

➤ **Vendor:** Microsoft

➤ **Exam Code:** DP-201

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

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

**[Visit Braindump2go and Download Full Version DP-201 Exam Dumps](#)**

**QUESTION 104**

You have a MongoDB database that you plan to migrate to an Azure Cosmos DB account that uses the MongoDB API.

During testing, you discover that the migration takes longer than expected.

You need to recommend a solution that will reduce the amount of time it takes to migrate the data.

What are two possible recommendations to achieve this goal? Each correct answer presents a complete solution. NOTE:

Each correct selection is worth one point.

- A. Increase the Request Units (RUs).
- B. Turn off indexing.
- C. Add a write region.
- D. Create unique indexes.
- E. Create compound indexes.

**Correct Answer:** AB

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**Explanation:**

A: Increase the throughput during the migration by increasing the Request Units (RUs).

For customers that are migrating many collections within a database, it is strongly recommend to configure database-level throughput. You must make this choice when you create the database. The minimum database-level throughput capacity is 400 RU/sec. Each collection sharing database-level throughput requires at least 100 RU/sec.

B: By default, Azure Cosmos DB indexes all your data fields upon ingestion. You can modify the indexing policy in Azure Cosmos DB at any time. In fact, it is often recommended to turn off indexing when migrating data, and then turn it back on when the data is already in Cosmos DB.

References:

<https://docs.microsoft.com/bs-latn-ba/Azure/cosmos-db/mongodb-pre-migration>

**QUESTION 105**

You need to recommend a storage solution for a sales system that will receive thousands of small files per minute. The files will be in JSON, text, and CSV formats. The files will be processed and transformed before they are loaded into an Azure data warehouse. The files must be stored and secured in folders.

Which storage solution should you recommend?

- A. Azure Data Lake Storage Gen2
- B. Azure Cosmos DB
- C. Azure SQL Database
- D. Azure Blob storage

**Correct Answer:** A **Section:**

**(none) Explanation**

**Explanation/Reference:**

**Explanation:**

Azure provides several solutions for working with CSV and JSON files, depending on your needs. The primary landing place for these files is either Azure Storage or Azure Data Lake Store.1 Azure Data Lake Storage is an optimized storage for big data analytics workloads.

IncoD: Azure Blob Storage containers is a general purpose object store for a wide variety of storage scenarios. Blobs are stored in containers, which are similar to folders.

References:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/scenarios/csv-and-json>

**QUESTION 106**

You are designing an Azure Cosmos DB database that will support vertices and edges.

Which Cosmos DB API should you include in the design?

- A. SQL
- B. Cassandra
- C. Gremlin
- D. Table

**Correct Answer:** C **Section:**

**(none) Explanation**

**Explanation/Reference:**

**Explanation:**

The Azure Cosmos DB Gremlin API can be used to store massive graphs with billions of vertices and edges.

References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/graph-introduction>

**[DP-201 Exam Dumps](#)** **[DP-201 Exam Questions](#)** **[DP-201 PDF Dumps](#)** **[DP-201 VCE Dumps](#)**

**<https://www.braindump2go.com/dp-201.html>**

**QUESTION 107**

You are designing a big data storage solution. The solution must meet the following requirements:

- Provide unlimited account sizes.
- Support a hierarchical file system.
- Be optimized for parallel analytics workloads.

Which storage solution should you use?

- A. Azure Data Lake Storage Gen2
- B. Azure Blob storage
- C. Apache HBase in Azure HDInsight
- D. Azure Cosmos DB

**Correct Answer: A Section:**  
**(none) Explanation**

**Explanation/Reference:****Explanation:**

Azure Data Lake Storage is optimized performance for parallel analytics workloads A key mechanism that allows Azure Data Lake Storage Gen2 to provide file system performance at object storage scale and prices is the addition of a hierarchical namespace. This allows the collection of objects/ files within an account to be organized into a hierarchy of directories and nested subdirectories in the same way that the file system on your computer is organized.

References:

<https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-namespace>

**QUESTION 108**

**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 store delimited text files in an Azure Data Lake Storage account that will be organized into department folders. You

need to configure data access so that users see only the files in their respective department folder.

Solution: From the storage account, you enable a hierarchical namespace, and you use RBAC.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B Section:**  
**(none) Explanation**

**Explanation/Reference:****Explanation:**

Disable the hierarchical namespace. And instead of RBAC use access control lists (ACLs).

Note: Azure Data Lake Storage implements an access control model that derives from HDFS, which in turn derives from the POSIX access control model. Blob container ACLs does not support the hierarchical namespace, so it must be disabled.

References:

<https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-known-issues>

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

**QUESTION 109**

**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 store delimited text files in an Azure Data Lake Storage account that will be organized into department folders. You

need to configure data access so that users see only the files in their respective department folder.

Solution: From the storage account, you disable a hierarchical namespace, and you use RBAC.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B Section:**  
**(none) Explanation**

**Explanation/Reference:****Explanation:**

Instead of RBAC use access control lists (ACLs).

Note: Azure Data Lake Storage implements an access control model that derives from HDFS, which in turn derives from the POSIX access control model. Blob container ACLs does not support the hierarchical namespace, so it must be disabled.

References:

<https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-known-issues>

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

**QUESTION 110**

**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-201 Exam Dumps](#) [DP-201 Exam Questions](#) [DP-201 PDF Dumps](#) [DP-201 VCE Dumps](#)**

**<https://www.braindump2go.com/dp-201.html>**

You plan to store delimited text files in an Azure Data Lake Storage account that will be organized into department folders. You need to configure data access so that users see only the files in their respective department folder.

Solution: From the storage account, you disable a hierarchical namespace, and you use access control lists (ACLs). Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: A Section:**  
**(none) Explanation**

**Explanation/Reference:**

**Explanation:**

Azure Data Lake Storage implements an access control model that derives from HDFS, which in turn derives from the POSIX access control model. Blob container ACLs does not support the hierarchical namespace, so it must be disabled.

References:

<https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-known-issues>

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

#### **QUESTION 111**

You plan to store 100 GB of data used by a line-of-business (LOB) app.

You need to recommend a data storage solution for the data. The solution must meet the following requirements:

- Minimize storage costs.
- Natively support relational queries.
- Provide a recovery time objective (RTO) of less than one minute.

What should you include in the recommendation?

- A. Azure Cosmos DB
- B. Azure SQL Database
- C. Azure SQL Data Warehouse
- D. Azure Blob storage

**Correct Answer: D Section:**  
**(none) Explanation**

**Explanation/Reference:**

**Explanation:**

Incorrect Answers:

A: Azure Cosmos DB would require an SQL API.

#### **QUESTION 112**

You are designing a data storage solution for a database that is expected to grow to 50 TB. The usage pattern is singleton inserts, singleton updates, and reporting.

Which storage solution should you use?

- A. Azure SQL Database elastic pools
- B. Azure SQL Data Warehouse
- C. Azure Cosmos DB that uses the Gremlin API
- D. Azure SQL Database Hyperscale

**Correct Answer: D Section:**  
**(none) Explanation**

**Explanation/Reference:**

**Explanation:**

A Hyperscale database is an Azure SQL database in the Hyperscale service tier that is backed by the Hyperscale scale-out storage technology. A Hyperscale

database supports up to 100 TB of data and provides high throughput and performance, as well as rapid scaling to adapt to the workload requirements.

Scaling is transparent to the application connectivity, query processing, etc. work like any other Azure SQL database.

Incorrect Answers:

A: SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single Azure SQL Database server and share a set number of resources at a set price. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.

B: Rather than SQL Data Warehouse, consider other options for operational (OLTP) workloads that have large numbers of singleton selects.

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-service-tier-hyperscale-faq>

#### **QUESTION 113**

A company plans to use Apache Spark Analytics to analyze intrusion detection data

You need to recommend a solution to monitor network and system activities for malicious activities and policy violations. Reports must be produced in an electronic format and sent to management. The solution must minimize administrative efforts.

What should you recommend?

- A. Azure Data Factory
- B. Azure Data Lake
- C. Azure Databricks
- D. Azure HDInsight

**Correct Answer: D Section:**  
**(none) Explanation**

**Explanation/Reference:**

**Explanation:**

With Azure HDInsight you can set up Azure Monitor alerts that will trigger when the value of a metric or the results of a query meet certain conditions. You can condition on a query returning a record with a value that is greater than or less than a certain threshold, or even on the number of results returned by a query. For example, you could create an alert to send an email if a Spark job fails or if a Kafka disk usage becomes over 90 percent full.

References:

<https://azure.microsoft.com/en-us/blog/monitoring-on-azure-hdinsight-part-4-workload-metrics-and-logs/>

**QUESTION 114**

You are designing an Azure Databricks interactive cluster. The cluster will be used infrequently and will be configured for auto-termination.

You need to ensure that the cluster configuration is retained indefinitely after the cluster is terminated. The solution must minimize costs.

What should you do?

- A. Clone the cluster after it is terminated.
- B. Terminate the cluster manually when processing completes.
- C. Create an Azure runbook that starts the cluster every 90 days.
- D. Pin the cluster.

**Correct Answer: D Section:**  
**(none) Explanation**

**Explanation/Reference:**

**Explanation:**

To keep an interactive cluster configuration even after it has been terminated for more than 30 days, an administrator can pin a cluster to the cluster list.

References:

<https://docs.azuredatabricks.net/clusters/clusters-manage.html#automatic-termination>