

QUESTION 89

You plan to deploy an Azure Cosmos DB database that supports multi-master replication.

You need to select a consistency level for the database to meet the following requirements:

- Provide a recovery point objective (RPO) of less than 15 minutes.
- Provide a recovery time objective (RTO) of zero minutes.

What are three possible consistency levels that you can select? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Strong
- B. Bounded Staleness
- C. Eventual
- D. Session
- E. Consistent Prefix

Correct Answer: CDE

Explanation:

[Download Full Version DP-200 Exam Dumps\(Updated in Feb/2023\)](#)

Region(s)	Replication mode	Consistency level	RPO	RTO
1	Single or Multi-Master	Any Consistency Level	< 240 Minutes	<1 Week
>1	Single Master	Session, Consistent Prefix, Eventual	< 15 minutes	< 15 minutes
>1	Single Master	Bounded Staleness	K & T	< 15 minutes
>1	Single Master	Strong	0	< 15 minutes
>1	Multi-Master	Session, Consistent Prefix, Eventual	< 15 minutes	0
>1	Multi-Master	Bounded Staleness	K & T	0

<https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels-choosing>

QUESTION 90

You are creating a new notebook in Azure Databricks that will support R as the primary language but will also support Scala and SQL. Which switch should you use to switch between languages?

- A. %<language>
- B. \[<language>]
- C. \(<language>)
- D. @<Language>

Correct Answer: A

Explanation:


You can override the primary language by specifying the language magic command %<language> at the beginning of a cell. The supported magic commands are: %python, %r, %scala, and %sql.

References:

<https://docs.databricks.com/user-guide/notebooks/notebook-use.html#mix-languages>

QUESTION 91

[DP-200 Exam Dumps](#) [DP-200 PDF Dumps](#) [DP-200 VCE Dumps](#) [DP-200 Q&As](#)
<https://www.ensurepass.com/DP-200.html>



Use the following login credentials as needed:

Azure Username: xxxxx

Azure Password: xxxxx

The following information is for technical support purposes only:

Lab Instance: 10277521

You need to ensure that you can recover any blob data from an Azure Storage account named storage 10277521 up to 30 days after the data is deleted.

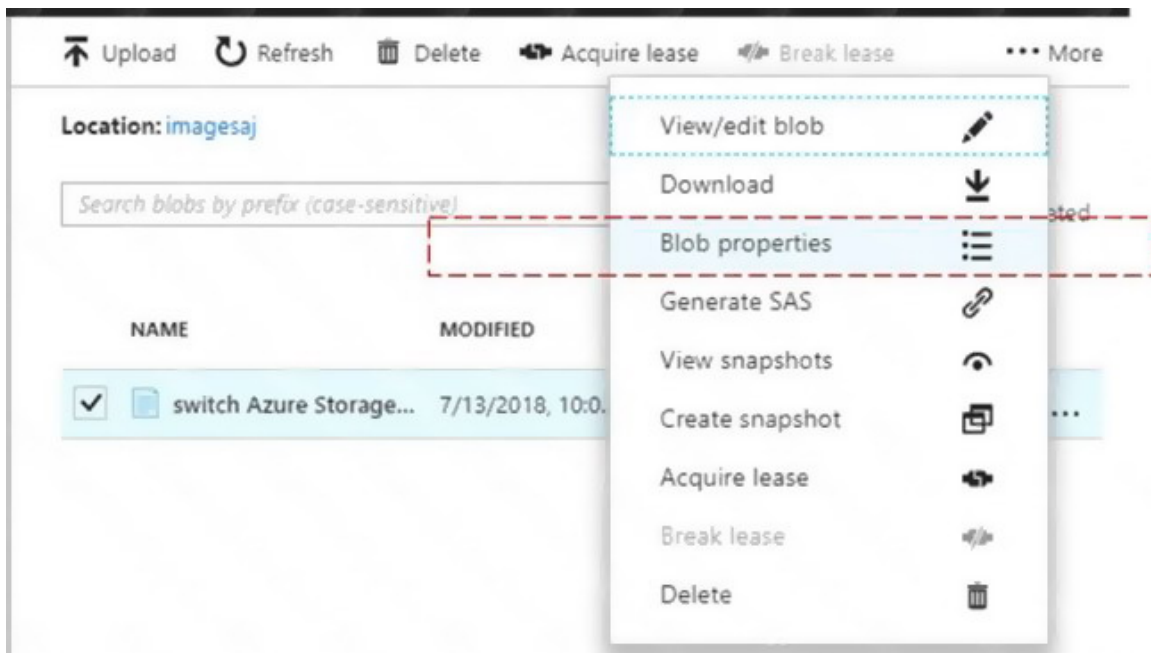
To complete this task, sign in to the Azure portal.

Correct Answer: See the explanation below.

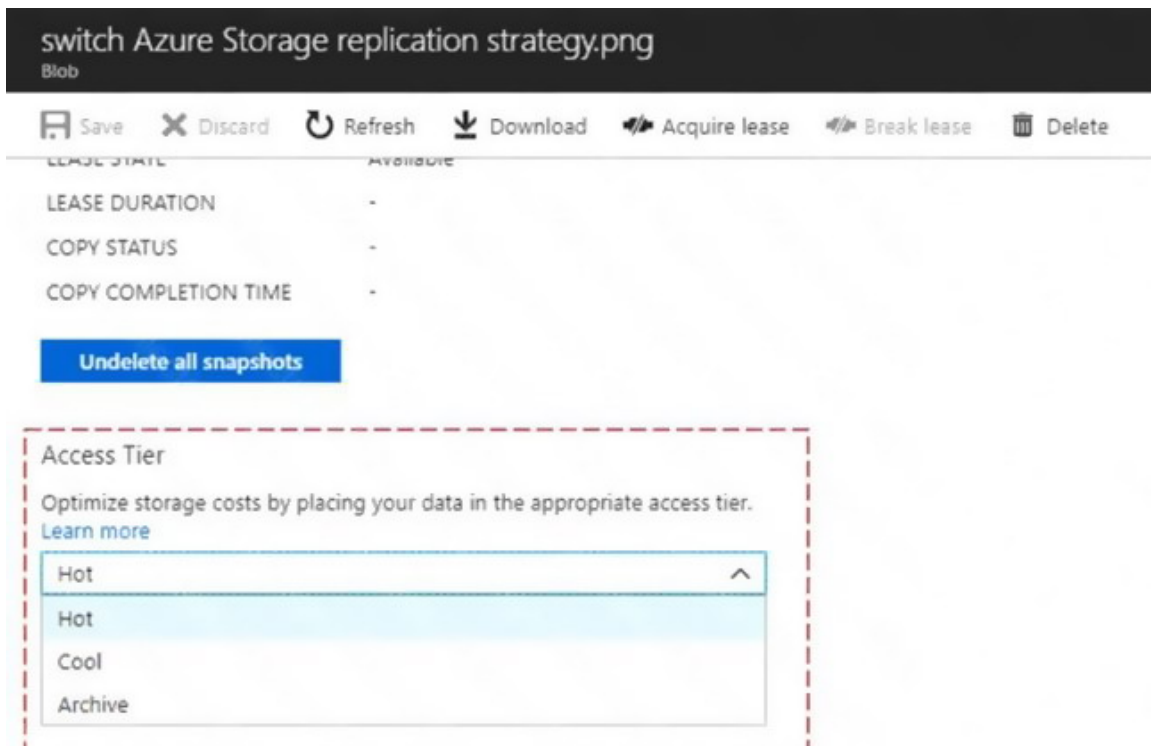
Explanation:

1. Open Azure Portal and open the Azure Blob storage account named storage10277521.
2. Right-click and select Blob properties

[Download Full Version DP-200 Exam Dumps\(Updated in Feb/2023\)](#)



3. From the properties window, change the access tier for the blob to Cool.



Note:

The cool access tier has lower storage costs and higher access costs compared to hot storage. This tier is intended for data that will remain in the cool tier for at least 30 days.

Reference:

[DP-200 Exam Dumps](#) [DP-200 PDF Dumps](#) [DP-200 VCE Dumps](#) [DP-200 Q&As](#)

<https://www.ensurepass.com/DP-200.html>

[Download Full Version DP-200 Exam Dumps\(Updated in Feb/2023\)](#)

<https://dailydotnettips.com/how-to-update-access-tier-in-azure-storage-blob-level/>

QUESTION 92

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 scenario, 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, Spark, Scala, and SQL
- 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 Databrick 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

Correct Answer: B

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.

References:

<https://docs.azuredatabricks.net/clusters/configure.html>

QUESTION 93

[DP-200 Exam Dumps](#) [DP-200 PDF Dumps](#) [DP-200 VCE Dumps](#) [DP-200 Q&As](#)

<https://www.ensurepass.com/DP-200.html>