

100% Pass Quiz 2025 Authoritative Microsoft DP-203: Data Engineering on Microsoft Azure Quiz



DOWNLOAD the newest ITexamReview DP-203 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1RejrRYKBp9rcxdzWzmOeuNV6-uREzh0I>

Data Engineering on Microsoft Azure (DP-203) practice test software is another great way to reduce your stress level when preparing for the Microsoft Exam Questions. With our software, you can practice your excellence and improve your competence on the Data Engineering on Microsoft Azure (DP-203) exam dumps. Each Microsoft DP-203 practice exam, composed of numerous skills, can be measured by the same model used by real examiners.

Microsoft DP-203 (Data Engineering on Microsoft Azure) Certification Exam is a highly sought-after certification for professionals in the data engineering field. DP-203 exam is designed to test the candidate's knowledge and skills in implementing data solutions using Microsoft Azure technologies. DP-203 exam covers a wide range of topics, including data storage, data processing, data analysis, and data visualization. Candidates who pass DP-203 exam are recognized as experts in designing, implementing, and maintaining data solutions on the Azure platform.

Microsoft DP-203 Exam is a comprehensive exam that requires candidates to have hands-on experience in implementing data engineering solutions on the Azure platform. DP-203 exam is designed for individuals who have a good understanding of Azure services and are familiar with data engineering concepts. DP-203 exam is intended for data engineers, data architects, and developers who specialize in building data pipelines, data warehousing, and data analytics solutions on the Azure platform.

>> DP-203 Quiz <<

With Our Information-Packed PDF, Prepare for Microsoft DP-203 Exam Questions

To create the most popular DP-203 exam questions in the market, we have been working hard on the compiling the content and design the displays. And our professional experts have been studying and doing research on the DP-203 study materials for a long time. These experts spent a lot of time before the DP-203 Study Materials officially met with everyone. They spent a lot of time to collate data and carefully studied the characteristics of the stocks to make sure every detail is perfect.

Microsoft Data Engineering on Microsoft Azure Sample Questions (Q35-Q40):

NEW QUESTION # 35

You are implementing an Azure Stream Analytics solution to process event data from devices.

The devices output events when there is a fault and emit a repeat of the event every five seconds until the fault is resolved. The devices output a heartbeat event every five seconds after a previous event if there are no faults present.

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

DeviceID	EventType	EventTime
78cc5ht9-w357-684r-w4fr-kr16h6p9874e	HeartBeat	2020-12-01T19:00.000Z
78cc5ht9-w357-684r-w4fr-kr16h6p9874e	HeartBeat	2020-12-01T19:05.000Z
78cc5ht9-w357-684r-w4fr-kr16h6p9874e	TemperatureSensorFault	2020-12-01T19:07.000Z

You need to calculate the uptime between the faults.

How should you complete the Stream Analytics SQL query? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

SELECT

DeviceID,

MIN(EventTime) as StartTime,

MAX(EventTime) as EndTime,

DATEDIFF(second, MIN(EventTime), MAX(EventTime)) AS duration_in_seconds

FROM input TIMESTAMP BY EventTime

WHERE EventType='HeartBeat'

WHERE LAG(EventType, 1) OVER (LIMIT DURATION(second,5)) <> EventType

WHERE IsFirst(second,5) = 1

GROUP BY

DeviceID

,SessionWindow(second, 5, 50000) OVER (PARTITION BY DeviceID)

,TumblingWindow(second,5)

HAVING DATEDIFF(second, MIN(EventTime), MAX(EventTime)) > 5

Answer:

Explanation:

```

SELECT

DeviceID,

MIN(EventTime) as StartTime,

MAX(EventTime) as EndTime,

DATEDIFF(second, MIN(EventTime), MAX(EventTime)) AS duration_in_seconds

FROM input TIMESTAMP BY EventTime

```

```

WHERE EventType='HeartBeat'
WHERE LAG(EventType, 1) OVER (LIMIT DURATION(second,5)) <> EventType
WHERE IsFirst(second,5) = 1

```

GROUP BY

DeviceID

```

,SessionWindow(second, 5, 50000) OVER (PARTITION BY DeviceID)
,TumblingWindow(second,5)
HAVING DATEDIFF(second, MIN(EventTime), MAX(EventTime)) > 5

```

Reference:

<https://docs.microsoft.com/en-us/stream-analytics-query/session-window-azure-stream-analytics>

<https://docs.microsoft.com/en-us/stream-analytics-query/tumbling-window-azure-stream-analytics>

NEW QUESTION # 36

You have an Azure subscription that contains an Azure Databricks workspace. The workspace contains a notebook named Notebook1. In Notebook1, you create an Apache Spark DataFrame named df_sales that contains the following columns:

- * Customer
- * Salesperson
- * Region
- * Amount

You need to identify the three top performing salespersons by amount for a region named HQ.

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.

Values

agg(col('SalesPerson'))

filter(col('SalesPerson'))

groupBy(col('SalesPerson'))

groupBy(col('TotalAmount'))

orderBy(col('TotalAmount'))

orderBy(desc('TotalAmount'))

Answer Area

```

df_sales.filter(col('Region')=='HQ').
    .agg(sum('Amount').alias('TotalAmount')).
    .limit(3)

```

Answer:

Explanation:

Values

agg(col('SalesPerson'))

filter(col('SalesPerson'))

groupBy(col('SalesPerson'))

groupBy(col('TotalAmount'))

orderBy(col('TotalAmount'))

orderBy(desc('TotalAmount'))

Answer Area

```
df_sales.filter(col('Region')=='HQ').filter(col('SalesPerson'))
.agg(sum('Amount').alias('TotalAmount')).orderBy(desc('TotalAmount')).limit(3)
```

Explanation

Values

agg(col('SalesPerson'))

filter(col('SalesPerson'))

groupBy(col('SalesPerson'))

groupBy(col('TotalAmount'))

orderBy(col('TotalAmount'))

orderBy(desc('TotalAmount'))

Answer Area

```
df_sales.filter(col('Region')=='HQ').filter(col('SalesPerson'))
.agg(sum('Amount').alias('TotalAmount')).orderBy(desc('TotalAmount')).limit(3)
```

NEW QUESTION # 37

You have an Azure Active Directory (Azure AD) tenant that contains a security group named Group1. You have an Azure Synapse Analytics dedicated SQL pool named dw1 that contains a schema named schema1. You need to grant Group1 read-only permissions to all the tables and views in schema1. The solution must use the principle of least privilege.

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.

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

Actions

Create a database role named Role1 and grant Role1 SELECT permissions to schema1.

Create a database role named Role1 and grant Role1 SELECT permissions to dw1.

Assign the Azure role-based access control (Azure RBAC) Reader role for dw1 to Group1.

Create a database user in dw1 that represents Group1 and uses the FROM EXTERNAL PROVIDER clause.

Assign Role1 to the Group1 database user.

Answer Area

Answer:

Explanation:

Actions

Create a database role named Role1 and grant Role1 SELECT permissions to schema1.

Create a database role named Role1 and grant Role1 SELECT permissions to dw1.

Assign the Azure role-based access control (Azure RBAC) Reader role for dw1 to Group1.

Create a database user in dw1 that represents Group1 and uses the FROM EXTERNAL PROVIDER clause.

Assign Role1 to the Group1 database user.

Answer Area

Create a database role named Role1 and grant Role1 SELECT permissions to schema1.

Assign Role1 to the Group1 database user.

Assign the Azure role-based access control (Azure RBAC) Reader role for dw1 to Group1.



Reference:

<https://docs.microsoft.com/en-us/azure/data-share/how-to-share-from-sql>

NEW QUESTION # 38

You are developing a solution using a Lambda architecture on Microsoft Azure.

The data at test layer must meet the following requirements:

Data storage:

- * Serve as a repository (or high volumes of large files in various formats.
- * Implement optimized storage for big data analytics workloads.
- * Ensure that data can be organized using a hierarchical structure.

Batch processing:

- * Use a managed solution for in-memory computation processing.
- * Natively support Scala, Python, and R programming languages.
- * Provide the ability to resize and terminate the cluster automatically.

Analytical data store:

- * Support parallel processing.
- * Use columnar storage.
- * Support SQL-based languages.

You need to identify the correct technologies to build the Lambda architecture.

Which technologies should you use? To answer, select the appropriate options in the answer area NOTE: Each correct selection is worth one point.

Architecture requirement	Microsoft Technology
Data storage	▼ Azure SQL Database Azure Blob Storage Azure Cosmos DB Azure Data Lake Store
Batch processing	▼ HDInsight Spark HDInsight Hadoop Azure Databricks HDInsight Interactive Query
Analytical data store	▼ HDInsight HBase Azure SQL Data Warehouse Azure Analysis Services Azure Cosmos DB

Answer:

Explanation:

Architecture requirement

Technology

Data storage

	▼
Azure SQL Database	
Azure Blob Storage	
Azure Cosmos DB	
Azure Data Lake Store	

Batch processing

	▼
HDInsight Spark	
HDInsight Hadoop	
Azure Databricks	
HDInsight Interactive Query	

Analytical data store

	▼
HDInsight HBase	
Azure SQL Data Warehouse	
Azure Analysis Services	
Azure Cosmos DB	

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-namespaces>

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/batch-processing>

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-overview-what-is>

NEW QUESTION # 39

You have the following Azure Stream Analytics query.

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 combines two streams of partitioned data.	☐	☐
The stream scheme key and count must match the output scheme.	☐	☐
Providing 60 streaming units will optimize the performance of the query.	☐	☐

Answer:

Explanation:

