

➤ **Vendor: Microsoft**

➤ **Exam Code: DP-200**

➤ **Exam Name: Implementing an Azure Data Solution**

➤ **New Updated Questions from [Braindump2go](#) (Updated in [May/2020](#))**

Visit Braindump2go and Download Full Version DP-200 Exam Dumps

QUESTION 175

You have to access Azure Blob Storage from Azure Databricks using secrets stored in a key vault. You already have the storage account, the blob container and Azure key vault in place.

You decide to implement the following steps:

- Add the secret to the key vault
- Create a Databricks workspace and add the access keys
- Access the blob container from Azure Databricks

Would these steps fulfil the requirement?

- A. Yes
- B. No

Answer: B

Explanation:

You are supposed to add a secret scope to the Databricks workspace and not the access keys.
<https://docs.microsoft.com/en-us/azure/azure-databricks/store-secrets-azure-key-vault>

QUESTION 176

A company is planning on creating an Azure SQL database to support a mission critical application. The application needs to be highly available and not have any performance degradation during maintenance windows. Which of the following technologies can be used to implement this solution? (Choose 3)

- A. Premium Service Tier
- B. Virtual Machine Scale Sets
- C. Basic Service Tier
- D. SQL Data Sync
- E. Always On Availability Groups
- F. Zone-redundant configuration

Answer: AEF

Explanation:

Option B is incorrect since this is used to scale virtual machines based on the load or demand.

Option C is incorrect since this service tier does not provide high availability options.

Option D is incorrect since this feature is used to keep multiple databases in sync.

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-high-availability>

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sql/virtual-machines-windows-portal-sql-availability-group-tutorial>

QUESTION 177

A company has an Azure SQL data warehouse. They want to use PolyBase to retrieve data from an Azure Blob storage account and ingest into the Azure SQL data warehouse. The files are stored in parquet format. The data needs to be loaded into a table called lead2pass_sales.

Which of the following actions need to be performed to implement this requirement? (Choose 4)

[DP-200 Exam Dumps](#) [DP-200 Exam Questions](#) [DP-200 PDF Dumps](#) [DP-200 VCE Dumps](#)

<https://www.braindump2go.com/dp-200.html>

- A. Create an external file format that would map to the parquet-based files
- B. Load the data into a staging table
- C. Create an external table called lead2pass_sales_details
- D. Create an external data source for the Azure Blob storage account
- E. Create a master key on the database
- F. Configure Polybase to use the Azure Blob storage account

Answer: BCDE

Explanation:

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/sql-data-warehouse/load-data-from-azure-blob-storage-using-polybase.md>

QUESTION 178

You have to deploy resources on Azure HD Insight for a batch processing job. The batch processing must run daily and must scale to minimize costs. You also be able to monitor cluster performance.

You need to decide on a tool that will monitor the clusters and provide information on suggestions on how to scale.

You decide to download the Azure HDInsight cluster logs by using Azure PowerShell.

Would this fulfil the requirement?

- A. Yes
- B. No

Answer: B

Explanation:

This will not give you a complete picture and give you the ability to decide on how to scale the cluster.

You have to use Azure HDInsight cluster management solutions

<https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-hadoop-oms-log-analytics-tutorial#install-hdinsight-cluster-management-solutions>

QUESTION 179

A company has a storage account named whizlabstore2020. They want to ensure that they can recover a blob object if it was deleted in the last 10 days. Which of the following would they implement for this requirement?

- A. Firewalls and virtual networks
- B. CORS
- C. Soft Delete
- D. Access Keys

Answer: C

Explanation:

They can use the Soft Delete feature.

How soft delete works

When enabled, soft delete enables you to save and recover your data when blobs or blob snapshots are deleted. This protection extends to blob data that is erased as the result of an overwrite.

When data is deleted, it transitions to a soft deleted state instead of being permanently erased. When soft delete is on and you overwrite data, a soft deleted snapshot is generated to save the state of the overwritten data. Soft deleted objects are invisible unless explicitly listed. You can configure the amount of time soft deleted data is recoverable before it is permanently expired.

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-soft-delete?tabs=azure-portal>

QUESTION 180

A company wants to use the Azure Databricks service. There is need to create clusters based on the following configuration

- ClusterA - Here the cluster need to be configured to terminate automatically after 120 minutes
- ClusterB - Here an environment needs to be created for each notebook
- ClusterC - Here a group of data engineers will be sharing the same cluster

Which of the following cluster type would you set for ClusterB?

- A. Standard
- B. Basic
- C. Job
- D. High concurrency

Answer: D

Explanation:

Here you have to use the cluster type as "High concurrency".

High concurrency clusters

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. This sharing is accomplished with:

- **Preemption:** Proactively preempts Spark tasks from over-committed users to ensure all users get their fair share of cluster time and their jobs complete in a timely manner even when contending with dozens of other users. This uses Spark [Task Preemption for High Concurrency](#).
- **Fault isolation:** Creates an environment for each notebook, effectively isolating them from one another.

<https://docs.microsoft.com/en-us/azure/databricks/clusters/configure>

QUESTION 181

You need to design the disaster recovery solution for customer sales data analytics.

Which three actions should you recommend? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Provision multiple Azure Databricks workspaces in separate Azure regions.
- B. Migrate users, notebooks, and cluster configurations from one workspace to another in the same region.
- C. Use zone redundant storage.
- D. Migrate users, notebooks, and cluster configurations from one region to another.
- E. Use Geo-redundant storage.
- F. Provision a second Azure Databricks workspace in the same region.

Answer: ADE

Explanation:

Scenario: The analytics solution for customer sales data must be available during a regional outage.

To create your own regional disaster recovery topology for databricks, follow these requirements:

1. Provision multiple Azure Databricks workspaces in separate Azure regions
2. Use Geo-redundant storage.

3. Once the secondary region is created, you must migrate the users, user folders, notebooks, cluster configuration, jobs configuration, libraries, storage, init scripts, and reconfigure access control.

Note: Geo-redundant storage (GRS) is designed to provide at least 99.99999999999999% (16 9's) durability of objects over a given year by replicating your data to a secondary region that is hundreds of miles away from the primary region. If your storage account has GRS enabled, then your data is durable even in the case of a complete regional outage or a disaster in which the primary region isn't recoverable.

References:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy-grs>

QUESTION 182

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 are designing an HDInsight/Hadoop cluster solution that uses Azure Data Lake Gen1 Storage.

The solution requires POSIX permissions and enables diagnostics logging for auditing.

You need to recommend solutions that optimize storage.

Proposed Solution: Ensure that files stored are larger than 250MB.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

Depending on what services and workloads are using the data, a good size to consider for files is 256 MB or greater. If the file sizes cannot be batched when landing in Data Lake Storage Gen1, you can have a separate compaction job that combines these files into larger ones.

Note: POSIX permissions and auditing in Data Lake Storage Gen1 comes with an overhead that becomes apparent when working with numerous small files. As a best practice, you must batch your data into larger files versus writing thousands or millions of small files to Data Lake Storage Gen1. Avoiding small file sizes can have multiple benefits, such as:

Lowering the authentication checks across multiple files
Reduced open file connections

Faster copying/replication

Fewer files to process when updating Data Lake Storage Gen1 POSIX permissions

References:
<https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-best-practices>

QUESTION 183

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.

[DP-200 Exam Dumps](#) [DP-200 Exam Questions](#) [DP-200 PDF Dumps](#) [DP-200 VCE Dumps](#)

<https://www.braindump2go.com/dp-200.html>

You are designing an HDInsight/Hadoop cluster solution that uses Azure Data Lake Gen1 Storage. The solution requires POSIX permissions and enables diagnostics logging for auditing. You need to recommend solutions that optimize storage. Proposed Solution: Implement compaction jobs to combine small files into larger files. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

Depending on what services and workloads are using the data, a good size to consider for files is 256 MB or greater. If the file sizes cannot be batched when landing in Data Lake Storage Gen1, you can have a separate compaction job that combines these files into larger ones.

Note: POSIX permissions and auditing in Data Lake Storage Gen1 comes with an overhead that becomes apparent when working with numerous small files. As a best practice, you must batch your data into larger files versus writing thousands or millions of small files to Data Lake Storage Gen1. Avoiding small file sizes can have multiple benefits, such as:

Lowering the authentication checks across multiple files
Reduced open file connections

Faster copying/replication

Fewer files to process when updating Data Lake Storage Gen1 POSIX permissions

References:

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

QUESTION 184

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 are designing an HDInsight/Hadoop cluster solution that uses Azure Data Lake Gen1 Storage. The solution requires POSIX permissions and enables diagnostics logging for auditing.

You need to recommend solutions that optimize storage.

Proposed Solution: Ensure that files stored are smaller than 250MB.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Ensure that files stored are larger, not smaller than 250MB.

You can have a separate compaction job that combines these files into larger ones.

Note: The file POSIX permissions and auditing in Data Lake Storage Gen1 comes with an overhead that becomes apparent when working with numerous small files. As a best practice, you must batch your data into larger files versus writing thousands or millions of small files to Data Lake Storage Gen1. Avoiding small file sizes can have multiple benefits, such as:

Lowering the authentication checks across multiple files
Reduced open file connections

Faster copying/replication

Fewer files to process when updating Data Lake Storage Gen1 POSIX permissions

References:

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

QUESTION 185

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.

[DP-200 Exam Dumps](#) [DP-200 Exam Questions](#) [DP-200 PDF Dumps](#) [DP-200 VCE Dumps](#)

<https://www.braindump2go.com/dp-200.html>

You are designing an Azure SQL Database that will use elastic pools. You plan to store data about customers in a table. Each record uses a value for CustomerID.

You need to recommend a strategy to partition data based on values in CustomerID.

Proposed Solution: Separate data into customer regions by using vertical partitioning.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Vertical partitioning is used for cross-database queries. Instead we should use Horizontal Partitioning, which also is called sharding.

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-query-overview>