

DP-420^{Q&As}

Designing and Implementing Cloud-Native Applications Using Microsoft Azure Cosmos DB

Pass Microsoft DP-420 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass2lead.com/dp-420.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



QUESTION 1

You have operational data in an Azure Cosmos DB for NoSQL database.

Database users report that the performance of the database degrades significantly when a business analytics team runs large Apache Spark-based queries against the database.

You need to reduce the impact that running the Spark-based queries has on the database users.

What should you implement?

- A. Azure Synapse Link
- B. a default consistency level of Consistent Prefix
- C. a default consistency level of Strong
- D. the Spark connector

Correct Answer: A

QUESTION 2

You are designing an Azure Cosmos DB Core (SQL) API solution to store data from IoT devices. Writes from the devices will occur every second. The following is a sample of the data.

```
{
  "id" : "03c1ca5a-db18-4231-908f-09a9bc7a7c3e",
  "deviceManufacturer" : "Contoso, Ltd",
  "deviceId" : "f460df85-799f-4d58-b051-67561b4993c6",
  "timestamp" : "2021-09-19T13:47:45",
  "sensor1Value" : true,
  "sensor2Value" : "75",
  "sensor3Value" : "4554",
  "sensor4Value" : "454",
  "sensor5Value" : "42128"
}
```

You need to select a partition key that meets the following requirements for writes:

1.
Minimizes the partition skew
2.
Avoids capacity limits

3.

Avoids hot partitions What should you do?

- A. Use timestamp as the partition key.
- B. Create a new synthetic key that contains deviceId and sensor1Value.
- C. Create a new synthetic key that contains deviceId and deviceManufacturer.
- D. Create a new synthetic key that contains deviceId and a random number.

Correct Answer: D

Use a partition key with a random suffix. Distribute the workload more evenly is to append a random number at the end of the partition key value. When you distribute items in this way, you can perform parallel write operations across partitions.

Incorrect Answers:

A: You will also not like to partition the data on "DateTime", because this will create a hot partition. Imagine you have partitioned the data on time, then for a given minute, all the calls will hit one partition. If you need to retrieve the data for a customer, then it will be a fan-out query because data may be distributed on all the partitions.

B: Sensor1Value has only two values.

C: All the devices could have the same manufacturer.

Reference: <https://docs.microsoft.com/en-us/azure/cosmos-db/sql/synthetic-partition-keys>

QUESTION 3

You have a global ecommerce application that stores data in an Azure Cosmos DB for NoSQL account. The account is configured for multi-region writes.

You need to create a stored procedure for a custom conflict resolution policy for a new container.

In the event of a conflict caused by a deletion the deletion must always take priority.

Which parameter should you check in the stored procedure function?

- A. conflictingItems
- B. is Tombstone
- C. existingItem
- D. incomingItem

Correct Answer: B

QUESTION 4

HOTSPOT

You have the indexing policy shown in the following exhibit.

The screenshot shows the SQL Server Enterprise Manager interface. On the left, a tree view shows the hierarchy: Test > Scale > families. The 'families' folder is selected. The main pane shows the 'Indexing Policy' for the 'families' table. The policy is displayed as a JSON object with the following structure:

```
1 {
2   "indexingMode": "consistent",
3   "automatic": true,
4   "includedPaths": [
5     {
6       "path": "/surname/?"
7     }
8   ],
9   "excludedPaths": [
10    {
11      "path": "/*"
12    }
13  ],
14  "compositeIndexes": [
15    [
16      {
17        "path": "/name"
18      },
19      {
20        "path": "/age"
21      }
22    ]
23  ]
24 }
```

Use the drop-down menus to select the answer choice that answers each question based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

When creating a query, which ORDER BY statement will execute successfully?

	▼
ORDER BY c.age ASC, c.name ASC	
ORDER BY c.age DESC, c.name DESC	
ORDER BY c.name ASC, c.age DESC	
ORDER BY c.name DESC, c.age ASC	
ORDER BY c.name DESC, c.age DESC	

During the creation of an item, when will the index update?

	▼
Never	
At a scheduled interval	
At the same time as the item creation	
After the item appears in the change feed	

Correct Answer:

Answer Area

When creating a query, which ORDER BY statement will execute successfully?

	▼
ORDER BY c.age ASC, c.name ASC	
ORDER BY c.age DESC, c.name DESC	
ORDER BY c.name ASC, c.age DESC	
ORDER BY c.name DESC, c.age ASC	
ORDER BY c.name DESC, c.age DESC	

During the creation of an item, when will the index update?

	▼
Never	
At a scheduled interval	
At the same time as the item creation	
After the item appears in the change feed	

Box 1: ORDER BY c.name DESC, c.age DESC

Queries that have an ORDER BY clause with two or more properties require a composite index.

The following considerations are used when using composite indexes for queries with an ORDER BY clause with two or more properties:

1.

If the composite index paths do not match the sequence of the properties in the ORDER BY clause, then the composite index can't support the query.

2.

The order of composite index paths (ascending or descending) should also match the order in the ORDER BY clause.

3.

The composite index also supports an ORDER BY clause with the opposite order on all paths.

Box 2: At the same time as the item creation Azure Cosmos DB supports two indexing modes:

1.

Consistent: The index is updated synchronously as you create, update or delete items. This means that the consistency of your read queries will be the consistency configured for the account.

2.

None: Indexing is disabled on the container.

Reference: <https://docs.microsoft.com/en-us/azure/cosmos-db/index-policy>

QUESTION 5

You have a database named db1 in an Azure Cosmos DB Core (SQL) API account.

You have a third-party application that is exposed through a REST API.

You need to migrate data from the application to a container in db1 on a weekly basis.

What should you use?

- A. Database Migration Assistant
- B. Azure Data Factory
- C. Azure Migrate

Correct Answer: B

You can use Copy Activity in Azure Data Factory to copy data from and to Azure Cosmos DB (SQL API).

The Azure Cosmos DB (SQL API) connector is supported for the following activities:

Copy activity with supported source/sink matrix

Mapping data flow

Lookup activity

Incorrect:

Not A: Azure Migrate provides a centralized hub to assess and migrate on-premises servers, infrastructure, applications, and data to Azure. It assesses on-premises databases and migrate them to Azure SQL Database or to SQL Managed

Instance.

Not C: Data Migration Assistant (DMA) enables you to upgrade to a modern data platform by detecting compatibility issues that can impact database functionality on your new version of SQL Server. It recommends performance and reliability

improvements for your target environment.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-cosmos-db>

[Latest DP-420 Dumps](#)

[DP-420 PDF Dumps](#)

[DP-420 Exam Questions](#)