

# Microsoft

## Exam Questions DP-420

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



### NEW QUESTION 1

- (Exam Topic 1)

You are troubleshooting the current issues caused by the application updates.

Which action can address the application updates issue without affecting the functionality of the application?

- A. Enable time to live for the con-product container.
- B. Set the default consistency level of account1 to strong.
- C. Set the default consistency level of account1 to bounded staleness.
- D. Add a custom indexing policy to the con-product container.

**Answer: C**

#### Explanation:

Bounded staleness is frequently chosen by globally distributed applications that expect low write latencies but require total global order guarantee. Bounded staleness is great for applications featuring group collaboration and sharing, stock ticker, publish-subscribe/queueing etc.

Scenario: Application updates in con-product frequently cause HTTP status code 429 "Too many requests". You discover that the 429 status code relates to excessive request unit (RU) consumption during the updates.

Reference:

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

### NEW QUESTION 2

- (Exam Topic 1)

You configure multi-region writes for account1.

You need to ensure that App1 supports the new configuration for account1. The solution must meet the business requirements and the product catalog requirements.

What should you do?

- A. Set the default consistency level of account1 to bounded staleness.
- B. Create a private endpoint connection.
- C. Modify the connection policy of App1.
- D. Increase the number of request units per second (RU/s) allocated to the con-product and con-productVendor containers.

**Answer: D**

#### Explanation:

App1 queries the con-product and con-productVendor containers.

Note: Request unit is a performance currency abstracting the system resources such as CPU, IOPS, and memory that are required to perform the database operations supported by Azure Cosmos DB.

Scenario:

Develop an app named App1 that will run from all locations and query the data in account1.

Once multi-region writes are configured, maximize the performance of App1 queries against the data in account1.

Whenever there are multiple solutions for a requirement, select the solution that provides the best performance, as long as there are no additional costs associated.

Reference:

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

### NEW QUESTION 3

- (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 4

- (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 SELECT * FROM c where c.orderDate = "2019-05-03"	<input type="radio"/>	<input type="radio"/>
If you run the following query, the query will run as a cross-partition query SELECT * FROM c where c.customerId = "5425"	<input type="radio"/>	<input type="radio"/>
If you run the following query, the query will run as a cross-partition query SELECT * FROM c where c.orderDate = "2019-05-03" and c.orderId = "9d7816e6-f401-42ba-ad05-0e03de35c0b8"	<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 5

- (Exam Topic 2)  
You have an Azure Cosmos DB Core (SQL) API account named account1.  
In account1, you run the following query in a container that contains 100GB of data. SELECT \*  
FROM c  
WHERE LOWER(c.categoryid) = "hockey"  
You view the following metrics while performing the query.

Retrieved Document Count	:	45,654
Retrieved Document Size	:	543,765,234 bytes
Output Document Count	:	12
Output Document Size	:	451 bytes
Index Utilization	:	0.00 %
Total Query Execution Time	:	2,400.34 milliseconds
Query Preparation Times		
Query Compilation Time	:	0.09 milliseconds
Logical Plan Build Time	:	0.04 milliseconds
Physical Plan Build Time	:	0.03 milliseconds
Query Optimization Time	:	0.01 milliseconds
Index Lookup Time	:	0.00 milliseconds
Document Load Time	:	3,167.26 milliseconds
Runtime Execution Times		
Query Engine Times	:	299.16 milliseconds
System Function Execution Time	:	79.34 milliseconds
User-defined Function Execution Time	:	0.00 milliseconds
Document Write Time	:	0.01 milliseconds
Client Side Metrics		
Retry Count	:	0
Request Charge	:	3,898.95 RUs

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
The query performs a cross-partition query	<input type="radio"/>	<input type="radio"/>
The query uses an index	<input type="radio"/>	<input type="radio"/>
Recreating the container with the partition key set to /categoryId will improve the performance of the query	<input type="radio"/>	<input type="radio"/>

- A. Mastered  
B. Not Mastered

**Answer:** A

#### Explanation:

Box 1: No

Each physical partition should have its own index, but since no index is used, the query is not cross-partition.

Box 2: No

Index utilization is 0% and Index Look up time is also zero.

Box 3: Yes

A partition key index will be created, and the query will perform across the partitions. Reference: <https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-query-container>

### NEW QUESTION 6

- (Exam Topic 2)

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

You plan to create a container that will store employee data for 5,000 small businesses. Each business will have up to 25 employees. Each employee item will have an emailAddress value.

You need to ensure that the emailAddress value for each employee within the same company is unique.

To what should you set the partition key and the unique key? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

### Answer Area

Partition key	<div><div></div><div>▼</div><div>companyId</div><div>companyId+emailAddress</div><div>emailAddress</div><div>employeeId</div></div>
Unique key	<div><div></div><div>▼</div><div>companyId</div><div>emailAddress</div><div>employeeId</div></div>

- A. Mastered  
B. Not Mastered

**Answer:** A

#### Explanation:

Box 1: CompanyID

After you create a container with a unique key policy, the creation of a new or an update of an existing item resulting in a duplicate within a logical partition is prevented, as specified by the unique key constraint. The partition key combined with the unique key guarantees the uniqueness of an item within the scope of the container.

For example, consider an Azure Cosmos container with Email address as the unique key constraint and CompanyID as the partition key. When you configure the user's email address with a unique key, each item has a unique email address within a given CompanyID. Two items can't be created with duplicate email addresses and with the same partition key value.

Box 2: emailAddress

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

### NEW QUESTION 7

- (Exam Topic 2)

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')]",
      "type": "String"
    }
  },
  "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.

### 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"/>

- A. Mastered  
 B. Not Mastered

**Answer: A**

### Explanation:

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>

### NEW QUESTION 8



- (Exam Topic 2)

You have an Azure Cosmos DB Core (SQL) API account.

You configure the diagnostic settings to send all log information to a Log Analytics workspace.

You need to identify when the provisioned request units per second (RU/s) for resources within the account were modified.

You write the following query. AzureDiagnostics

| where Category == "ControlPlaneRequests" What should you include in the query?

A. | where OperationName startswith "AccountUpdateStart"

B. | where OperationName startswith "SqlContainersDelete"

C. | where OperationName startswith "MongoCollectionsThroughputUpdate"

D. | where OperationName startswith "SqlContainersThroughputUpdate"

**Answer: A**

**Explanation:**

The following are the operation names in diagnostic logs for different operations: RegionAddStart, RegionAddComplete

RegionRemoveStart, RegionRemoveComplete AccountDeleteStart, AccountDeleteComplete RegionFailoverStart, RegionFailoverComplete AccountCreateStart, AccountCreateComplete

\*AccountUpdateStart\*, AccountUpdateComplete VirtualNetworkDeleteStart, VirtualNetworkDeleteComplete DiagnosticLogUpdateStart,

DiagnosticLogUpdateComplete

Reference: <https://docs.microsoft.com/en-us/azure/cosmos-db/audit-control-plane-logs>

**NEW QUESTION 9**

- (Exam Topic 2)

You have a container in an Azure Cosmos DB Core (SQL) API account. The container stores telemetry data from IoT devices. The container uses telemetryId as the partition key and has a throughput of 1,000 request units per second (RU/s). Approximately 5,000 IoT devices submit data every five minutes by using the same telemetryId value.

You have an application that performs analytics on the data and frequently reads telemetry data for a single IoT device to perform trend analysis.

The following is a sample of a document in the container.

```
{
  "id" : "9ccf1906-2a30-4dc0-9644-2185f5dcbbd7",
  "deviceId" : "bba6fe24-6d97-4935-8d58-36baa4b8a0e1",
  "telemetryId" : "9d7816e6-f401-42ba-ad05-0e03de35c0b8",
  "date" : "2019-05-03",
  "time" : "13:05",
  "temp" : "21"
}
```

You need to reduce the amount of request units (RUs) consumed by the analytics application. What should you do?

A. Decrease the offerThroughput value for the container.

B. Increase the offerThroughput value for the container.

C. Move the data to a new container that has a partition key of deviceId.

D. Move the data to a new container that uses a partition key of date.

**Answer: C**

**Explanation:**

The partition key is what will determine how data is routed in the various partitions by Cosmos DB and needs to make sense in the context of your specific scenario. The IoT Device ID is generally the "natural" partition key for IoT applications.

Reference: <https://docs.microsoft.com/en-us/azure/architecture/solution-ideas/articles/iot-using-cosmos-db>

**NEW QUESTION 10**

- (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

Last Write Wins (default) mode

Merge Procedures (custom) mode

An application that reads from the conflicts feed

## Answer Area

Container1:

Container2:

- 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 10

- (Exam Topic 2)

You have three containers in an Azure Cosmos DB Core (SQL) API account as shown in the following table.

Name	Database	Time to Live
cn1	db1	On (no default)
cn2	db1	Off
cn3	db1	On (no default)

You have the following Azure functions:

A function named Fn1 that reads the change feed of cn1 A function named Fn2 that reads the change feed of cn2 A function named Fn3 that reads the change feed of cn3

You perform the following actions: Delete an item named item1 from cn1. Update an item named item2 in cn2.

For an item named item3 in cn3, update the item time to live to 3,600 seconds.

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

Fn1 will receive item1 from the change feed

☐
☐

Fn2 can check the \_etag of item2 to see whether the item is an update or an insert

☐
☐

Fn3 will receive item3 from the change feed

☐
☐

- A. Mastered  
 B. Not Mastered

**Answer:** A

### Explanation:

Box 1: No

Azure Cosmos DB's change feed is a great choice as a central data store in event sourcing architectures where all data ingestion is modeled as writes (no updates or deletes).

Note: The change feed does not capture deletes. If you delete an item from your container, it is also removed from the change feed. The most common method of handling this is adding a soft marker on the items that are being deleted. You can add a property called "deleted" and set it to "true" at the time of deletion. This document update will show up in the change feed. You can set a TTL on this item so that it can be automatically deleted later.

Box 2: No

The \_etag format is internal and you should not take dependency on it, because it can change anytime.

Box 3: Yes

Change feed support in Azure Cosmos DB works by listening to an Azure Cosmos container for any changes. Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/sql/change-feed-design-patterns> <https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed>

## NEW QUESTION 12

- (Exam Topic 2)

You have the following query.

SELECT \* FROM

WHERE c.sensor = "TEMP1"

AND c.value < 22

AND c.timestamp >= 1619146031231

You need to recommend a composite index strategy that will minimize the request units (RUs) consumed by the query.

What should you recommend?

- A. a composite index for (sensor ASC, value ASC) and a composite index for (sensor ASC, timestamp ASC)
- B. a composite index for (sensor ASC, value ASC, timestamp ASC) and a composite index for (sensor DESC, value DESC, timestamp DESC)
- C. a composite index for (value ASC, sensor ASC) and a composite index for (timestamp ASC, sensor ASC)
- D. a composite index for (sensor ASC, value ASC, timestamp ASC)

**Answer:** A

**Explanation:**

If a query has a filter with two or more properties, adding a composite index will improve performance. Consider the following query:

SELECT \* FROM c WHERE c.name = "Tim" and c.age > 18

In the absence of a composite index on (name ASC, and age ASC), we will utilize a range index for this query. We can improve the efficiency of this query by creating a composite index for name and age.

Queries with multiple equality filters and a maximum of one range filter (such as >, <, <=, >=, !=) will utilize the composite index.

Reference:

<https://azure.microsoft.com/en-us/blog/three-ways-to-leverage-composite-indexes-in-azure-cosmos-db/>

**NEW QUESTION 15**

- (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 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 Data Factory 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


**Answer:** B

**Explanation:**

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.

The following diagram represents the data flow and components involved in the solution:




C:\Users\Admin\Desktop\Data\Odt data\Untitled.jpg

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

**NEW QUESTION 17**

- (Exam Topic 2)

You have an Azure Cosmos DB Core (SQL) API account used by an application named App1. You open the Insights pane for the account and see the following chart.



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

The HTTP 404 status code is caused by [answer choice]

incorrect connection URLs

an intermittent firewall issue

incorrectly formatted partition keys

requesting resources that do not exist

There are [answer choice] successful resource creations in the account during the time period of the chart

zero

6 thousand

6.01 thousand

30.03 thousand

36.03 thousand

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: incorrect connection URLs  
400 Bad Request: Returned when there is an error in the request URI, headers, or body. The response body will contain an error message explaining what the specific problem is.  
The HyperText Transfer Protocol (HTTP) 400 Bad Request response status code indicates that the server cannot or will not process the request due to something that is perceived to be a client error (for example, malformed request syntax, invalid request message framing, or deceptive request routing).  
Box 2: 6 thousand  
201 Created: Success on PUT or POST. Object created or updated successfully. Note:  
200 OK: Success on GET, PUT, or POST. Returned for a successful response.  
404 Not Found: Returned when a resource does not exist on the server. If you are managing or querying an index, check the syntax and verify the index name is specified correctly.  
Reference: <https://docs.microsoft.com/en-us/rest/api/searchservice/http-status-codes>

NEW QUESTION 21


- (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 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 function that uses Azure Cosmos DB Core (SQL) API change feed as a trigger and Azure event hub as the output.  
Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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.  
The following diagram represents the data flow and components involved in the solution:



C:\Users\Admin\Desktop\Data\Odt

data\Untitled.jpg  
Reference: <https://docs.microsoft.com/en-us/azure/cosmos-db/sql/changefeed-ecommerce-solution>

NEW QUESTION 25

- (Exam Topic 2)  
The settings for a container in an Azure Cosmos DB Core (SQL) API account are configured as shown in the following exhibit.

## Settings

## Indexing Policy

### Time to Live

- ☐ Off
- ☒ On (no default)
- ☐ On

### Geospatial Configuration

- ☒ Geography
- ☐ Geometry

### Partition key

/productName

Which statement describes the configuration of the container?

- A. All items will be deleted after one year.
- B. Items stored in the collection will be retained always, regardless of the items time to live value.
- C. Items stored in the collection will expire only if the item has a time to live value.
- D. All items will be deleted after one hour.

**Answer: C**

#### Explanation:

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-live>

### NEW QUESTION 30

- (Exam Topic 2)

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.

#### Actions

#### Answer Area

- 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



- A. Mastered
- B. Not Mastered

**Answer: A**

#### Explanation:

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>

### NEW QUESTION 33

- (Exam Topic 2)

You plan to create an Azure Cosmos DB Core (SQL) API account that will use customer-managed keys stored in Azure Key Vault. You need to configure an access policy in Key Vault to allow Azure Cosmos DB access to the keys. Which three permissions should you enable in the access policy? Each correct answer presents part of the solution.  
NOTE: Each correct selection is worth one point.

- A. Wrap Key
- B. Get
- C. List
- D. Update
- E. Sign
- F. Verify
- G. Unwrap Key

**Answer:** ABG

**Explanation:**  
Reference: <https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-setup-cmk>

**NEW QUESTION 35**

- (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 39**

- (Exam Topic 2)  
You need to implement a trigger in Azure Cosmos DB Core (SQL) API that will run before an item is inserted into a container. Which two actions should you perform to ensure that the trigger runs? Each correct answer presents part of the solution.  
NOTE: Each correct selection is worth one point.

- A. Append pre to the name of the JavaScript function trigger.
- B. For each create request, set the access condition in RequestOptions.
- C. Register the trigger as a pre-trigger.
- D. For each create request, set the consistency level to session in RequestOptions.
- E. For each create request, set the trigger name in RequestOptions.

**Answer:** C

**Explanation:**

C: When triggers are registered, you can specify the operations that it can run with.

F: When executing, pre-triggers are passed in the RequestOptions object by specifying PreTriggerInclude and then passing the name of the trigger in a List object.

Reference: <https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-use-stored-procedures-triggers-udfs>

**NEW QUESTION 44**

- (Exam Topic 2)

**HOTSPOT**

You configure Azure Cognitive Search to index a container in an Azure Cosmos DB Core (SQL) API account as shown in the following exhibit.

+ Add field + Add subfield Delete								
Field Name	Type	Retrievable	Filterable	Sortable	Facetable	Searchable	Analyzer	Suggester
id	Edm.String	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		<input type="checkbox"/>
name	Edm.String	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	Standard - Lucene	<input type="checkbox"/>
▼headquarters	Edm.ComplexType							
country	Edm.String	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
iso	Edm.String	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
employees	Edm.Int32	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>			
id	Edm.String	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		

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

NOTE: Each correct selection is worth one point.

**Answer Area**

The [answer choice] field is limited to exact match comparisons

country

id

name

The [answer choice] field is hidden form the search results

country

id

name

- A. Mastered  
B. Not Mastered

**Answer: A**

**Explanation:**

Box 1: country

The country field is filterable.

Note: filterable: Indicates whether to enable the field to be referenced in \$filter queries. Filterable differs from searchable in how strings are handled. Fields of type Edm.String or Collection(Edm.String) that are filterable do not undergo lexical analysis, so comparisons are for exact matches only.

Box 2: name

The name field is not Retrievable.

Retrievable: Indicates whether the field can be returned in a search result. Set this attribute to false if you want to use a field (for example, margin) as a filter, sorting, or scoring mechanism but do not want the field to be visible to the end user.

Note: searchable: Indicates whether the field is full-text searchable and can be referenced in search queries. Reference: <https://docs.microsoft.com/en-us/rest/api/searchservice/create-index>

**NEW QUESTION 47**

- (Exam Topic 2)

You have an Azure Cosmos DB Core (SQL) API account that is used by 10 web apps.

You need to analyze the data stored in the account by using Apache Spark to create machine learning models. The solution must NOT affect the performance of the web apps.

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. In an Apache Spark pool in Azure Synapse, create a table that uses cosmos.olap as the data source.  
B. Create a private endpoint connection to the account.  
C. In an Azure Synapse Analytics serverless SQL pool, create a view that uses OPENROWSET and the CosmosDB provider.  
D. Enable Azure Synapse Link for the account and Analytical store on the container.  
E. In an Apache Spark pool in Azure Synapse, create a table that uses cosmos.oltp as the data source.

**Answer: AD**

**Explanation:**

Reference:

<https://github.com/microsoft/MCW-Cosmos-DB-Real-Time-Advanced-Analytics/blob/main/Hands-on%20lab/H>

**NEW QUESTION 50**

- (Exam Topic 2)

You have an Azure Cosmos DB Core (SQL) account that has a single write region in West Europe. You run the following Azure CLI script.



```
az cosmosdb update -n $accountName -g $resourceGroupName \
  --locations regionName='West Europe' failoverPriority=0 isZoneRedundant=False \
  --locations regionName='North Europe' failoverPriority=1 isZoneRedundant=False

az cosmosdb failover-priority-change -n $accountName -g $resourceGroupName \
  --failover-policies 'North Europe=0' 'West Europe=1'
```

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
After running the script, there will be an instance of Azure Cosmos DB in North Europe that is writable	<input type="radio"/>	<input type="radio"/>
After running the script, the Azure Cosmos DB instance in West Europe will be writable	<input type="radio"/>	<input type="radio"/>
The cost of the Azure Cosmos DB account is unaffected by running the script	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:

Box 1: Yes

The Automatic failover option allows Azure Cosmos DB to failover to the region with the highest failover priority with no user action should a region become unavailable.

Box 2: No

West Europe is used for failover. Only North Europe is writable. To Configure multi-region set UseMultipleWriteLocations to true.

Box 3: Yes

Provisioned throughput with single write region costs \$0.008/hour per 100 RU/s and provisioned throughput with multiple writable regions costs \$0.016/per hour per 100 RU/s.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-multi-master> <https://docs.microsoft.com/en-us/azure/cosmos-db/optimize-cost-regions>

#### NEW QUESTION 51

- (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 52

- (Exam Topic 2)

You have an Azure Cosmos DB Core (SQL) API account named account1. You have the Azure virtual networks and subnets shown in the following table.

Subnet	Network	IP address range	Virtual machine
subnet1	vnet1	10.0.0.0/24	VM1
subnet2	vnet1	10.0.1.0/24	VM2
subnet3	vnet2	10.1.0.0/24	VM3

The vnet1 and vnet2 networks are connected by using a virtual network peer.

The Firewall and virtual network settings for account1 are configured as shown in the exhibit.

Allow access from

☐ All networks ☒ Selected networks

Configure network security for your Azure Cosmos DB account. [Learn more.](#)

Virtual networks

Secure your Azure Cosmos DB account with virtual networks. [+ Add existing virtual network](#) [+ Add new virtual network](#)

Virtual Network	Subnet	Address range	Endpoint Status
▼ vnet1	1	10.0.0.0/16	
	vnet1.subnet1	10.0.1.0/24	✓ Enabled

Firewall

Add IP ranges to allow access from the internet or your on-premises networks. [+Add my current IP](#) ⓘ

IP(Single IPv4 or CIDR range)

Exceptions

- ☐ Accept connections from within public Azure datacenters ⓘ
- ☒ Allow access from Azure Portal ⓘ

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

Statements	Yes	No
VM1 can access account 1	<input type="radio"/>	<input type="radio"/>
VM2 can access account 1	<input type="radio"/>	<input type="radio"/>
VM3 can access account 1	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

VM1 is on vnet1.subnet1 which has the Endpoint Status enabled.

Box 2: No

Only virtual network and their subnets added to Azure Cosmos account have access. Their peered VNets cannot access the account until the subnets within peered virtual networks are added to the account.

Box 3: No

Only virtual network and their subnets added to Azure Cosmos account have access.

Reference: <https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-configure-vnet-service-endpoint>

NEW QUESTION 55

.....

## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### 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

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The DP-420 Practice Test Here](#)**