

Test DP-203 Result, Exam DP-203 Cram



BTW, DOWNLOAD part of DumpsTests DP-203 dumps from Cloud Storage: <https://drive.google.com/open?id=1np5kOd1ULYQBDwfYByCAqjDX4kccWHhN>

Our DP-203 practice quiz will be the optimum resource. Many customers claimed that our study materials made them at once enlightened after using them for review. If you are still tentative about our DP-203 exam dumps, and some exam candidate remain ambivalent to the decision of whether to choose our DP-203 Training Materials, there are free demos for your reference for we understand your hesitation.

The DP-203 exam tests the candidate's knowledge and skills in various areas related to data engineering on Azure, such as designing data storage solutions, implementing data processing solutions, creating data pipelines, and monitoring and optimizing data solutions. DP-203 exam also covers topics related to Azure data services such as Azure Data Factory, Azure Databricks, Azure Stream Analytics, and Azure Synapse Analytics.

The DP-203 exam consists of multiple choice questions that cover a range of topics, including data storage, data processing, data security, and data monitoring. DP-203 Exam is timed and lasts for about 150 minutes. Microsoft recommends that candidates have at least two years of experience working with Azure data services before taking the DP-203 exam. Data Engineering on Microsoft Azure certification is valid for two years, after which candidates must retake the exam or earn a different Azure certification to maintain their credentials.

>> Test DP-203 Result <<

Exam DP-203 Cram & DP-203 Valid Exam Topics

Dare to pursue, we will have a good future. Do you want to be successful people? Do you want to be IT talent? Do you want to pass Microsoft DP-203 certification? DumpsTests will provide you with high quality dumps. It includes real questions and answers, which is useful to the candidates. DumpsTests Microsoft DP-203 Exam Dumps is ordered, finished, and to the point. Only DumpsTests can perfect to show its high quality, however, not every website has high quality exam dumps. Than cardiac operations a rush to purchase our Microsoft DP-203 Oh! The successful rate is 100%.

The DP-203 exam is designed for data engineers, architects, and developers who want to demonstrate their expertise in Azure data technologies. It is intended for professionals who are responsible for designing and implementing data solutions on Azure, and who have experience working with Azure data services. Candidates who pass the DP-203 Exam will be able to demonstrate their ability to create scalable and secure data solutions that meet the needs of their organization.

Microsoft Data Engineering on Microsoft Azure Sample Questions (Q34-Q39):

NEW QUESTION # 34

You have an Azure Databricks workspace named workspace! in the Standard pricing tier. Workspace! contains an all-purpose cluster named cluster). You need to reduce the time it takes for cluster 1 to start and scale up. The solution must minimize costs. What should you do first?

- A. Create a cluster policy in workspace1.
- B. Upgrade workspace1 to the Premium pricing tier.
- C. Configure a global init script for workspace1.
- D. Create a pool in workspace1.

Answer: D

Explanation:

Topic 2, Contoso

Transactional Date

Contoso has three years of customer, transactional, operation, sourcing, and supplier data comprised of 10 billion records stored across multiple on-premises Microsoft SQL Server servers. The SQL server instances contain data from various operational systems. The data is loaded into the instances by using SQL server integration Services (SSIS) packages.

You estimate that combining all product sales transactions into a company-wide sales transactions dataset will result in a single table that contains 5 billion rows, with one row per transaction.

Most queries targeting the sales transactions data will be used to identify which products were sold in retail stores and which products were sold online during different time period. Sales transaction data that is older than three years will be removed monthly.

You plan to create a retail store table that will contain the address of each retail store. The table will be approximately 2 MB.

Queries for retail store sales will include the retail store addresses.

You plan to create a promotional table that will contain a promotion ID. The promotion ID will be associated to a specific product. The product will be identified by a product ID. The table will be approximately 5 GB.

Streaming Twitter Data

The ecommerce department at Contoso develops and Azure logic app that captures trending Twitter feeds referencing the company's products and pushes the products to Azure Event Hubs.

Planned Changes

Contoso plans to implement the following changes:

- * Load the sales transaction dataset to Azure Synapse Analytics.
- * Integrate on-premises data stores with Azure Synapse Analytics by using SSIS packages.
- * Use Azure Synapse Analytics to analyze Twitter feeds to assess customer sentiments about products.

Sales Transaction Dataset Requirements

Contoso identifies the following requirements for the sales transaction dataset:

- * Partition data that contains sales transaction records. Partitions must be designed to provide efficient loads by month. Boundary values must belong to the partition on the right.
- * Ensure that queries joining and filtering sales transaction records based on product ID complete as quickly as possible.
- * Implement a surrogate key to account for changes to the retail store addresses.
- * Ensure that data storage costs and performance are predictable.
- * Minimize how long it takes to remove old records.

Customer Sentiment Analytics Requirement

Contoso identifies the following requirements for customer sentiment analytics:

- * Allow Contoso users to use PolyBase in an Azure Synapse Analytics dedicated SQL pool to query the content of the data records that host the Twitter feeds. Data must be protected by using row-level security (RLS). The users must be authenticated by using their own AzureAD credentials.
- * Maximize the throughput of ingesting Twitter feeds from Event Hubs to Azure Storage without purchasing additional throughput or capacity units.
- * Store Twitter feeds in Azure Storage by using Event Hubs Capture. The feeds will be converted into Parquet files.
- * Ensure that the data store supports Azure AD-based access control down to the object level.
- * Minimize administrative effort to maintain the Twitter feed data records.
- * Purge Twitter feed data records if they are older than two years.

Data Integration Requirements

Contoso identifies the following requirements for data integration:

Use an Azure service that leverages the existing SSIS packages to ingest on-premises data into datasets stored in a dedicated SQL pool of Azure Synapse Analytics and transform the data.

Identify a process to ensure that changes to the ingestion and transformation activities can be version controlled and developed independently by multiple data engineers.

NEW QUESTION # 35

You have an Azure Synapse Analytics pipeline named Pipeline1 that contains a data flow activity named Dataflow1.

Pipeline1 retrieves files from an Azure Data Lake Storage Gen 2 account named storage1.

Dataflow1 uses the AutoResolveIntegrationRuntime integration runtime configured with a core count of 128.

You need to optimize the number of cores used by Dataflow1 to accommodate the size of the files in storage1.

What should you configure? To answer, select the appropriate options in the answer area.



To Pipeline1, add:

- A custom activity
- A Get Metadata activity
- An If Condition activity

For Dataflow1, set the core count by using:

- Dynamic content
- Parameters
- User properties

Answer:

Explanation:



To Pipeline1, add:

- A custom activity
- A Get Metadata activity
- An If Condition activity

For Dataflow1, set the core count by using:

- Dynamic content
- Parameters
- User properties

Explanation

Box 1: A Get Metadata activity

Dynamically size data flow compute at runtime

The Core Count and Compute Type properties can be set dynamically to adjust to the size of your incoming source data at runtime. Use pipeline activities like Lookup or Get Metadata in order to find the size of the source dataset data. Then, use Add Dynamic Content in the Data Flow activity properties.

Box 2: Dynamic content

Reference: <https://docs.microsoft.com/en-us/azure/data-factory/control-flow-execute-data-flow-activity>

NEW QUESTION # 36

You need to design an analytical storage solution for the transactional data. The solution must meet the sales transaction dataset requirements.

What should you include in the solution? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Microsoft

Table type to store retail store data:

Hash
Replicated
Round-robin

Table type to store promotional data:

Hash
Replicated
Round-robin

Answer:

Explanation:

Table type to store retail store data:





Table type to store promotional data:



Explanation:

Graphical user interface, text, application, table Description automatically generated

Table type to store retail store data:

Table type to store promotional data:

Box 1: Round-robin

Round-robin tables are useful for improving loading speed.

Scenario: Partition data that contains sales transaction records. Partitions must be designed to provide efficient loads by month.

Box 2: Hash

Hash-distributed tables improve query performance on large fact tables.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribution>

Topic 1, Contoso Case Study Transactional Data

Contoso has three years of customer, transactional, operation, sourcing, and supplier data comprised of 10 billion records stored across multiple on-premises Microsoft SQL Server servers. The SQL server instances contain data from various operational systems. The data is loaded into the instances by using SQL server integration Services (SSIS) packages.

You estimate that combining all product sales transactions into a company-wide sales transactions dataset will result in a single table that contains 5 billion rows, with one row per transaction.

Most queries targeting the sales transactions data will be used to identify which products were sold in retail stores and which products were sold online during different time period. Sales transaction data that is older than three years will be removed monthly.

You plan to create a retail store table that will contain the address of each retail store. The table will be approximately 2 MB.

Queries for retail store sales will include the retail store addresses.

You plan to create a promotional table that will contain a promotion ID. The promotion ID will be associated to a specific product.

The product will be identified by a product ID. The table will be approximately 5 GB.

Streaming Twitter Data

The ecommerce department at Contoso develops and Azure logic app that captures trending Twitter feeds referencing the company's products and pushes the products to Azure Event Hubs.

Planned Changes

Contoso plans to implement the following changes:

- * Load the sales transaction dataset to Azure Synapse Analytics.
- * Integrate on-premises data stores with Azure Synapse Analytics by using SSIS packages.
- * Use Azure Synapse Analytics to analyze Twitter feeds to assess customer sentiments about products.

Sales Transaction Dataset Requirements

Contoso identifies the following requirements for the sales transaction dataset:

- * Partition data that contains sales transaction records. Partitions must be designed to provide efficient loads by month. Boundary values must belong to the partition on the right.
- * Ensure that queries joining and filtering sales transaction records based on product ID complete as quickly as possible.
- * Implement a surrogate key to account for changes to the retail store addresses.

* Ensure that data storage costs and performance are predictable.

* Minimize how long it takes to remove old records.

Customer Sentiment Analytics Requirement

Contoso identifies the following requirements for customer sentiment analytics:

* Allow Contoso users to use PolyBase in an Azure Synapse Analytics dedicated SQL pool to query the content of the data records that host the Twitter feeds. Data must be protected by using row-level security (RLS). The users must be authenticated by using their own AzureAD credentials.

* Maximize the throughput of ingesting Twitter feeds from Event Hubs to Azure Storage without purchasing additional throughput or capacity units.

* Store Twitter feeds in Azure Storage by using Event Hubs Capture. The feeds will be converted into Parquet files.

* Ensure that the data store supports Azure AD-based access control down to the object level.

* Minimize administrative effort to maintain the Twitter feed data records.

* Purge Twitter feed data records if they are older than two years.

Data Integration Requirements

Contoso identifies the following requirements for data integration:

Use an Azure service that leverages the existing SSIS packages to ingest on-premises data into datasets stored in a dedicated SQL pool of Azure Synapse Analytics and transform the data.

Identify a process to ensure that changes to the ingestion and transformation activities can be version controlled and developed independently by multiple data engineers.

NEW QUESTION # 37

You are implementing Azure Stream Analytics windowing functions.

Which windowing function should you use for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

 Microsoft

Segment the data stream into distinct time segments that repeat but do not overlap:

Hopping

Sliding

Tumbling

Segment the data stream into distinct time segments that repeat and can overlap:

Hopping

Sliding

Tumbling

Segment the data stream to produce an output only when an event occurs:

Hopping

Sliding

Tumbling

Answer:

Explanation:

Answer Area

 Microsoft

Segment the data stream into distinct time segments that repeat but do not overlap:

Hopping

Sliding

Tumbling

Segment the data stream into distinct time segments that repeat and can overlap:

Hopping

Sliding

Tumbling

Segment the data stream to produce an output only when an event occurs:

Hopping

Sliding

Tumbling

NEW QUESTION # 38

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have

a correct solution.

After you answer a question in this scenario, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure Storage account that contains 100 GB of files. The files contain text and numerical values.

75% of the rows contain description data that has an average length of 1.1 MB.

You plan to copy the data from the storage account to an Azure SQL

You need to prepare the files to ensure that the data copies quickly.

Solution: You modify the

- A. Yes

4

Evaluation

Explanation:

Instead modify the files to ensure that each row is less than 1 MB.

Instead mod
References:

References:
<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/guidance-for-loading-data>

NEW QUESTION # 39

Exam DP-203 Cram: <https://www.dumpstests.com/DP-203-latest-test-dumps.html>