DP-203 Real Questions | Reliable DP-203 Exam Price

BTW, DOWNLOAD part of BraindumpStudy DP-203 dumps from Cloud Storage: https://drive.google.com/open?id=1_m0vRwpxtk3bMileVoFTIQQEDKhpR8pp

The Data Engineering on Microsoft Azure (DP-203) exam preparation material is available in three different formats for the customers. The formats are PDF format, web-based software, and Microsoft DP-203 desktop practice exam software. The portable PDF format means customers can access real Data Engineering on Microsoft Azure (DP-203) exam questions on their smartphones, tablets, and laptops. The PDF format can be printed and customers can also make proper DP-203 exam notes.

The DP-203 Exam covers a broad range of topics, including data storage, data processing, data integration, and data security. Successful candidates will be able to design and implement data solutions using Azure services such as Azure SQL Database, Azure Cosmos DB, Azure Data Factory, and Azure Stream Analytics. They will also be able to implement advanced analytics solutions using Azure Databricks and Azure Synapse Analytics.

>> DP-203 Real Questions <<

Reliable DP-203 Exam Price & Sure DP-203 Pass

You can save a lot of time for collecting real-time information if you choose our DP-203 study guide. Because our professionals have done all of these collections for you and they are more specialized in the field. So the keypoints are all contained in the DP-203 Exam Questions. Besides, in order to ensure that you can see the updated DP-203 practice prep as soon as possible, our system will send the updated information to your email address as soon as possible.

Microsoft DP-203 (Data Engineering on Microsoft Azure) Certification Exam is one of the most popular certification exams in the data engineering domain. It is designed for professionals who want to validate their skills in working with data engineering technologies on the Microsoft Azure platform. Data Engineering on Microsoft Azure certification exam is intended to showcase the proficiency of candidates in designing, building, and maintaining data processing systems on Azure.

To pass the DP-203 Exam, candidates must demonstrate their expertise in designing and implementing data solutions using Azure data services. They must also have a solid understanding of data management principles and be able to apply them in a cloud environment. Additionally, candidates must have experience working with programming languages such as SQL, Python, and PowerShell.

Microsoft Data Engineering on Microsoft Azure Sample Questions (Q100-Q105):

NEW QUESTION # 100
You have an Azure Data Lake Storage account named account1.
You use an Azure Synapse Analytics serverless SQL pool to access sales data stored in account1.
You need to create a bar chart that displays sales by product. The solution must minimize development effort.
In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order

Answer:

Explanation:
Explanation:
To create a bar chart that displays sales by product while minimizing development effort, follow these steps in the correct order:
Correct Order
* Create a SQL script by using Synapse Studio.
* Start by writing a SQL query in Synapse Studio that retrieves the sales by product data from the data stored in Azure Data Lake.
* Add a SELECT statement that will return the sales by product data.
* Write the query in the SQL script, such as:
sql
Copy
SELECT ProductName, SUM(SalesAmount) AS TotalSales
FROM OPENROWSET(
BULK 'https://account1.dfs.core.windows.net/salesdata/',
FORMAT='CSV',
DATA_SOURCE='MyDataSource'
) AS SalesData
GROUP BY ProductName
ORDER BY TotalSales DESC;
* Execute the script.
* Run the SQL script to retrieve the data and ensure it returns the correct results.
* Switch to the Chart view.
* In Synapse Studio, switch to the Chart view to visualize the query results.
* Modify the Chart settings.
* Customize the chart settings to display a bar chart with "ProductName" on the x-axis and
"TotalSales" on the y-axis.
Explanation:
* The sequence ensures that you first retrieve the data correctly, verify it by executing the query, and then switch to the Chart view to visualize the data with minimal development effort.

NEW QUESTION # 101
You have an Azure Data Lake Storage Gen2 container.
Data is ingested into the container, and then transformed by a data integration application. The data is NOT modified after that.
Users can read files in the container but cannot modify the files.
You need to design a data archiving solution that meets the following requirements:
New data is accessed frequently and must be available as quickly as possible.
Data that is older than five years is accessed infrequently but must be available within one second when requested.
Data that is older than seven years is NOT accessed. After seven years, the data must be persisted at the lowest cost possible.
Costs must be minimized while maintaining the required availability.
How should you manage the data? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point

Answer:

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers Box 1: Replicated Replicated tables are ideal for

small star-schema dimension tables, because the fact table is often distributed on a column that is not compatible with the connected dimension tables. If this case applies to your schema, consider changing small dimension tables currently implemented as round-robin to replicated.
Box 2: Replicated
Box 3: Replicated
Box 4: Hash-distributed
For Fact tables use hash-distribution with clustered columnstore index. Performance improves when two hash tables are joined on the same distribution column.
https://azure.microsoft.com/en-us/updates/reduce-data-movement-and-make-your-queries-more-efficient-with-the-general-availability-of-replicated-tables/
https://azure.microsoft.com/en-us/blog/replicated-tables-now-generally-available-in-azure-sql-data-warehouse/

NEW QUESTION # 102
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 section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You plan to create an Azure Databricks workspace that has a tiered structure. The workspace will contain the following three workloads:
* A workload for data engineers who will use Python and SQL.
* A workload for jobs that will run notebooks that use Python, Scala, and SOL.
* A workload that data scientists will use to perform ad hoc analysis in Scala and R.
The enterprise architecture team at your company identifies the following standards for Databricks environments:
* The data engineers must share a cluster.
* The job cluster will be managed by using a request process whereby data scientists and data engineers provide packaged notebooks for deployment to the cluster.
* All the data scientists must be assigned their own cluster that terminates automatically after 120 minutes of inactivity. Currently, there are three data scientists.
You need to create the Databricks clusters for the workloads.
Solution: You create a Standard cluster for each data scientist, a High Concurrency cluster for the data engineers, and a Standard cluster for the jobs.
Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:
Explanation
We would need a High Concurrency cluster for the jobs.
Note:
Standard clusters are recommended for a single user. Standard can run workloads developed in any language:
Python, R, Scala, and SQL.
A high concurrency cluster is a managed cloud resource. The key benefits of high concurrency clusters are that they provide Apache Spark-native fine-grained sharing for maximum resource utilization and minimum query latencies.
Reference:
https://docs.azuredatabricks.net/clusters/configure.html

NEW QUESTION # 103
You have an Azure Data Factory instance that contains two pipelines named Pipeline1 and Pipeline2.
Pipeline1 has the activities shown in the following exhibit.
Pipeline2 has the activities shown in the following exhibit.
You execute Pipeline2, and Stored procedure1 in Pipeline1 fails.
What is the status of the pipeline runs?

- A. Pipeline1 and Pipeline2 failed.

- B. Pipeline1 succeeded and Pipeline2 failed.
- C. Pipeline1 failed and Pipeline2 succeeded.
- D. Pipeline1 and Pipeline2 succeeded.

Answer: D

Explanation:
Explanation
Activities are linked together via dependencies. A dependency has a condition of one of the following:
Succeeded, Failed, Skipped, or Completed.
Consider Pipeline1:
If we have a pipeline with two activities where Activity2 has a failure dependency on Activity1, the pipeline will not fail just because Activity1 failed. If Activity1 fails and Activity2 succeeds, the pipeline will succeed.
This scenario is treated as a try-catch block by Data Factory.
Waterfall chart Description automatically generated with medium confidence
The failure dependency means this pipeline reports success.
Note:
If we have a pipeline containing Activity1 and Activity2, and Activity2 has a success dependency on Activity1, it will only execute if Activity1 is successful. In this scenario, if Activity1 fails, the pipeline will fail.
Reference:
https://datasavvy.me/category/azure-data-factory/

NEW QUESTION # 104
You have an on-premises data warehouse that includes the following fact tables. Both tables have the following columns: DateKey, ProductKey, RegionKey. There are 120 unique product keys and 65 unique region keys.
Queries that use the data warehouse take a long time to complete.
You plan to migrate the solution to use Azure Synapse Analytics. You need to ensure that the Azure-based solution optimizes query performance and minimizes processing skew.
What should you recommend? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point

Answer:

Explanation:
Reference:
https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-tables-distribute

NEW QUESTION # 105
.....

- DP-203 New Exam Camp □ Latest DP-203 Test Pdf □ DP-203 Latest Braindumps Book □ Go to website （ www.vce4dumps.com ） open and search for □ DP-203 □ to download for free □New DP-203 Exam Name
- Valid Exam DP-203 Practice □ Latest DP-203 Test Pdf □ Latest DP-203 Test Pdf◂ Enter ➤ www.pdfvce.com □ and search for （ DP-203 ） to download for free □Latest DP-203 Test Pdf
- Latest DP-203 Exam Pass4sure □ Test DP-203 Quiz □ DP-203 Latest Braindumps Book □ Easily obtain free download of □ DP-203 □ by searching on 《 www.pdfdumps.com 》 □Valid Exam DP-203 Registration
- www.stes.tyc.edu.tw, myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, building.lv, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, Disposable vapes

P.S. Free 2026 Microsoft DP-203 dumps are available on Google Drive shared by BraindumpStudy: https://drive.google.com/open?id=1_m0vRwpxtk3bMileVoFTIQQEDKhpR8pp