



**Free Questions for DP-600 by certsdeals**

**Shared by Davis on 22-07-2024**

**For More Free Questions and Preparation Resources**

**Check the Links on Last Page**

# Question 1

---

## Question Type: DragDrop

---

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
```

## Question 2

Question Type: MultipleChoice

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?

### Options:

---

- 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. Reference = You can find more about DMVs in the Microsoft documentation for Analytics Platform System.

## Question 3

---

### Question Type: Hotspot

---

You have a Fabric workspace that uses the default Spark starter pool and runtime version 1,2.

You plan to read a CSV file named Sales.raw.csv in a lakehouse, select columns, and save the data as a Delta table to the managed area of the lakehouse. Sales\_raw.csv contains 12 columns.

You have the following code.

```
from pyspark.sql.functions import year

(spark
 .read
 .format("csv")
 .option("header", 'true')
 .load("Files/sales_raw.csv")
 .select('SalesOrderNumber', 'OrderDate', 'CustomerName', 'UnitPrice')
 .withColumn("Year", year("OrderDate"))
 .write
 .partitionBy('Year')
 .saveAsTable("sales")
)
```

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

The Spark engine will read only the 'SalesOrderNumber', 'OrderDate', 'CustomerName', 'UnitPrice' columns from Sales\_raw.csv.

Yes

No

Removing the partition will reduce the execution time of the query.

Adding `inferSchema='true'` to the options will increase the execution time of the query.

## Question 4

---

### Question Type: MultipleChoice

---

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?

### Options:

---

- 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. Reference = The use of the Monitoring hub for performance management and troubleshooting is detailed in the Azure Synapse Analytics documentation.

## Question 5

---

**Question Type:** MultipleChoice

---

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?

**Options:**

---

- 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. Reference = Scheduling options for data pipelines are available in the Azure Data Factory documentation, which includes details on configuring recurring triggers.



## Question 6

### Question Type: MultipleChoice

You have a Fabric workspace named Workspace1 that contains a data flow named Dataflow1. Dataflow1 contains a query that returns the data shown in the following exhibit.

The screenshot shows a data query interface with a table of data. The table has columns for VendorID and dates from 2015-01-01 to 2015-01-08. The data is as follows:

VendorID	2015-01-01	2015-01-02	2015-01-03	2015-01-04	2015-01-05	2015-01-06	2015-01-07	2015-01-08
1	16	15	0	9	12	9	9	
2	20	15	7	17	20	6	6	

The interface also shows a 'Query settings' panel on the right with the following details:

- Properties: Name: Generated-NYC-Taxi-Gre...
- Entity type: Custom
- Applied steps: Source, Promoted..., Changed c...
- Data destination: No data destination

At the bottom, there is a 'Publish' button.

You need to transform the date columns into attribute-value pairs, where columns become rows.

You select the VendorID column.

Which transformation should you select from the context menu of the VendorID column?

### Options:

---

- A- Group by
- B- Unpivot columns
- C- Unpivot other columns
- D- Split column
- E- Remove other columns

### Answer:

---

B

### Explanation:

---

The transformation you should select from the context menu of the VendorID column to transform the date columns into attribute-value pairs, where columns become rows, is Unpivot columns (B). This transformation will turn the selected columns into rows with two new columns, one for the attribute (the original column names) and one for the value (the data from the cells). Reference = Techniques for unpivoting columns are covered in the Power Query documentation, which explains how to use the transformation in data modeling.

## Question 7

---

### Question Type: Hotspot

---

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"/>

## Explanation:

---

DataFrame write partitionBy

Apache Spark optimization with partitioning

## Question 8

---

**Question Type:** MultipleChoice

---

You have source data in a folder on a local computer.

You need to create a solution that will use Fabric to populate a data store. The solution must meet the following requirements:

- \* Support the use of dataflows to load and append data to the data store.
- \* Ensure that Delta tables are V-Order optimized and compacted automatically.

Which type of data store should you use?

## Options:

---

**A-** a lakehouse

**B-** an Azure SQL database

**C-** a warehouse

**D-** a KQL database

**Answer:**

---

A

**Explanation:**

---

A lakehouse (A) is the type of data store you should use. It supports dataflows to load and append data and ensures that Delta tables are Z-Order optimized and compacted automatically. Reference = The capabilities of a lakehouse and its support for Delta tables are described in the lakehouse and Delta table documentation.

**To Get Premium Files for DP-600 Visit**

**<https://www.p2pexams.com/products/dp-600>**

**For More Free Questions Visit**

**<https://www.p2pexams.com/microsoft/pdf/dp-600>**

