



Free Questions for DP-203 by ebraindumps

Shared by Pugh on 20-10-2022

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You are designing a folder structure for the files in an Azure Data Lake Storage Gen2 account. The account has one container that contains three years of data.

You need to recommend a folder structure that meets the following requirements:

- * Supports partition elimination for queries by Azure Synapse Analytics serverless SQL pool
- * Supports fast data retrieval for data from the current month
- * Simplifies data security management by department

Which folder structure should you recommend?

Options:

- A-** \YYY\MM\DD\Department\DataSource\DataFile_YYMMDD.parquet
- B-** \Department\DataSource\YYY\MM\DataFile_YYMMDD.parquet
- C-** \DD\MM\YYYY\Department\DataSource\DataFile_DDMMYY.parquet
- D-** \DataSource\Department\YYYYMM\DataFile_YYMMDD.parquet

Answer:

C

Question 2

Question Type: MultipleChoice

You have an Azure Data Lake Storage Gen2 account that contains two folders named Folder and Folder2.

You use Azure Data Factory to copy multiple files from Folder1 to Folder2.

```
Operation on target Copy_sks failed: Failure happened on 'Sink' side.  
ErrorCode=DelimitedTextMoreColumnsThanDefined,  
'Type=Microsoft.DataTransfer.Common.Shared.HybridDeliveryException,  
Message=Error found when processing 'Csv/Tsv Format Text' source  
'@_2020_11_09_11_43_32.avro' with row number 53: found more columns  
than expected column count 27., Source=Microsoft.DataTransfer.Common,'
```

You receive the following error.

What should you do to resolve the error.

Options:

- A- Add an explicit mapping.
- B- Enable fault tolerance to skip incompatible rows.
- C- Lower the degree of copy parallelism
- D- Change the Copy activity setting to Binary Copy

Answer:

B

Question 3

Question Type: MultipleChoice

You have an Azure data factory named ADF1.

You currently publish all pipeline authoring changes directly to ADF1.

You need to implement version control for the changes made to pipeline artifacts. The solution must ensure that you can apply version control to the resources currently defined in the UX Authoring canvas for ADF1.

Which two actions should you perform? Each correct answer presents part of the solution

NOTE: Each correct selection is worth one point.

Options:

- A- Create an Azure Data Factory trigger
- B- From the UX Authoring canvas, select Set up code repository
- C- Create a GitHub action
- D- From the UX Authoring canvas, run Publish All.
- E- Create a Git repository
- F- From the UX Authoring canvas, select Publish

Answer:

B, D

Question 4

Question Type: MultipleChoice

You plan to create a dimension table in Azure Synapse Analytics that will be less than 1 GB.

You need to create the table to meet the following requirements:

* Provide the fastest Query time.

* Minimize data movement during queries.

Which type of table should you use?

Options:

A- hash distributed

B- heap

C- replicated

D- round-robin

Answer:

C

Question 5

Question Type: MultipleChoice

You have an Azure Synapse Analytics dedicated SQL pool named Pool1. Pool1 contains a table named table1.

You load 5 TB of data into table1.

You need to ensure that columnstore compression is maximized for table1.

Which statement should you execute?

Options:

- A- ALTER INDEX ALL on table1 REORGANIZE
- B- ALTER INDEX ALL on table1 REBUILD
- C- DBCC DBREINDEX (table1)
- D- DBCC INDEXDEFRAG (pool1,table1)

Answer:

C

Question 6

Question Type: MultipleChoice

You have an Azure Synapse Analytics dedicated SQL pool named pool1.

You need to perform a monthly audit of SQL statements that affect sensitive data. The solution must minimize administrative effort.

What should you include in the solution?

Options:

- A- Microsoft Defender for SQL
- B- dynamic data masking
- C- sensitivity labels
- D- workload management

Answer:

B

Question 7

Question Type: MultipleChoice

You have an Azure subscription that contains an Azure Data Lake Storage account named myaccount1. The myaccount1 account contains two containers named container1 and contained. The subscription is linked to an Azure Active Directory (Azure AD) tenant that contains a security group named Group1.

You need to grant Group1 read access to container1. The solution must use the principle of least privilege. Which role should you assign to Group1?

Options:

- A- Storage Blob Data Reader for container1
- B- Storage Table Data Reader for container1
- C- Storage Blob Data Reader for myaccount1
- D- Storage Table Data Reader for myaccount1

Answer:

A

Question 8

Question Type: MultipleChoice

You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and a database named DB1. DB1 contains a fact table named Table1.

You need to identify the extent of the data skew in Table1.

What should you do in Synapse Studio?

Options:

- A- Connect to the built-in pool and query sysdm_pdw_sys_info.
- B- Connect to Pool1 and run DBCC CHECKALLOC.
- C- Connect to the built-in pool and run DBCC CHECKALLOC.
- D- Connect to Pool! and query sys.dm_pdw_nodes_db_partition_stats.

Answer:

D

Explanation:

Microsoft recommends use of sys.dm_pdw_nodes_db_partition_stats to analyze any skewness in the data.

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

Question 9

Question Type: MultipleChoice

You are designing an Azure Synapse solution that will provide a query interface for the data stored in an Azure Storage account. The storage account is only accessible from a virtual network.

You need to recommend an authentication mechanism to ensure that the solution can access the source data.

What should you recommend?

Options:

- A- a managed identity
- B- anonymous public read access
- C- a shared key

Answer:

A

Explanation:

Managed Identity authentication is required when your storage account is attached to a VNet.

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/quickstart-bulk-load-copy-tsql-examples>

Question 10

Question Type: MultipleChoice

You have an Azure Data Lake Storage Gen2 account named adls2 that is protected by a virtual network.

You are designing a SQL pool in Azure Synapse that will use adls2 as a source.

What should you use to authenticate to adls2?

Options:

- A- a shared access signature (SAS)
- B- a managed identity
- C- a shared key
- D- an Azure Active Directory (Azure AD) user

Answer:

B

Explanation:

Managed identity for Azure resources is a feature of Azure Active Directory. The feature provides Azure services with an automatically managed identity in Azure AD. You can use the Managed Identity capability to authenticate to any service that support Azure AD authentication.

Managed Identity authentication is required when your storage account is attached to a VNet.

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/quickstart-bulk-load-copy-tsql-examples>

To Get Premium Files for DP-203 Visit

<https://www.p2pexams.com/products/dp-203>

For More Free Questions Visit

<https://www.p2pexams.com/microsoft/pdf/dp-203>

