

# Microsoft

## Exam Questions DP-600

Implementing Analytics Solutions Using Microsoft Fabric



### NEW QUESTION 1

- (Topic 1)

You need to ensure the data loading activities in the AnalyticsPOC workspace are executed in the appropriate sequence. The solution must meet the technical requirements.

What should you do?

- A. Create a pipeline that has dependencies between activities and schedule the pipeline.
- B. Create and schedule a Spark job definition.
- C. Create a dataflow that has multiple steps and schedule the dataflow.
- D. Create and schedule a Spark notebook.

**Answer:** A

#### Explanation:

To meet the technical requirement that data loading activities must ensure the raw and cleansed data is updated completely before populating the dimensional model, you would need a mechanism that allows for ordered execution. A pipeline in Microsoft Fabric with dependencies set between activities can ensure that activities are executed in a specific sequence. Once set up, the pipeline can be scheduled to run at the required intervals (hourly or daily depending on the data source).

### NEW QUESTION 2

- (Topic 1)

What should you recommend using to ingest the customer data into the data store in the AnalyticsPOC workspace?

- A. a stored procedure
- B. a pipeline that contains a KQL activity
- C. a Spark notebook
- D. a dataflow

**Answer:** D

#### Explanation:

For ingesting customer data into the data store in the AnalyticsPOC workspace, a dataflow (D) should be recommended. Dataflows are designed within the Power BI service to ingest, cleanse, transform, and load data into the Power BI environment. They allow for the low-code ingestion and transformation of data as needed by Litware's technical requirements. References = You can learn more about dataflows and their use in Power BI environments in Microsoft's Power BI documentation.

### NEW QUESTION 3

- (Topic 2)

You have a Fabric tenant that contains 30 CSV files in OneLake. The files are updated daily.

You create a Microsoft Power BI semantic model named Model1 that uses the CSV files as a data source. You configure incremental refresh for Model 1 and publish the model to a Premium capacity in the Fabric tenant.

When you initiate a refresh of Model1, the refresh fails after running out of resources. What is a possible cause of the failure?

- A. Query folding is occurring.
- B. Only refresh complete days is selected.
- C. XMLA Endpoint is set to Read Only.
- D. Query folding is NOT occurring.
- E. The data type of the column used to partition the data has changed.

**Answer:** E

#### Explanation:

A possible cause for the failure is that query folding is NOT occurring (D). Query folding helps optimize refresh by pushing down the query logic to the source system, reducing the amount of data processed and transferred, hence conserving resources. References = The Power BI documentation on incremental refresh and query folding provides detailed information on this topic.

### NEW QUESTION 4

- (Topic 2)

You have a Fabric tenant that uses a Microsoft tower BI Premium capacity. You need to enable scale-out for a semantic model. What should you do first?

- A. At the semantic model level, set Large dataset storage format to Off.
- B. At the tenant level, set Create and use Metrics to Enabled.
- C. At the semantic model level, set Large dataset storage format to On.
- D. At the tenant level, set Data Activator to Enabled.

**Answer:** C

#### Explanation:

To enable scale-out for a semantic model, you should first set Large dataset storage format to On (C) at the semantic model level. This configuration is necessary to handle larger datasets effectively in a scaled-out environment. References = Guidance on configuring large dataset storage formats for scale-out is available in the Power BI documentation.

### NEW QUESTION 5

DRAG DROP - (Topic 2)

You have a Fabric tenant that contains a semantic model. The model contains data about retail stores.

You need to write a DAX query that will be executed by using the XMLA endpoint. The query must return a table of stores that have opened since December 1, 2023.

How should you complete the DAX expression? To answer, drag the appropriate values to the correct 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:**

The correct order for the DAX expression would be:

- ? DEFINE VAR \_SalesSince = DATE ( 2023, 12, 01 )
- ? EVALUATE
- ? FILTER (
- ? SUMMARIZE ( Store, Store[Name], Store[OpenDate] ),
- ? Store[OpenDate] >= \_SalesSince )

In this DAX query, you're defining a variable `_SalesSince` to hold the date from which you want to filter the stores. `EVALUATE` starts the definition of the query. The `FILTER` function is used to return a table that filters another table or expression. `SUMMARIZE` creates a summary table for the stores, including the `Store[Name]` and `Store[OpenDate]` columns, and the filter expression `Store[OpenDate] >= _SalesSince` ensures only stores opened on or after December 1, 2023, are included in the results.

References =

- ? DAX FILTER Function
- ? DAX SUMMARIZE Function

**NEW QUESTION 6**

HOTSPOT - (Topic 2)

You have a data warehouse that contains a table named `Stage.Customers`. `Stage.Customers` contains all the customer record updates from a customer relationship management (CRM) system. There can be multiple updates per customer. You need to write a T-SQL query that will return the customer ID, name, postal code, and the last updated time of the most recent row for each customer ID. How should you complete the code? To answer, select the appropriate options in the answer area.  
 NOTE Each correct selection is worth one point.

Answer Area

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

? In the `ROW_NUMBER()` function, choose `OVER (PARTITION BY CustomerID ORDER BY LastUpdated DESC)`.

? In the `WHERE` clause, choose `WHERE X = 1`.

To select the most recent row for each customer ID, you use the `ROW_NUMBER()` window function partitioned by `CustomerID` and ordered by `LastUpdated` in descending order. This will assign a row number of 1 to the most recent update for each customer. By selecting rows where the row number (`X`) is 1, you get the latest update per customer. References =

- ? Use the `OVER` clause to aggregate data per partition
- ? Use window functions

**NEW QUESTION 7**

HOTSPOT - (Topic 2)

You have a Fabric warehouse that contains a table named `Sales.Orders`. `Sales.Orders` contains the following columns.

Name	Data type	Nullable
OrderID	Integer	No
CustomerID	Integer	No
OrderDate	Date	No
Quantity	Integer	Yes
Weight	Decimal(18, 3)	Yes
ListPrice	Decimal(18, 2)	No
SalePrice	Decimal(18, 2)	Yes

You need to write a T-SQL query that will return the following columns.

Name	Description
OrderID	Returns OrderID
CustomerID	Returns CustomerID
PeriodDate	Returns a date representing the first day of the month for OrderDate
DayName	Returns the name of the day for OrderDate, such as Wednesday

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

**Answer Area**

```
SELECT OrderID, CustomerID,
DATEFROMPARTS
DATE_BUCKET
DATEFROMPARTS
DATEPART
DATETRUNC
DATENAME(
weekday
day
dayofyear
weekday
, OrderDate) AS DayName
FROM Sales.Ord
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

For the PeriodDate that returns the first day of the month for OrderDate, you should use DATEFROMPARTS as it allows you to construct a date from its individual components (year, month, day).

For the DayName that returns the name of the day for OrderDate, you should use

DATENAME with the weekday date part to get the full name of the weekday. The complete SQL query should look like this:

```
SELECT OrderID, CustomerID,
DATEFROMPARTS(YEAR(OrderDate), MONTH(OrderDate), 1) AS PeriodDate, DATENAME(weekday, OrderDate) AS DayName
FROM Sales.Orders
```

Select DATEFROMPARTS for the PeriodDate and weekday for the DayName in the answer area.

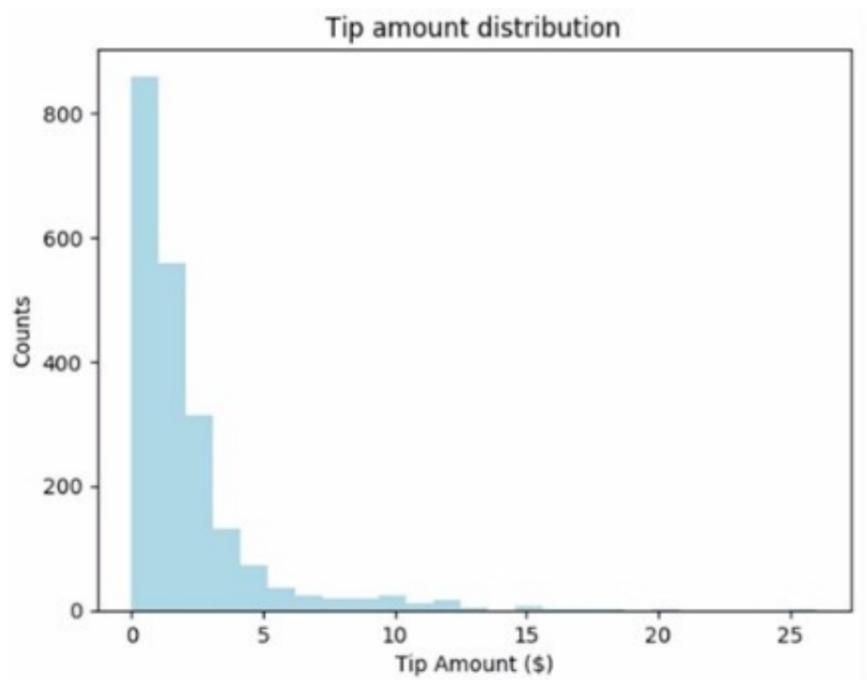
**NEW QUESTION 8**

- (Topic 2)

You have a Fabric notebook that has the Python code and output shown in the following exhibit.

```
# Look at a histogram of tips by count by using Matplotlib

ax1 = sampled_taxi_pd_df['tipAmount'].plot(kind='hist', bins=25, facecolor='lightblue')
ax1.set_title('Tip amount distribution')
ax1.set_xlabel('Tip Amount ($)')
ax1.set_ylabel('Counts')
plt.suptitle('')
plt.show()
```



Which type of analytics are you performing?

- A. predictive
- B. descriptive
- C. prescriptive
- D. diagnostic

**Answer: B**

**Explanation:**

The Python code and output shown in the exhibit display a histogram, which is a representation of the distribution of data. This kind of analysis is descriptive analytics, which is used to describe or summarize the features of a dataset. Descriptive analytics answers the question of "what has happened" by providing insight into past data through tools such as mean, median, mode, standard deviation, and graphical representations like histograms. References: Descriptive analytics and the use of histograms as a way to visualize data distribution are basic concepts in data analysis, often covered in introductory analytics and Python programming resources.

**NEW QUESTION 9**

- (Topic 2)

You have a Fabric tenant that contains a warehouse.

Several times a day, the performance of all warehouse queries degrades. You suspect that Fabric is throttling the compute used by the warehouse. What should you use to identify whether throttling is occurring?

- A. the Capacity settings
- B. the Monitoring hub
- C. dynamic management views (DMVs)
- D. the Microsoft Fabric Capacity Metrics app

**Answer: B**

**Explanation:**

To identify whether throttling is occurring, you should use the Monitoring hub (B). This provides a centralized place where you can monitor and manage the health, performance, and reliability of your data estate, and see if the compute resources are being throttled. References = The use of the Monitoring hub for performance management and troubleshooting is detailed in the Azure Synapse Analytics documentation.

**NEW QUESTION 10**

DRAG DROP - (Topic 2)

You are creating a dataflow in Fabric to ingest data from an Azure SQL database by using a T-SQL statement.

You need to ensure that any foldable Power Query transformation steps are processed by the Microsoft SQL Server engine.

How should you complete the code? To answer, drag the appropriate values to the correct 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.

**Values**

- EnableFolding
- NativeQuery
- Optimize
- Record
- StopFolding
- Table
- Value

**Answer Area**

```

let
    Source = Sql.Databases(
        "server.database.windows.net"
    ),
    Database = Source[[Name = "db"]][Data],
    Query = [ ] - [ ] (
        Database,
        " SELECT * FROM customer WHERE country IN ('USA', 'UK')",
        null,
        [ [ ] = true]
    )
in
    Query
                    
```

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

You should complete the code as follows:

- ? Table
- ? NativeQuery
- ? EnableFolding

In Power Query, using Table before the SQL statement ensures that the result of the SQL query is treated as a table. NativeQuery allows a native database query to be passed through from Power Query to the source database. The EnableFolding option ensures that any subsequent transformations that can be folded will be sent back and executed at the source database (Microsoft SQL Server engine in this case).

**NEW QUESTION 10**

- (Topic 2)

You have a Fabric tenant that contains a semantic model. The model uses Direct Lake mode.

You suspect that some DAX queries load unnecessary columns into memory. You need to identify the frequently used columns that are loaded into memory.

What are two ways to achieve the goal? Each correct answer presents a complete solution. NOTE: Each correct answer is worth one point.

- A. Use the Analyze in Excel feature.
- B. Use the Vertipaq Analyzer tool.
- C. Query the \$system.discovered\_STORAGE\_TABLE\_COLUMN-IN\_SEGMENTS dynamic management view (DMV).
- D. Query the discover\_hehory6Rant dynamic management view (DMV).

**Answer: BC**

**Explanation:**

The Vertipaq Analyzer tool (B) and querying the

\$system.discovered\_STORAGE\_TABLE\_COLUMNS\_IN\_SEGMENTS dynamic management view (DMV) (C) can help identify which columns are frequently loaded into memory. Both methods provide insights into the storage and retrieval aspects of the semantic model. References = The Power BI documentation on Vertipaq Analyzer and DMV queries offers detailed guidance on how to use these tools for performance analysis.

**NEW QUESTION 11**

- (Topic 2)

You have a Microsoft Power BI report named Report1 that uses a Fabric semantic model. Users discover that Report1 renders slowly.

You open Performance analyzer and identify that a visual named Orders By Date is the slowest to render. The duration breakdown for Orders By Date is shown in the following table.

Name	Duration (ms)
DAX query	27
Visual display	39
Other	1047

What will provide the greatest reduction in the rendering duration of Report1?

- A. Change the visual type of Orders By Dale.
- B. Enable automatic page refresh.
- C. Optimize the DAX query of Orders By Date by using DAX Studio.
- D. Reduce the number of visuals in Report1.

**Answer: C**

**Explanation:**

Based on the duration breakdown provided, the major contributor to the rendering duration is categorized as "Other," which is significantly higher than DAX Query and Visual display times. This suggests that the issue is less likely with the DAX calculation or visual rendering times and more likely related to model performance or the complexity of the visual. However, of the options provided, optimizing the DAX query can be a crucial step, even if "Other" factors are dominant. Using DAX Studio, you can analyze and optimize the DAX queries that power your visuals for performance improvements. Here's how you might proceed:

- ? Open DAX Studio and connect it to your Power BI report.
- ? Capture the DAX query generated by the Orders By Date visual.
- ? Use the Performance Analyzer feature within DAX Studio to analyze the query.
- ? Look for inefficiencies or long-running operations.
- ? Optimize the DAX query by simplifying measures, removing unnecessary calculations, or improving iterator functions.
- ? Test the optimized query to ensure it reduces the overall duration.

References: The use of DAX Studio for query optimization is a common best practice for improving Power BI report performance as outlined in the Power BI documentation.

**NEW QUESTION 12**

- (Topic 2)

You have a Fabric tenant that contains a lakehouse.

You plan to query sales data files by using the SQL endpoint. The files will be in an Amazon Simple Storage Service (Amazon S3) storage bucket.

You need to recommend which file format to use and where to create a shortcut. Which two actions should you include in the recommendation? Each correct answer

presents part of the solution.

NOTE: Each correct answer is worth one point.

- A. Create a shortcut in the Files section.
- B. Use the Parquet format
- C. Use the CSV format.
- D. Create a shortcut in the Tables section.
- E. Use the delta format.

**Answer: BD**

**Explanation:**

You should use the Parquet format (B) for the sales data files because it is optimized for performance with large datasets in analytical processing and create a shortcut in the Tables section (D) to facilitate SQL queries through the lakehouse's SQL endpoint. References = The best practices for working with file formats and shortcuts in a lakehouse environment are covered in the lakehouse and SQL endpoint documentation provided by the cloud data platform services.

**NEW QUESTION 17**

HOTSPOT - (Topic 2)

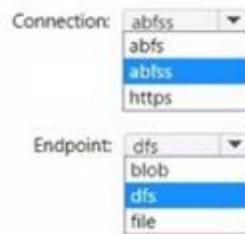
You have a Fabric workspace named Workspace1 and an Azure Data Lake Storage Gen2 account named storage!. Workspace1 contains a lakehouse named Lakehouse1.

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

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

NOTE: Each correct selection is worth one point.

Answer Area



- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

When creating a shortcut to an Azure Data Lake Storage Gen2 account in a lakehouse, you should use the abfss (Azure Blob File System Secure) connection string and the dfs (Data Lake File System) endpoint. The abfss is used for secure access to Azure Data Lake Storage, and the dfs endpoint indicates that the Data Lake Storage Gen2 capabilities are to be used.

**NEW QUESTION 21**

DRAG DROP - (Topic 2)

You are implementing two dimension tables named Customers and Products in a Fabric warehouse.

You need to use slowly changing dimension (SCD) to manage the versioning of data. The solution must meet the requirements shown in the following table.

Table	Change action
Customers	Create a new version of the row.
Products	Overwrite the existing value in the latest row.

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

NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

For the Customers table, where the requirement is to create a new version of the row, you would use:

? Type 2 SCD: This type allows for the creation of a new record each time a change occurs, preserving the history of changes over time.

For the Products table, where the requirement is to overwrite the existing value in the latest row, you would use:

? Type 1 SCD: This type updates the record directly, without preserving historical data.

**NEW QUESTION 24**

- (Topic 2)

You have a Fabric tenant tha1 contains a lakehouse named Lakehouse1. Lakehouse1 contains a Delta table named Customer.

When you query Customer, you discover that the query is slow to execute. You suspect that maintenance was NOT performed on the table.

You need to identify whether maintenance tasks were performed on Customer. Solution: You run the following Spark SQL statement:

REFRESH TABLE customer Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

No, the REFRESH TABLE statement does not provide information on whether maintenance tasks were performed. It only updates the metadata of a table to reflect any changes on the data files. References = The use and effects of the REFRESH TABLE command are explained in the Spark SQL documentation.

**NEW QUESTION 27**

- (Topic 2)

You have a Fabric tenant that contains a lakehouse named Lakehouse1.

You need to prevent new tables added to Lakehouse1 from being added automatically to the default semantic model of the lakehouse.

What should you configure? (5)

- A. the semantic model settings
- B. the Lakehouse1 settings
- C. the workspace settings
- D. the SQL analytics endpoint settings

**Answer:** A

**Explanation:**

To prevent new tables added to Lakehouse1 from being automatically added to the default semantic model, you should configure the semantic model settings. There should be an option within the settings of the semantic model to include or exclude new tables by default. By adjusting these settings, you can control the automatic inclusion of new tables.

References: The management of semantic models and their settings would be covered under the documentation for the semantic layer or modeling features of the Fabric tenant's lakehouse solution.

**NEW QUESTION 29**

- (Topic 2)

You have an Azure Repos Git repository named Repo1 and a Fabric-enabled Microsoft Power BI Premium capacity. The capacity contains two workspaces named Workspace1 and Workspace2. Git integration is enabled at the workspace level.

You plan to use Microsoft Power BI Desktop and Workspace1 to make version-controlled changes to a semantic model stored in Repo1. The changes will be built and deployed to Workspace2 by using Azure Pipelines.

You need to ensure that report and semantic model definitions are saved as individual text files in a folder hierarchy. The solution must minimize development and maintenance effort.

In which file format should you save the changes?

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

**Answer:** C

**Explanation:**

When working with Power BI Desktop and Git integration for version control, report and semantic model definitions should be saved in the PBIX format. PBIX is the Power BI Desktop file format that contains definitions for reports, data models, and queries, and it can be easily saved and tracked in a version-controlled environment. The solution should minimize development and maintenance effort, and saving in PBIX format allows for the easiest transition from development to deployment, especially when using Azure Pipelines for CI/CD (continuous integration/continuous deployment) practices.

References: The use of PBIX files with Power BI Desktop and Azure Repos for version control is discussed in Microsoft's official Power BI documentation, particularly in the sections covering Power BI Desktop files and Azure DevOps integration.

**NEW QUESTION 33**

- (Topic 2)

You have a Fabric tenant that contains a data pipeline.

You need to ensure that the pipeline runs every four hours on Mondays and Fridays. To what should you set Repeat for the schedule?

- A. Daily
- B. By the minute
- C. Weekly
- D. Hourly

**Answer:** C

**Explanation:**

You should set Repeat for the schedule to Weekly (C). This allows you to specify the pipeline to run on specific days of the week, in this case, every four hours on Mondays and Fridays. References = Scheduling options for data pipelines are available in the Azure Data Factory documentation, which includes details on configuring recurring triggers.

**NEW QUESTION 36**

DRAG DROP - (Topic 2)

You are implementing a medallion architecture in a single Fabric workspace.

You have a lakehouse that contains the Bronze and Silver layers and a warehouse that contains the Gold layer.

You create the items required to populate the layers as shown in the following table.

Layer	Data integration tool
Bronze	Pipelines with Copy activities
Silver	Dataflows
Gold	Stored procedures

You need to ensure that the layers are populated daily in sequential order such that Silver is populated only after Bronze is complete, and Gold is populated only after Silver is complete. The solution must minimize development effort and complexity. What should you use to execute each set of items? To answer, drag the appropriate options to the correct items. Each option 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.

**Execution Methods**

- A pipeline Copy activity
- A pipeline Dataflow activity
- A pipeline Stored procedure activity
- A schedule
- A Spark job definition
- An Invoke pipeline activity

**Answer Area**

Orchestration pipeline:

Bronze layer:

Silver layer:

Gold layer:

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

To execute each set of items in sequential order with minimized development effort and complexity, you should use the following options:  
 ? Orchestration pipeline: Use a pipeline with an Invoke pipeline activity. This allows for orchestrating and scheduling the execution of other pipelines, ensuring they run in the correct sequence.  
 ? Bronze layer: Implement a pipeline Copy activity. This aligns with the table indicating that the Bronze layer uses pipelines with Copy activities for data integration.  
 ? Silver layer: Implement a pipeline Dataflow activity. The table specifies that Dataflows are used for the Silver layer.  
 ? Gold layer: Implement a pipeline Stored procedure activity. Stored procedures are specified for the Gold layer according to the table.

**NEW QUESTION 40**

HOTSPOT - (Topic 2)

You have a Fabric tenant that contains a lakehouse. You are using a Fabric notebook to save a large DataFrame by using the following code.

```
df.write.partitionBy("year", "month", "day").mode("overwrite").parquet("Files/SalesOrder")
```

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

**Answer Area**

Statements	Yes	No
The results will form a hierarchy of folders for each partition key.	<input type="radio"/>	<input type="radio"/>
The resulting file partitions can be read in parallel across multiple nodes.	<input type="radio"/>	<input type="radio"/>
The resulting file partitions will use file compression.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

? The results will form a hierarchy of folders for each partition key. - Yes  
 ? The resulting file partitions can be read in parallel across multiple nodes. - Yes  
 ? The resulting file partitions will use file compression. - No  
 Partitioning data by columns such as year, month, and day, as shown in the DataFrame write operation, organizes the output into a directory hierarchy that reflects the partitioning structure. This organization can improve the performance of read operations, as queries that filter by the partitioned columns can scan only the relevant directories. Moreover, partitioning facilitates parallelism because each partition can be processed independently across different nodes in a distributed system like Spark. However, the code snippet provided does not explicitly specify that file compression should be used, so we cannot assume that the output will be compressed without additional context.  
 References =  
 ? DataFrame write partitionBy  
 ? Apache Spark optimization with partitioning

**NEW QUESTION 42**

- (Topic 2)

You have a Fabric tenant that contains a warehouse. A user discovers that a report that usually takes two minutes to render has been running for 45 minutes and has still not rendered. You need to identify what is preventing the report query from completing. Which dynamic management view (DMV) should you use?

- A. sys.dm-exec\_requests
- B. sys.dm\_exec\_sessions
- C. sys.dm\_exec\_connections

D. sys.dm\_pdw\_exec\_requests

**Answer:** D

**Explanation:**

The correct DMV to identify what is preventing the report query from completing is sys.dm\_pdw\_exec\_requests (D). This DMV is specific to Microsoft Analytics Platform System (previously known as SQL Data Warehouse), which is the environment assumed to be used here. It provides information about all queries and load commands currently running or that have recently run. References = You can find more about DMVs in the Microsoft documentation for Analytics Platform System.

**NEW QUESTION 44**

- (Topic 2)

You have a Fabric warehouse that contains a table named Staging.Sales. Staging.Sales contains the following columns.

Name	Data type	Nullable
ProductID	Integer	No
ProductName	Varchar(30)	No
SalesDate	Datetime2(6)	No
WholesalePrice	Decimal(18, 2)	Yes
Amount	Decimal(18, 2)	Yes

You need to write a T-SQL query that will return data for the year 2023 that displays ProductID and ProductName and has a summarized Amount that is higher than 10,000. Which query should you use?

A)

```
SELECT ProductID, ProductName, SUM(Amount) AS TotalAmount
FROM Staging.Sales
WHERE DATEPART(YEAR,SaleDate) = '2023'
GROUP BY ProductID, ProductName
HAVING SUM(Amount) > 10000
```

B)

```
SELECT ProductID, ProductName, SUM(Amount) AS TotalAmount
FROM Staging.Sales
GROUP BY ProductID, ProductName
HAVING DATEPART(YEAR,SaleDate) = '2023' AND SUM(Amount) > 10000
```

C)

```
SELECT ProductID, ProductName, SUM(Amount) AS TotalAmount
FROM Staging.Sales
WHERE DATEPART(YEAR,SaleDate) = '2023' AND SUM(Amount) > 10000
```

D)

```
SELECT ProductID, ProductName, SUM(Amount) AS TotalAmount
FROM Staging.Sales
WHERE DATEPART(YEAR,SaleDate) = '2023'
GROUP BY ProductID, ProductName
HAVING TotalAmount > 10000
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

**Answer:** B

**Explanation:**

The correct query to use in order to return data for the year 2023 that displays ProductID, ProductName, and has a summarized Amount greater than 10,000 is Option B. The reason is that it uses the GROUP BY clause to organize the data by ProductID and ProductName and then filters the result using the HAVING clause to only include groups where the sum of Amount is greater than 10,000. Additionally, the DATEPART(YEAR, SaleDate) = '2023' part of the HAVING clause ensures that only records from the year 2023 are included. References = For more information, please visit the official documentation on T-SQL queries and the GROUP BY clause at T-SQL GROUP BY.

**NEW QUESTION 49**

.....

## **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-600 Practice Exam Features:**

- \* DP-600 Questions and Answers Updated Frequently
- \* DP-600 Practice Questions Verified by Expert Senior Certified Staff
- \* DP-600 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* DP-600 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The DP-600 Practice Test Here](#)**