

Vendor: Microsoft

Exam Code: DP-420

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

Certification: Microsoft Certifications

Total Questions: 176 Q&A (View Details)

Updated on: Apr 15, 2026

<https://www.leads4pass.com/dp-420.html>

Question 1:

HOTSPOT

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     {
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 2:

You have a database in an Azure Cosmos DB for NoSQL account.

The database contains a container named container1. The indexing mode container1 is set to none.

You configure Azure Cognitive Search to extract data from container1 and make the data searchable.

You discover that the Cognitive Search index is missing all the data from the Azure Cosmos DB index.

What should you do to resolve the issue?

- A. Modify The index attributes in Cognitive Search to searchable.
- B. Modify the index attributes in Cognitive Search to Retrievable.
- C. Change the indexing mode of container 1 to consistent-
- D. Modify the indexing policy of container 1 to exclude the / * path

Correct Answer: C

Question 3:

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
Bookauthorlnk	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?

- A. Create a container for Author and a container for Book. In each Author document, embed bookId for each book by the author. In each Book document embed authorId of each author.
- B. Create Author, Book, and Bookauthorlnk 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 Bookauthorlnk.

Correct Answer: A

Store multiple entity types in the same container.

Question 4:

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 application to use the change feed processor to read the change feed and you configure the application to trigger the function.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

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>

Question 5:

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?

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

Correct Answer: D

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.

Reference:

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

Question 6:

DRAG DROP

You have an app that stores data in an Azure Cosmos DB Core (SQL) API account. The app performs queries that return large result sets.

You need to return a complete result set to the app by using pagination. Each page of results must return 80 items.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

Actions

Configure `MaxItemCount` in `QueryRequestOptions`

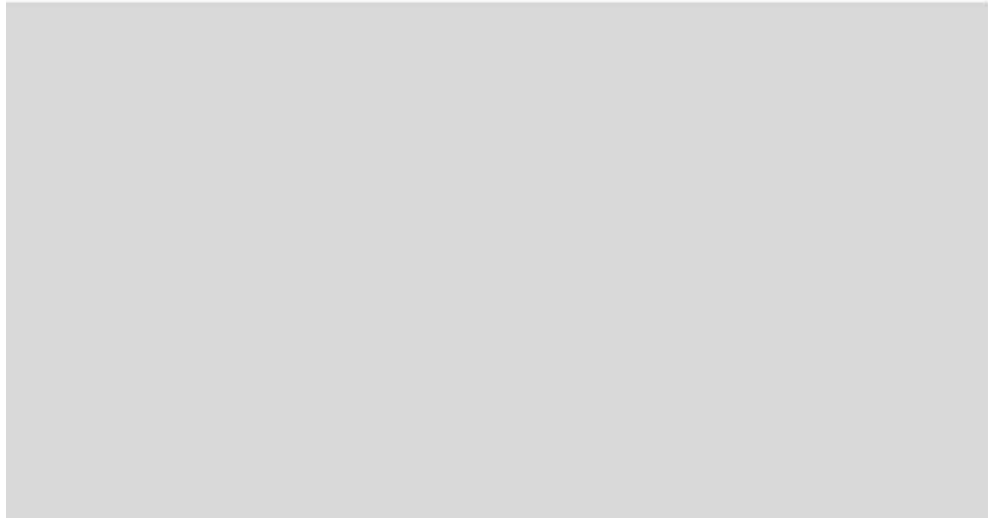
Run the query and provide a continuation token

Configure `MaxBufferedItemCount` in `QueryRequestOptions`

Append the results to a variable

Run the query and increment `MaxItemCount`

Answer Area



Correct Answer:

Actions

Configure MaxBufferedItemCount in QueryRequestOptions
Run the query and increment MaxItemCount

Answer Area

Configure MaxItemCount in QueryRequestOptions
Run the query and provide a continuation token
Append the results to a variable

When DefaultTimeToLive is -1 then your Time to Live setting is On (No default)

Time to Live on a container, if present and the value is set to "-1", it is equal to infinity, and items don't expire by default.

Time to Live on an item:

This Property is applicable only if DefaultTimeToLive is present and it is not set to null for the parent container.

If present, it overrides the DefaultTimeToLive value of the parent container.

Reference:

[https://docs.microsoft.com/en-us/azure/cosmos-db/sql/time-to-liveExplanation:](https://docs.microsoft.com/en-us/azure/cosmos-db/sql/time-to-liveExplanation)

Step 1: Configure the MaxItemCount in QueryRequestOptions

You can specify the maximum number of items returned by a query by setting the `MaxItemCount`. The `MaxItemCount` is specified per request and tells the query engine to return that number of items or fewer.

Box 2: Run the query and provide a continuation token

In the .NET SDK and Java SDK you can optionally use continuation tokens as a bookmark for your query's progress. Azure Cosmos DB query executions are stateless at the server side and can be resumed at any time using the continuation

token.

If the query returns a continuation token, then there are additional query results.

Step 3: Append the results to a variable

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/sql/sql-query-pagination>

Question 7:

HOTSPOT

You provision Azure resources by using the following Azure Resource Manager (ARM) template.

```

{
  "$schema": "https://schema.management.azure.com/schemas/2019-04-01/deploymentTemplate.json#",
  "contentVersion": "1.0.0.0",
  "parameters": {
    "db": {
      "defaultValue": "[resourceId('Microsoft.DocumentDB/databaseAccounts', 'prod1')]",
      "type": "String"
    },
    "sms": {
      "defaultValue": "[resourceId('microsoft.insights/actionGroups', 'sms')]"
    }
  },
  "variables": {},
  "resources": [
    {
      "type": "microsoft.insights/actionGroups",
      "apiVersion": "2019-06-01",
      "name": "sms",
      "location": "Global",
      "properties": {
        "groupShortName": "Send message",
        "enabled": true,
        "emailReceivers": [],
        "smsReceivers": [
          {
            "name": "Action-SMS",
            "countryCode": "44",
            "phoneNumber": "7111111111"
          }
        ]
      }
    },
    {
      "type": "microsoft.insights/activityLogAlerts",
      "apiVersion": "2020-10-01",
      "name": "Alert1",
      "location": "Global",
      "dependsOn": ["sms"],
      "properties": {
        "scopes": [ "[parameters('db')]" ],
        "condition": {
          "allOf": [
            {
              "field": "category",
              "equals": "Administrative"
            },
            {
              "field": "operationName",
              "equals": "Microsoft.DocumentDB/databaseAccounts/regenerateKey/action"
            }
          ]
        },
        "actions": {
          "actionGroups": [
            {
              "actionGroupId": "[parameters('sms')]",
              "webhookProperties": {}
            }
          ]
        },
        "enabled": true
      }
    }
  ]
}

```

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
The alert will be triggered when an Azure Cosmos DB key is used	<input type="radio"/>	<input type="radio"/>
Two alert actions will be performed when the alert is triggered	<input type="radio"/>	<input type="radio"/>
The alert will be triggered when an item that has a new partition key value is created	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Answer Area

Statements	Yes	No
The alert will be triggered when an Azure Cosmos DB key is used	<input type="radio"/>	<input checked="" type="radio"/>
Two alert actions will be performed when the alert is triggered	<input type="radio"/>	<input checked="" type="radio"/>
The alert will be triggered when an item that has a new partition key value is created	<input checked="" type="radio"/>	<input type="radio"/>

Box 1: No

An alert is triggered when the DB key is regenerated, not when it is used.

Note: The az cosmosdb keys regenerate command regenerates an access key for a Azure Cosmos DB database account.

Box 2: No

Only an SMS action will be taken.

Emailreceivers is empty so no email action is taken.

Box 3: Yes

Yes, an alert is triggered when the DB key is regenerated.

Reference:

<https://docs.microsoft.com/en-us/cli/azure/cosmosdb/keys>

Question 8:

You have an Azure Cosmos DB database named databaset contains a container named container1. The container1 container store product data and has the following indexing policy.

```

{
  "indexingMode": "consistent",
  "includedPaths":
  [
    {
      "path": "/product/category/?"
    },
    {
      "path": "/product/brand/?"
    }
  ],
  "excludedPaths":
  [
    {
      "path": "/*"
    },
    {
      "path": "/product/brand"
    }
  ]
}

```

Which path will be indexed?

- A. /product/brand
- B. /product/category
- C. /product/[]/category
- D. /product/brand/tailspin

Correct Answer: A

The indexing policy has an includedPaths array that contains only one path:

/product/brand/. This means that only the properties under /product/brand will be indexed.

The symbol indicates that only scalar values will be indexed, not arrays or objects¹.

The excludedPaths array contains a single path: /* .

This means that all other properties will be excluded from indexing.

The * symbol indicates a wildcard that matches any property name¹.

Therefore, the paths /product/category , /product/[]/category , and /product/brand/tailspin will not be indexed.

Question 9:

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"

Correct Answer: CDF

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"
```

Incorrect Answers:

B: JSON is plain text.

Note, full example:

```
{ "name": "cosmosdb-sink-connector", "config": {  
  "connector.class": "com.azure.cosmos.kafka.connect.sink.CosmosDBSinkConnector",  
  "tasks.max": "1",  
  "topics": [  
    "hotels"  
  ],  
  "value.converter": "org.apache.kafka.connect.json.AvroConverter",  
  "value.converter.schemas.enable": "false",  
  "key.converter": "org.apache.kafka.connect.json.AvroConverter",  
  "key.converter.schemas.enable": "false",  
  "connect.cosmos.connection.endpoint": "https://documents.azure.com:443/",  
  "connect.cosmos.master.key": "",  
  "connect.cosmos.databasename": "kafkaconnect",  
  "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/>

Question 10:

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>

Question 11:

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

Correct Answer: C

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>

Question 12:

You have an Azure Cosmos DB for NoSQL account1 that is configured for automatic failover. The account1 account has a single read-write region in West US and a read region in East US.

You run the following PowerShell command.

```
Update-AzCosmosDBAccountFailoverPriority -ResourceGroupName `rg1` -Name `account1` -FailoverPolicy @( `East US`, `West US` )
```

What is the effect of running the command?

- A. A manual failover will occur.
- B. The account will be unavailable to writes during the change.
- C. The provisioned throughput for account1 will increase.
- D. The account will be configured for multi-region writes.

Correct Answer: D

You can use the Set-AzCosmosDBAccountRegion cmdlet to update the regions that an Azure Cosmos DB account uses. You can use this cmdlet to add a region or change the region failover order. The cmdlet requires a resource group name, an Azure Cosmos DB account name, and a list of regions in desired failover order¹. For your scenario, based on the PowerShell command, you are using the Set-AzCosmosDBAccountRegion cmdlet to update the regions for an Azure Cosmos DB account named account1 that is configured for automatic failover. The command specifies two regions: West US and East US. The effect of running the command is that the account will be configured for multi-region writes.

Multi-region writes is a feature of Azure Cosmos DB that allows you to write data to any region in your account and have it automatically replicated to all other regions. This feature provides high availability and low latency for write operations across multiple regions. To enable multi-region writes, you need to specify at least two regions in your account and set them as write regions². In your command, you are setting both West US and East US as write regions by using the -IsZoneRedundant parameter with a value of \$true for both regions.

Question 13:

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" : "2021-09-29",
  "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 evaluating whether to use orderDate as the partition key.

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

NOTE: Each correct selection is worth one point.

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

Correct Answer: BD

Question 14:

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

You need to make the contents of container1 available as reference data for an Azure Stream Analytics job.

Solution: You create an Azure Synapse pipeline that uses Azure Cosmos DB Core (SQL) API as the input and Azure Blob Storage as the output.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Instead create an Azure function that uses Azure Cosmos DB Core (SQL) API change feed as a trigger and Azure event hub as the output.

The Azure Cosmos DB change feed is a mechanism to get a continuous and incremental feed of records from an Azure Cosmos container as those records are being created or modified. Change feed support works by listening to container for any changes. It then outputs the sorted list of documents that were changed in the order in which they were modified.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/sql/changefeed-ecommerce-solution>

Question 15:

You have an Azure Cosmos DB for NoSQL account named account1 that has a single read-write region and one additional read region.

Account1 uses the strong default consistency level.

You have an application that uses the eventual consistency level when submitting requests to account1.

How will writes from the application be handled?

- A. Writes will use the strong consistency level.
- B. Azure Cosmos DB will reject writes from the application.
- C. The write order is not guaranteed during replication.
- D. Writes will use the eventual consistency level.

Correct Answer: A

This is because the write concern is mapped to the default consistency level configured on your Azure Cosmos DB account, which is strong in this case. Strong consistency ensures that every write operation is synchronously committed to every region associated with your Azure Cosmos DB account. The eventual consistency level that the application uses only applies to the read operations. Eventual consistency offers higher availability and better performance, but it does not guarantee the order or latency of the reads.