Microsoft DP-420 Mock Exam

Shared by Bailey on 17-06-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You are implementing an Azure Data Factory data flow that will use an Azure Cosmos DB (SQL API) sink to write a dataset. The data flow will use 2,000 Apache Spark partitions.

You need to ensure that the ingestion from each Spark partition is balanced to optimize throughput.

Which sink setting should you configure?

Options:

A- Throughput
B- Write throughput budget
C- Batch size
D- Collection action

Answer:

C

Explanation:

Batch size: An integer that represents how many objects are being written to Cosmos DB collection in each batch. Usually, starting with the default batch size is sufficient. To further tune this value, note:

Cosmos DB limits single request's size to 2MB. The formula is 'Request Size = Single Document Size * Batch Size'. If you hit error saying 'Request size is too large', reduce the batch size value.

The larger the batch size, the better throughput the service can achieve, while make sure you allocate enough RUs to empower your workload.

Incorrect Answers:

A: Throughput: Set an optional value for the number of RUs you'd like to apply to your CosmosDB collection for each execution of this data flow. Minimum is 400.

B: Write throughput budget: An integer that represents the RUs you want to allocate for this Data Flow write operation, out of the total throughput allocated to the collection.

D: Collection action: Determines whether to recreate the destination collection prior to writing.

None: No action will be done to the collection.

Recreate: The collection will get dropped and recreated

Question 2

Question Type: MultipleChoice

You maintain a relational database for a book publisher. The database contains the following tables.

Name	Column
Author	authorId **(primary key)**
	fullname
	address
	contactinfo
Book	bookId **(primary key)**
	isbn
	title
	genre
BookauthorInk	authorId **(foreign key)**
	bookId **(foreign key)**

The most common query lists the books for a given authorId.

You need to develop a non-relational data model for Azure Cosmos DB Core (SQL) API that will replace the relational database. The solution must minimize latency and read operation costs.

What should you include in the solution?

Options:

A- Create a container for Author and a container for Book. In each Author document, embed booked for each book by the author. In each Book document embed author of each author.
B- Create Author, Book, and BookauthorInk documents in the same container.
C- Create a container that contains a document for each Author and a document for each Book.

In each Book document, embed authorId.
D- Create a container for Author and a container for Book. In each Author document and Book document embed the data from BookauthorInk.

Answer:

A

Explanation:

Store multiple entity types in the same container.

Question 3

Question Type: MultipleChoice

You have an Azure Cosmos DB Core (SQL) API account that uses a custom conflict resolution policy. The account has a registered merge procedure that throws a runtime exception.

The runtime exception prevents conflicts from being resolved.

You need to use an Azure function to resolve the conflicts.

What should you use?

Options:

A- a function that pulls items from the conflicts feed and is triggered by a timer trigger
B- a function that receives items pushed from the change feed and is triggered by an Azure Cosmos DB trigger
C- a function that pulls items from the change feed and is triggered by a timer trigger
D- a function that receives items pushed from the conflicts feed and is triggered by an Azure Cosmos DB trigger

Answer:

D

Explanation:

The Azure Cosmos DB Trigger uses the Azure Cosmos DB Change Feed to listen for inserts and

updates across partitions. The change feed publishes inserts and updates, not deletions.

Question 4

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 Cosmos DB Core (SQL) API account named account 1 that uses autoscale throughput.

You need to run an Azure function when the normalized request units per second for a container in account1 exceeds a specific value.

Solution: You configure an Azure Monitor alert to trigger the function.

Does this meet the goal?

Options:

A- Yes
B- No

Answer:

A

Explanation:

You can set up alerts from the Azure Cosmos DB pane or the Azure Monitor service in the Azure portal.

Note: Alerts are used to set up recurring tests to monitor the availability and responsiveness of your Azure Cosmos DB resources. Alerts can send you a notification in the form of an email, or execute an Azure Function when one of your metrics reaches the threshold or if a specific event is logged in the activity log.

Question 5

Question Type: MultipleChoice

You have a container in an Azure Cosmos DB for NoSQL account that stores data about orders.

The following is a sample of an order document.

```
{
  "orderId" : "d4a9179b-5ead-43a3-b851-add9a71ac4b6",
  "customerId" : "f6e39103-bdc7-4346-9cfb-45daa4b2becf",
  "orderDate" : "2022-09-29",
  "orderItems" : [...],
  "total" : 12345
}
```

Documents are up to 2 KB.

You plan to receive one million orders daily.

Customers will frequently view their past order history.

You are the evaluating whether to use order-Date as the partition key.

What are two effects of using order-Date as the partition key? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

Options:

A- You will exceed the maximum number of partition key values.
B- You will exceed the maximum storage per partition.
C- There will always be a hot partition.
D- Queries will run cross-partition.

Answer:

C, D

Question 6

Question Type: MultipleChoice

Your company develops an application named App1 that uses the Azure Cosmos DB SDK and the Eventual consistency level.

App1 queries an Azure Cosmos DB for NoSQL account named account!

You need to identify which consistency level to assign to App1 to meet the following requirements:

* Maximize the throughput of the queries generated by App1 without increasing the number of request units currently used by the queries.

* Provide the highest consistency guarantees.

Which consistency level should you identify?

Options:

A- Strong
B- Bounded Staleness
C- Session
D- Consistent Prefix

Answer:

A