

## Exam Questions DP-420

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

<https://www.2passeasy.com/dumps/DP-420/>



**NEW QUESTION 1**

- (Exam Topic 1)

You need to identify which connectivity mode to use when implementing App2. The solution must support the planned changes and meet the business requirements.

Which connectivity mode should you identify?

- A. Direct mode over HTTPS
- B. Gateway mode (using HTTPS)
- C. Direct mode over TCP

**Answer: C**

**Explanation:**

Scenario: Develop an app named App2 that will run from the retail stores and query the data in account2. App2 must be limited to a single DNS endpoint when accessing account2.

By using Azure Private Link, you can connect to an Azure Cosmos account via a private endpoint. The private endpoint is a set of private IP addresses in a subnet within your virtual network.

When you're using Private Link with an Azure Cosmos account through a direct mode connection, you can use only the TCP protocol. The HTTP protocol is not currently supported.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-configure-private-endpoints>

**NEW QUESTION 2**

- (Exam Topic 2)

You have a container named container1 in an Azure Cosmos DB Core (SQL) API account. The container1 container has 120 GB of data. The following is a sample of a document in container1.

```
{
  "customerId" : "5425",
  "orderId" : "9d7816e6-f401-42ba-ad05-0e03de35c0b8",
  "orderDate" : "2019-05-03",
  "orderDetails" : []
}
```

The orderId property is used as the partition key.

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

**Answer Area**

Statements	Yes	No
If you run the following query, the query will run as a cross-partition query <pre>SELECT * FROM c where c.orderDate = "2019-05-03"</pre>	<input type="radio"/>	<input type="radio"/>
If you run the following query, the query will run as a cross-partition query <pre>SELECT * FROM c where c.customerId = "5425"</pre>	<input type="radio"/>	<input type="radio"/>
If you run the following query, the query will run as a cross-partition query <pre>SELECT * FROM c where c.orderDate = "2019-05-03" and c.orderId = "9d7816e6-f401-42ba-ad05-0e03de35c0b8"</pre>	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: Yes

Records with different OrderIDs will match.

Box 2: Yes

Records with different OrderIDs will match.

Box 3: No

Only records with one specific OrderId will match

**NEW QUESTION 3**

- (Exam Topic 2)

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 the function to have an Azure CosmosDB trigger. Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

Instead configure an Azure Monitor alert to trigger the function.  
 You can set up alerts from the Azure Cosmos DB pane or the Azure Monitor service in the Azure portal. Reference: <https://docs.microsoft.com/en-us/azure/cosmos-db/create-alerts>

**NEW QUESTION 4**

- (Exam Topic 2)

You have an Azure Cosmos DB Core (SQL) API account that is configured for multi-region writes. The account contains a database that has two containers named container1 and container2.

The following is a sample of a document in container1:

```
{
  "customerId": 1234, "firstName": "John",
  "lastName": "Smith", "policyYear": 2021
}
```

The following is a sample of a document in container2:

```
{
  "gpsId": 1234,
  "latitude": 38.8951,
  "longitude": -77.0364
}
```

You need to configure conflict resolution to meet the following requirements:

For container1 you must resolve conflicts by using the highest value for policyYear.

For container2 you must resolve conflicts by accepting the distance closest to latitude: 40.730610 and longitude: -73.935242.

Administrative effort must be minimized to implement the solution.

What should you configure for each container? To answer, drag the appropriate configurations to the correct containers. Each configuration may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Configurations	Answer Area
Last Write Wins (default) mode	Container1: 
Merge Procedures (custom) mode	Container2: 
An application that reads from the conflicts feed	

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Last Write Wins (LWW) (default) mode

Last Write Wins (LWW): This resolution policy, by default, uses a system-defined timestamp property. It's based on the time-synchronization clock protocol.

Box 2: Merge Procedures (custom) mode

Custom: This resolution policy is designed for application-defined semantics for reconciliation of conflicts. When you set this policy on your Azure Cosmos container, you also need to register a merge stored procedure. This procedure is automatically invoked when conflicts are detected under a database transaction on the server. The system provides exactly once guarantee for the execution of a merge procedure as part of the commitment protocol.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/conflict-resolution-policies> <https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-manage-conflicts>

**NEW QUESTION 5**

- (Exam Topic 2)

You need to configure an Apache Kafka instance to ingest data from an Azure Cosmos DB Core (SQL) API account. The data from a container named telemetry must be added to a Kafka topic named iot. The solution must store the data in a compact binary format.

Which three configuration items should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. "connector.class": "com.azure.cosmos.kafka.connect.source.CosmosDBSourceConnector"
- B. "key.converter": "org.apache.kafka.connect.json.JsonConverter"
- C. "key.converter": "io.confluent.connect.avro.AvroConverter"
- D. "connect.cosmos.containers.topicmap": "iot#telemetry"
- E. "connect.cosmos.containers.topicmap": "iot"
- F. "connector.class": "com.azure.cosmos.kafka.connect.source.CosmosDBSinkConnector"

**Answer:** CDF

**Explanation:**

C: Avro is binary format, while JSON is text.

F: Kafka Connect for Azure Cosmos DB is a connector to read from and write data to Azure Cosmos DB. The Azure Cosmos DB sink connector allows you to

export data from Apache Kafka topics to an Azure Cosmos DB database. The connector polls data from Kafka to write to containers in the database based on the topics subscription.

D: Create the Azure Cosmos DB sink connector in Kafka Connect. The following JSON body defines config for the sink connector.

Extract:

```
"connector.class": "com.azure.cosmos.kafka.connect.sink.CosmosDBSinkConnector", "key.converter": "org.apache.kafka.connect.json.AvroConverter"
"connect.cosmos.containers.topicmap": "hotels#kafka"
```

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/sql/kafka-connector-sink> <https://www.confluent.io/blog/kafka-connect-deep-dive-converters-serialization-explained/>

**NEW QUESTION 6**

- (Exam Topic 2)

You have the indexing policy shown in the following exhibit.

```

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      "path": "/name"
17    },
18    {
19      "path": "/age"
20    }
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.

**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	

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

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:

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.

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

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:

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.

None: Indexing is disabled on the container.

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

**NEW QUESTION 7**

- (Exam Topic 2)

You have a container named container1 in an Azure Cosmos DB Core (SQL) API account.

You need to provide a user named User1 with the ability to insert items into container1 by using role-based access control (RBAC). The solution must use the principle of least privilege.

Which roles should you assign to User1?

- A. CosmosDB Operator only
- B. DocumentDB Account Contributor and Cosmos DB Built-in Data Contributor
- C. DocumentDB Account Contributor only
- D. Cosmos DB Built-in Data Contributor only

**Answer:** A

**Explanation:**

Cosmos DB Operator: Can provision Azure Cosmos accounts, databases, and containers. Cannot access any data or use Data Explorer.

Reference: <https://docs.microsoft.com/en-us/azure/cosmos-db/role-based-access-control>

**NEW QUESTION 8**

- (Exam Topic 2)

You have a database in an Azure Cosmos DB SQL API Core (SQL) account that is used for development. The database is modified once per day in a batch process.

You need to ensure that you can restore the database if the last batch process fails. The solution must minimize costs.

How should you configure the backup settings? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

**Answer Area**

Backup interval

- 1 hour
- 24 hours
- 1 weeks

Backup retention

- 2 days
- 1 week
- 30 days

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

Backup interval

- 1 hour
- 24 hours
- 1 weeks

Backup retention

- 2 days
- 1 week
- 30 days

**NEW QUESTION 9**

- (Exam Topic 2)

You plan to deploy two Azure Cosmos DB Core (SQL) API accounts that will each contain a single database. The accounts will be configured as shown in the following table.

Name	Description
development	<ul style="list-style-type: none"> <li>• Supports the development of new application features</li> <li>• Used intermittently as needed during development</li> </ul>
shipments	<ul style="list-style-type: none"> <li>• Captures over 100,000 updates per second generated at unpredictable times throughout the business day</li> <li>• Used with Azure Synapse Link for analytics</li> </ul>

How should you provision the containers within each account to minimize costs? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

development:

shipments:

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

Box 1: Serverless capacity mode

Azure Cosmos DB serverless best fits scenarios where you expect intermittent and unpredictable traffic with long idle times. Because provisioning capacity in such situations isn't required and may be cost-prohibitive, Azure Cosmos DB serverless should be considered in the following use-cases:

Getting started with Azure Cosmos DB

Running applications with bursty, intermittent traffic that is hard to forecast, or low (<10%) average-to-peak traffic ratio

Developing, testing, prototyping and running in production new applications where the traffic pattern is unknown

Integrating with serverless compute services like Azure Functions

Box 2: Provisioned throughput capacity mode and autoscale throughput The use cases of autoscale include:

Variable or unpredictable workloads: When your workloads have variable or unpredictable spikes in usage, autoscale helps by automatically scaling up and down based on usage. Examples include retail websites that have different traffic patterns depending on seasonality; IOT workloads that have spikes at various times during the day; line of business applications that see peak usage a few times a month or year, and more. With autoscale, you no longer need to manually provision for peak or average capacity.

Reference:

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

<https://docs.microsoft.com/en-us/azure/cosmos-db/provision-throughput-autoscale#use-cases-of-autoscale>

**NEW QUESTION 10**

- (Exam Topic 2)

You have a database named telemetry in an Azure Cosmos DB Core (SQL) API account that stores IoT data. The database contains two containers named readings and devices.

Documents in readings have the following structure.

- id
- deviceid
- timestamp
- ownerid
- measures (array)
  - type
  - value
  - metricid

Documents in devices have the following structure.

- id
- deviceid
- owner
  - ownerid
  - emailaddress
  - name brand model

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

**Answer Area**

Statements	Yes	No
To return for all devices owned by a specific emailaddress, multiple queries must be performed	<input type="radio"/>	<input type="radio"/>
To return deviceid, ownerid, timestamp, and value for a specific metricid, a join must be performed	<input type="radio"/>	<input type="radio"/>
To return deviceid, ownerid, emailaddress, and model, a join must be performed	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

Box 1: Yes

Need to join readings and devices.

Box 2: No

Only readings is required. All required fields are in readings.

Box 3: No

Only devices is required. All required fields are in devices.

**NEW QUESTION 10**

- (Exam Topic 2)

You have a container named container1 in an Azure Cosmos DB Core (SQL) API account. Upserts of items in container1 occur every three seconds.

You have an Azure Functions app named function1 that is supposed to run whenever items are inserted or replaced in container1.

You discover that function1 runs, but not on every upsert.

You need to ensure that function1 processes each upsert within one second of the upsert. Which property should you change in the Function.json file of function1?

- A. checkpointInterval
- B. leaseCollectionsThroughput
- C. maxItemsPerInvocation
- D. feedPollDelay

**Answer:** D

**Explanation:**

With an upsert operation we can either insert or update an existing record at the same time.

FeedPollDelay: The time (in milliseconds) for the delay between polling a partition for new changes on the feed, after all current changes are drained. Default is 5,000 milliseconds, or 5 seconds.

Reference: <https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-cosmosdb-v2-trigger>

**NEW QUESTION 12**

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DP-420 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DP-420 Product From:

<https://www.2passeasy.com/dumps/DP-420/>

## Money Back Guarantee

### **DP-420 Practice Exam Features:**

- \* DP-420 Questions and Answers Updated Frequently
- \* DP-420 Practice Questions Verified by Expert Senior Certified Staff
- \* DP-420 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* DP-420 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year