

[Download Full Version DP-203 Exam Dumps\(Updated in Feb/2023\)](#)

a weekly (or more frequent) basis. No matter how much testing is done there is still a risk that an existing, running job may break due to the introduction of a bug. If you are running mission critical jobs, these risks need to be avoided. You can reduce this risk by following Azure's paired region model.

Scenario:

The application development team will create an Azure event hub to receive real-time sales data, including store number, date, time, product ID, customer loyalty number, price, and discount amount, from the point of sale (POS) system and output the data to data storage in Azure

Reference:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-job-reliability>

QUESTION 19

You are designing a partition strategy for a fact table in an Azure Synapse Analytics dedicated SQL pool. The table has the following specifications:

- Contain sales data for 20,000 products.
- Use hash distribution on a column named ProductID.
- Contain 2.4 billion records for the years 2019 and 2020.

Which number of partition ranges provides optimal compression and performance of the clustered columnstore index?

- A. 40
- B. 240
- C. 400
- D. 2,400

Correct Answer: B

QUESTION 20

HOTSPOT

You have an Azure Synapse Analytics dedicated SQL pool that contains the users shown in the following table.

Name	Role
User1	Server admin
User2	db_datareader

User1 executes a query on the database, and the query returns the results shown in the following exhibit.

[Download Full Version DP-203 Exam Dumps\(Updated in Feb/2023\)](#)

```
1. SELECT c.name,
2.       tbl.name as table_name,
3.       typ.name as datatype,
4.       c.is_masked,
5.       c.masking_function
6. FROM sys.masked_columns AS c
7. INNER JOIN sys.tables AS tbl ON c.[object_id] = tbl.[object_id]
8. INNER JOIN sys.types AS typ ON c.user_type_id = typ.user_type_id
9. WHERE is_masked = 1;
```

Results Messages

	name	table_name	datatype	is_masked	masking_function
1	BirthDate	DimCustomer	date	1	default()
2	Gender	DimCustomer	nvarchar	1	default()
3	EmailAddress	DimCustomer	nvarchar	1	email()
4	YearlyIncome	DimCustomer	money	1	default()

User1 is the only user who has access to the unmasked data.

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

Answer Area

When User2 queries the YearlyIncome column, the values returned will be [answer choice].

When User1 queries the BirthDate column, the values returned will be [answer choice].

Correct Answer:

Answer Area

When User2 queries the YearlyIncome column, the values returned will be [answer choice].

When User1 queries the BirthDate column, the values returned will be [answer choice].

QUESTION 21

HOTSPOT

You have the following Azure Stream Analytics query.

[Download Full Version DP-203 Exam Dumps\(Updated in Feb/2023\)](#)

WITH

```
step1 AS (SELECT *
           FROM input1
           PARTITION BY StateID
           INTO 10),
step2 AS (SELECT *
           FROM input2
           PARTITION BY StateID
           INTO 10)
```

```
SELECT *
INTO output
FROM step1
PARTITION BY StateID
UNION step2
  BY StateID
```

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Statements	Yes	No
The query joins two streams of partitioned data.	<input type="radio"/>	<input type="radio"/>
The stream scheme key and count must match the output scheme.	<input type="radio"/>	<input type="radio"/>
Providing 60 streaming units will optimize the performance of the query.	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Statements	Yes	No
The query joins two streams of partitioned data.	<input checked="" type="radio"/>	<input type="radio"/>
The stream scheme key and count must match the output scheme.	<input checked="" type="radio"/>	<input type="radio"/>
Providing 60 streaming units will optimize the performance of the query.	<input checked="" type="radio"/>	<input type="radio"/>

QUESTION 22

You are designing a statistical analysis solution that will use custom proprietary Python functions

[DP-203 Exam Dumps](#) [DP-203 PDF Dumps](#) [DP-203 VCE Dumps](#) [DP-203 Q&As](#)

<https://www.ensurepass.com/DP-203.html>

[Download Full Version DP-203 Exam Dumps\(Updated in Feb/2023\)](#)

on near real-time data from Azure Event Hubs.

You need to recommend which Azure service to use to perform the statistical analysis. The solution must minimize latency.

What should you recommend?

- A. Azure Stream Analytics
- B. Azure SQL Database
- C. Azure Databricks
- D. Azure Synapse Analytics

Correct Answer: A

QUESTION 23

You have an Azure Synapse Analytics dedicated SQL Pool1. Pool1 contains a partitioned fact table named dbo.Sales and a staging table named stg.Sales that has the matching table and partition definitions.

You need to overwrite the content of the first partition in dbo.Sales with the content of the same partition in stg.Sales. The solution must minimize load times.

What should you do?

- A. Switch the first partition from dbo.Sales to stg.Sales.
- B. Switch the first partition from stg.Sales to dbo. Sales.
- C. Update dbo.Sales from stg.Sales.
- D. Insert the data from stg.Sales into dbo.Sales.

Correct Answer: D

QUESTION 24

DRAG DROP

You have an Apache Spark DataFrame named temperatures. A sample of the data is shown in the following table.

Date	Temp
---	---
18-01-2021	3
19-01-2021	4
20-01-2021	2
21-01-2021	2
---	---

You need to produce the following table by using a Spark SQL query.

Year	JAN	FEB	MAR	APR	MAY
2019	2.3	4.1	5.2	7.6	9.2
2020	2.4	4.2	4.9	7.8	9.1
2021	2.6	5.3	3.4	7.9	9.5

[DP-203 Exam Dumps](#) [DP-203 PDF Dumps](#) [DP-203 VCE Dumps](#) [DP-203 Q&As](#)

<https://www.ensurepass.com/DP-203.html>

[Download Full Version DP-203 Exam Dumps\(Updated in Feb/2023\)](#)

How should you complete the query? 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.

The screenshot shows a SQL query editor with two panes. The left pane, titled 'Values', contains a list of functions: CAST, COLLATE, CONVERT, FLATTEN, PIVOT, and UNPIVOT. The right pane, titled 'Answer Area', contains the following SQL query:

```
SELECT * FROM (
    SELECT YEAR(Date) Year, MONTH(Date)
    FROM Temperatures
    WHERE date BETWEEN DATE '2019-01-01' AND DATE
    '2021-08-31'
    Value (
        Value (Temp AS DECIMAL(4, 1)))
    AVG (
        FOR Month in (
            1 JAN, 2 FEB, 3 MAR, 4 APR, 5 MAY, 6
            JUN, 7 JUL, 8 AUG, 9 SEP, 10 OCT, 11 NOV,
            12 DEC
        )
    )
    ORDER BY Year ASC
```

Correct Answer:

The screenshot shows the same SQL query editor as above, but with the correct answer highlighted. The 'Values' list on the left has green dashed boxes around each function. In the 'Answer Area', the 'CONVERT' function is highlighted in the query, and the 'COLLATE' function is highlighted in the 'Value' field of the inner query.

QUESTION 25

HOTSPOT

You are designing an application that will store petabytes of medical imaging data

When the data is first created, the data will be accessed frequently during the first week. After one month, the data must be accessible within 30 seconds, but files will be accessed infrequently. After one year, the data will be accessed infrequently but must be accessible within five minutes.

You need to select a storage strategy for the data. The solution must minimize costs.

Which storage tier should you use for each time frame? To answer, select the appropriate options

[DP-203 Exam Dumps](#) [DP-203 PDF Dumps](#) [DP-203 VCE Dumps](#) [DP-203 Q&As](#)

<https://www.ensurepass.com/DP-203.html>