User. So even if we go for Premium Per User, then we need to purchase it for 70 users.

Also, since we want to publish reports to web and not embed to the app, I would not consider Embedded capacity.

Another important thing is to minimize the cost. We can use Premium capacity for publishing the reports to web, supporting Text Analytics transformations and paginated reports and having dataset more than 1GB. For the users who will be actively busy with development we can purchase 35 pro licenses (IMO 20 pro account is enough).

DP-500 PDF Dumps   DP-500 VCE Dumps   DP-500 Practice Test