



Free Questions for DP-203 by go4braindumps

Shared by Herrera on 06-06-2022

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You have an Azure subscription linked to an Azure Active Directory (Azure AD) tenant that contains a service principal named ServicePrincipal1. The subscription contains an Azure Data Lake Storage account named adls1. Adls1 contains a folder named Folder2 that has a URI of <https://adls1.dfs.core.windows.net/container1/Folder1/Folder2/>.

ServicePrincipal1 has the access control list (ACL) permissions shown in the following table.

Resource	Permission
container1	Access – Execute
Folder1	Access – Execute
Folder2	Access – Read

You need to ensure that ServicePrincipal1 can perform the following actions:

Traverse child items that are created in Folder2.

Read files that are created in Folder2.

The solution must use the principle of least privilege.

Which two permissions should you grant to ServicePrincipal1 for Folder2? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

Options:

A- Access - Read

B- Access - Write

C- Access - Execute

D- Default-Read

E- Default - Write

F- Default - Execute

Answer:

D, F

Explanation:

Execute (X) permission is required to traverse the child items of a folder.

There are two kinds of access control lists (ACLs), Access ACLs and Default ACLs.

Access ACLs: These control access to an object. Files and folders both have Access ACLs.

Default ACLs: A 'template' of ACLs associated with a folder that determine the Access ACLs for any child items that are created under that folder. Files do not have Default ACLs.

<https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-access-control>

Question 2

Question Type: MultipleChoice

You are designing database for an Azure Synapse Analytics dedicated SQL pool to support workloads for detecting ecommerce transaction fraud.

Data will be combined from multiple ecommerce sites and can include sensitive financial information such as credit card numbers.

You need to recommend a solution that meets the following requirements:

Users must be able to identify potentially fraudulent transactions.

Users must be able to use credit cards as a potential feature in models.

Users must NOT be able to access the actual credit card numbers.

What should you include in the recommendation?

Options:

- A- Transparent Data Encryption (TDE)
- B- row-level security (RLS)
- C- column-level encryption
- D- Azure Active Directory (Azure AD) pass-through authentication

Answer:

C

Explanation:

Use Always Encrypted to secure the required columns. You can configure Always Encrypted for individual database columns containing your sensitive data. Always Encrypted is a feature designed to protect sensitive data, such as credit card numbers or national identification numbers (for example, U.S. social security numbers), stored in Azure SQL Database or SQL Server databases.

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

Question 3

Question Type: MultipleChoice

You have an activity in an Azure Data Factory pipeline. The activity calls a stored procedure in a data warehouse in Azure Synapse Analytics and runs daily.

You need to verify the duration of the activity when it ran last.

What should you use?

Options:

- A- activity runs in Azure Monitor
- B- Activity log in Azure Synapse Analytics
- C- the sys.dm_pdw_wait_stats data management view in Azure Synapse Analytics
- D- an Azure Resource Manager template

Answer:

A

Explanation:

<https://docs.microsoft.com/en-us/azure/data-factory/monitor-visually>

Question 4

Question Type: MultipleChoice

A company uses Azure Stream Analytics to monitor devices.

The company plans to double the number of devices that are monitored.

You need to monitor a Stream Analytics job to ensure that there are enough processing resources to handle the additional load.

Which metric should you monitor?

Options:

A- Early Input Events

B- Late Input Events

C- Watermark delay

D- Input Deserialization Errors

Answer:

C

Explanation:

There are a number of resource constraints that can cause the streaming pipeline to slow down. The watermark delay metric can rise due to:

Not enough processing resources in Stream Analytics to handle the volume of input events.

Not enough throughput within the input event brokers, so they are throttled.

Output sinks are not provisioned with enough capacity, so they are throttled. The possible solutions vary widely based on the flavor of output service being used.

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-time-handling>

Question 5

Question Type: MultipleChoice

You are designing an anomaly detection solution for streaming data from an Azure IoT hub. The solution must meet the following requirements:

Send the output to Azure Synapse.

Identify spikes and dips in time series data.

Minimize development and configuration effort.

Which should you include in the solution?

Options:

A- Azure Databricks

B- Azure Stream Analytics

C- Azure SQL Database

Answer:

B

Explanation:

You can identify anomalies by routing data via IoT Hub to a built-in ML model in Azure Stream Analytics.

<https://docs.microsoft.com/en-us/learn/modules/data-anomaly-detection-using-azure-iot-hub/>

Question 6

Question Type: MultipleChoice

You have an enterprise data warehouse in Azure Synapse Analytics named DW1 on a server named Server1.

You need to determine the size of the transaction log file for each distribution of DW1.

What should you do?

Options:

A- On DW1, execute a query against the sys.database_files dynamic management view.

B- From Azure Monitor in the Azure portal, execute a query against the logs of DW1.

C- Execute a query against the logs of DW1 by using the Get-AzOperationalInsightsSearchResult PowerShell cmdlet.

D- On the master database, execute a query against the sys.dm_pdw_nodes_os_performance_counters dynamic management view.

Answer:

A

Explanation:

For information about the current log file size, its maximum size, and the autogrow option for the file, you can also use the size, max_size, and growth columns for that log file in sys.database_files.

<https://docs.microsoft.com/en-us/sql/relational-databases/logs/manage-the-size-of-the-transaction-log-file>

Question 7

Question Type: MultipleChoice

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 have an Azure Data Lake Storage account that contains a staging zone.

You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.

Solution: You schedule an Azure Databricks job that executes an R notebook, and then inserts the data into the data warehouse.

Does this meet the goal?

Options:

A- Yes

B- No

Answer:

B

Explanation:

Must use an Azure Data Factory, not an Azure Databricks job.

<https://docs.microsoft.com/en-US/azure/data-factory/transform-data>

Question 8

Question Type: MultipleChoice

You have an Azure subscription that contains an Azure Blob Storage account named storage1 and an Azure Synapse Analytics dedicated SQL pool named Pool1.

You need to store data in storage1. The data will be read by Pool1. The solution must meet the following requirements:

Enable Pool1 to skip columns and rows that are unnecessary in a query.

Automatically create column statistics.

Minimize the size of files.

Which type of file should you use?

Options:

A- JSON

B- Parquet

C- Avro

D- CSV

Answer:

B

Explanation:

Automatic creation of statistics is turned on for Parquet files. For CSV files, you need to create statistics manually until automatic creation of CSV files statistics is supported.

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-statistics>

Question 9

Question Type: MultipleChoice

You have two fact tables named Flight and Weather. Queries targeting the tables will be based on the join between the following columns.

Table	Column
Flight	ArrivalAirportID ArrivalDateTime
Weather	AirportID ReportDateTime

You need to recommend a solution that maximizes query performance.

What should you include in the recommendation?

Options:

- A- In the tables use a hash distribution of ArrivalDateTime and ReportDateTime.
- B- In the tables use a hash distribution of ArrivalAirportID and AirportID.
- C- In each table, create an identity column.
- D- In each table, create a column as a composite of the other two columns in the table.

Answer:

B

Explanation:

Hash-distribution improves query performance on large fact tables.

Incorrect Answers:

A: Do not use a date column for hash distribution. All data for the same date lands in the same distribution. If several users are all filtering on the same date, then only 1 of the 60 distributions do all the processing work.

Question 10

Question Type: MultipleChoice

You are designing a highly available Azure Data Lake Storage solution that will include geo-zone-redundant storage (GZRS).

You need to monitor for replication delays that can affect the recovery point objective (RPO).

What should you include in the monitoring solution?

Options:

- A- availability
- B- Average Success E2E Latency
- C- 5xx: Server Error errors
- D- Last Sync Time

Answer:

D

Explanation:

Because geo-replication is asynchronous, it is possible that data written to the primary region has not yet been written to the secondary region at the time an outage occurs. The Last Sync Time property indicates the last time that data from the primary region was written successfully to the secondary region. All writes made to the primary region before the last sync time are available to be read from the secondary location. Writes made to the primary region after the last sync time property may or may not be available for reads yet.

<https://docs.microsoft.com/en-us/azure/storage/common/last-sync-time-get>

Question 11

Question Type: MultipleChoice

You have an Azure Databricks resource.

You need to log actions that relate to changes in compute for the Databricks resource.

Which Databricks services should you log?

Options:

A- clusters

B- workspace

C- DBFS

D- SSH

E lobs

Answer:

B

Explanation:

Cloud Provider Infrastructure Logs. Databricks logging allows security and admin teams to demonstrate conformance to data governance standards within or from a Databricks workspace. Customers, especially in the regulated industries, also need records on activities like: -- User access control to cloud data storage -- Cloud Identity and Access Management roles -- User access to cloud network and compute

Azure Databricks offers three distinct workloads on several VM Instances tailored for your data analytics workflow---the Jobs Compute and Jobs Light Compute workloads make it easy for data engineers to build and execute jobs, and the All-Purpose Compute workload makes it easy for data scientists to explore, visualize, manipulate, and share data and insights interactively.

Question 12

Question Type: MultipleChoice

You are developing an application that uses Azure Data Lake Storage Gen 2.

You need to recommend a solution to grant permissions to a specific application for a limited time period.

What should you include in the recommendation?

Options:

A- Azure Active Directory (Azure AD) identities

B- shared access signatures (SAS)

C- account keys

D- role assignments

Answer:

B

Explanation:

A shared access signature (SAS) provides secure delegated access to resources in your storage account. With a SAS, you have granular control over how a client can access your data. For example:

What resources the client may access.

What permissions they have to those resources.

How long the SAS is valid.

<https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview>

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>

