


Trustworthy DP-203 Exam Content - DP-203 Valid Examcollection



BTW, DOWNLOAD part of ITexamReview DP-203 dumps from Cloud Storage: https://drive.google.com/open?id=17ROVP12jdfZJWVpGpYEZ3CqPXotVQA_g

It is a popular belief that only professional experts can be the leading one to do some adept job. And similarly, only high quality and high accuracy DP-203 Exam Questions like ours can give you confidence and reliable backup to get the certificate smoothly because our experts have extracted the most frequent-tested points for your reference. Good practice materials like our Data Engineering on Microsoft Azure study question can educate exam candidates with the most knowledge. Do not make your decisions now will be a pity for good.

For more information about the Microsoft DP-203 Exam visit the following reference link:

Microsoft DP-203 Exam Reference link

Microsoft DP-203 (Data Engineering on Microsoft Azure) certification exam is designed for professionals who want to validate their skills in designing and implementing data solutions on Microsoft Azure. Data Engineering on Microsoft Azure certification exam measures your ability to work with different Azure data services like Azure Data Factory, Azure Stream Analytics, Azure Databricks, and more. As a data engineer, you will learn how to use these tools to transform raw data into meaningful insights that can help organizations make better decisions.

>> **Trustworthy DP-203 Exam Content** <<

Microsoft DP-203 Exam Practice Test To Gain Brilliant Result

We have security and safety guarantee, which mean that you cannot be afraid of virus intrusion and information leakage since we have data protection acts, even though you end up studying DP-203 test guide of our company, we will absolutely delete your personal information and never against ethic code to sell your message to the third parties. Our DP-203 Exam Questions will spare no effort to perfect after-sales services. Thirdly countless demonstration and customer feedback suggest that our Data Engineering on Microsoft Azure study question can help them get the certification as soon as possible, thus becoming the elite, getting a promotion and a raise and so forth.

Microsoft Data Engineering on Microsoft Azure Sample Questions (Q181-Q186):

NEW QUESTION # 181

You need to design a data retention solution for the Twitter feed data records. The solution must meet the customer sentiment analytics requirements.

Which Azure Storage functionality should you include in the solution?

- A. change feed
- **B. soft delete**
- C. time-based retention
- D. lifecycle management

Answer: B

Explanation:

Scenario: Purge Twitter feed data records that are older than two years.

Data sets have unique lifecycles. Early in the lifecycle, people access some data often. But the need for access often drops drastically as the data ages. Some data remains idle in the cloud and is rarely accessed once stored. Some data sets expire days or months after creation, while other data sets are actively read and modified throughout their lifetimes. Azure Storage lifecycle management offers a rule-based policy that you can use to transition blob data to the appropriate access tiers or to expire data at the end of the data lifecycle.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview>

Topic 1, Contoso

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 Azure AD 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 that 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 # 182

You need to design a data retention solution for the Twitter feed data records. The solution must meet the customer sentiment analytics requirements.

Which Azure Storage functionality should you include in the solution?

- A. soft delete
- **B. lifecycle management**
- C. change feed
- D. time-based retention

Answer: B

Explanation:

Scenario: Purge Twitter feed data records that are older than two years.

Data sets have unique lifecycles. Early in the lifecycle, people access some data often. But the need for access often drops drastically as the data ages. Some data remains idle in the cloud and is rarely accessed once stored. Some data sets expire days or months after creation, while other data sets are actively read and modified throughout their lifetimes. Azure Storage lifecycle management offers a rule-based policy that you can use to transition blob data to the appropriate access tiers or to expire data at the end of the data lifecycle.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview> This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

NEW QUESTION # 183

You have an Azure Data Lake Storage Gen2 account that contains a JSON file for customers. The file contains two attributes named FirstName and LastName.

You need to copy the data from the JSON file to an Azure Synapse Analytics table by using Azure Databricks. A new column must be created that concatenates the FirstName and LastName values.

You create the following components:

- * A destination table in Azure Synapse
- * An Azure Blob storage container
- * A service principal

Which five actions should you perform in sequence next in is Databricks notebook? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Mount the Data Lake Storage onto DBFS.	
Write the results to a table in Azure Synapse.	
Perform transformations on the file.	
Specify a temporary folder to stage the data.	
Write the results to Data Lake Storage.	
Read the file into a data frame.	
Drop the data frame.	
Perform transformations on the data frame.	

 Microsoft

Answer:

Explanation:

Actions

Mount the Data Lake Storage onto DBFS.
Write the results to a table in Azure Synapse.
Perform transformations on the file.
Specify a temporary folder to stage the data.
Write the results to Data Lake Storage.
Read the file into a data frame.
Drop the data frame.
Perform transformations on the data frame.

Answer Area

Mount the Data Lake Storage onto DBFS.
Read the file into a data frame.
Perform transformations on the data frame.
Specify a temporary folder to stage the data.
Write the results to a table in Azure Synapse.

Explanation:

- 1) mount onto DBFS
- 2) read into data frame
- 3) transform data frame
- 4) specify temporary folder
- 5) write the results to table in Azure Synapse

<https://docs.databricks.com/data/data-sources/azure/azure-datalake-gen2.html> <https://docs.microsoft.com/en-us/azure/databricks/scenarios/databricks-extract-load-sql-data-warehouse>

NEW QUESTION # 184

You have Azure Data Factory configured with Azure Repos Git integration. The collaboration branch and the publish branch are set to the default values.

You have a pipeline named pipeline 1.

You build a new version of pipeline1 in a branch named feature 1.

From the Data Factory Studio, you select Publish

The source code of which branch will be built, and which branch will contain the output of the Azure Resource Manager (ARM) template? To answer, select the appropriate options in the answer area.

Answer Area

Source code:

ARM template output:

Answer:

Explanation:



Explanation:



NEW QUESTION # 185

You have an Azure Synapse workspace named MyWorkspace that contains an Apache Spark database named mytestdb.

You run the following command in an Azure Synapse Analytics Spark pool in MyWorkspace.

```
CREATE TABLE mytestdb.myParquetTable(
```

```
EmployeeID int,
```

```
EmployeeName string,
```

```
EmployeeStartDate date)
```

```
USING Parquet
```

You then use Spark to insert a row into mytestdb.myParquetTable. The row contains the following data.

EmployeeName	EmployeeID	EmployeeStartDate
Alice	24	2020-01-25

One minute later, you execute the following query from a serverless SQL pool in MyWorkspace.

```
SELECT EmployeeID
```

```
FROM mytestdb.dbo.myParquetTable
```

```
WHERE name = 'Alice';
```

What will be returned by the query?

- A. 0
- B. a null value
- C. an error

Answer: C

Explanation:

Once a database has been created by a Spark job, you can create tables in it with Spark that use Parquet as the storage format.

Table names will be converted to lower case and need to be queried using the lower case name. These tables will immediately become available for querying by any of the Azure Synapse workspace Spark pools. They can also be used from any of the Spark jobs subject to permissions.

Note: For external tables, since they are synchronized to serverless SQL pool asynchronously, there will be a delay until they appear.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/metadata/table>

NEW QUESTION # 186

.....

If you want our DP-203 study materials to download and print, the PDF version is perfect for you since it has the function of being printable. And the PDF version of our DP-203 exam questions can be noted when you want to memory something as well as to indicate the keypoints. Also, our DP-203 Preparation exam is unlimited in number of devices, making it easy for you to learn

DP-203 Valid Examcollection: <https://www.itexamreview.com/DP-203-exam-dumps.html>

- What's more, part of that ITexamReview DP-203 dumps now are free: https://drive.google.com/open?id=17ROVP12jdfZJWVpGpYEZ3CqPXotVQA_g

What's more, part of that ITexamReview DP-203 dumps now are free: https://drive.google.com/open?id=17ROVP12jdfZJWVpGpYEZ3CqPXotVQA_g