DBCLOUD
EXAMS

QUESTIONS & ANSWERS

verified



Technology Provider: Microsoft
Exam: Designing and Implementing Cloud-Native Applications Using Microsoft Azure Cosmos DB
Number of questions in the database: 207
Exam Version: NOV.2026

Topic 1 - Question Set 1

Question #1*Topic 1*

HOTSPOT -
You have an Azure Cosmos DB Core (SQL) API account named account1 that has the disableKeyBasedMetadataWriteAccess property enabled.
You are developing an app named App1 that will be used by a user named DevUser1 to create containers in account1. DevUser1 has a non-privileged user account in the Azure Active Directory (Azure AD) tenant.
You need to ensure that DevUser1 can use App1 to create containers in account1.
What should you do? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Hot Area:

Answer Area

Grant permissions to create containers by using:
- Account keys
- Resource tokens
- Role-based access control (RBAC)

Create containers by using the:
- Azure AD Graph API
- Azure Resource Manager API
- SQL (Core) API

Correct Answer:

https://dbcloud.academy/

Answer Area

Grant permissions to create containers by using:

- Account keys
- Resource tokens
- Role-based access control (RBAC)

Create containers by using the:

- Azure AD Graph API
- Azure Resource Manager API
- SQL (Core) API

Box 1: Resource tokens -
Resource tokens provide access to the application resources within a database. Resource tokens: Provide access to specific containers, partition keys, documents, attachments, stored procedures, triggers, and UDFs.
Box 2: Azure Resource Manager API
You can use Azure Resource Manager to help deploy and manage your Azure Cosmos DB accounts, databases, and containers.
Incorrect Answers:
The Microsoft Graph API is a RESTful web API that enables you to access Microsoft Cloud service resources.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/secure-access-to-data
https://docs.microsoft.com/en-us/rest/api/resources/

Question #2*Topic 1*

HOTSPOT -
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.

https://dbcloud.academy/

Hot Area:

Answer Area

Statements	Yes	No
After running the script, there will be an instance of Azure Cosmos DB in North Europe that is writable	○	○
After running the script, the Azure Cosmos DB instance in West Europe will be writable	○	○
The cost of the Azure Cosmos DB account is unaffected by running the script	○	○

Correct Answer:

Answer Area

Statements	Yes	No
After running the script, there will be an instance of Azure Cosmos DB in North Europe that is writable	○	○
After running the script, the Azure Cosmos DB instance in West Europe will be writable	○	○
The cost of the Azure Cosmos DB account is unaffected by running the script	○	○

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

Question #3*Topic 1*

You are developing an application that will use an Azure Cosmos DB Core (SQL) API account as a data source.
You need to create a report that displays the top five most ordered fruits as shown in the following table.

https://dbcloud.academy/

Name	Type	Orders
apple	fruit	1,000
orange	fruit	600
banana	fruit, exotic	400
plum	fruit	300
mango	fruit, exotic	200

A collection that contains aggregated data already exists. The following is a sample document:

```
{
    "name": "apple",
    "type": ["fruit", "exotic"]
    "orders": 10000
}
```

Which two queries can you use to retrieve data for the report? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A.

```
SELECT TOP  i.name, i.types, i.orders
FROM items i
WHERE EXISTS(SELECT VALUE t FROM t IN i.types WHERE t.name = 'fruit')
ORDER BY i.orders,i.types
```

B.

```
SELECT TOP  i.name, i.types, i.orders
FROM items i
WHERE EXISTS(SELECT VALUE t FROM t IN i.types WHERE t.name = 'fruit')
ORDER BY i.orders DESC
```

C.

```
SELECT TOP  i.name, i.types, i.orders
FROM items i
WHERE EXISTS(SELECT VALUE t FROM t IN i.types WHERE t.name = 'fruit')
ORDER BY i.types DESC
```

D.

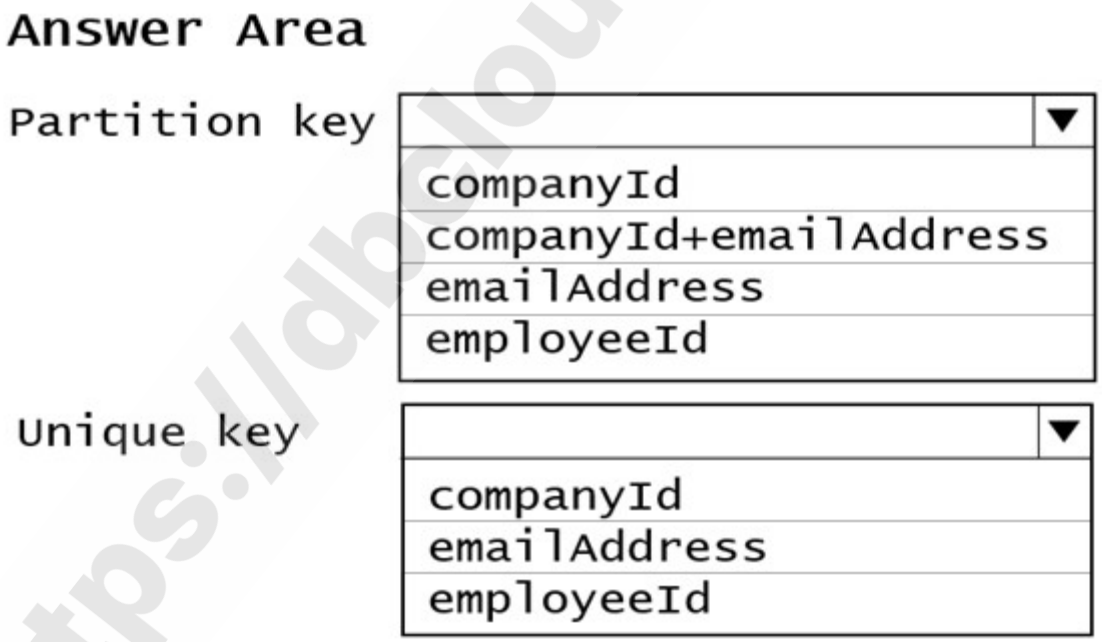
```
SELECT TOP  i.name, i.types, i.orders
FROM items i
WHERE ARRAY_CONTAINS(i.types, {name: 'fruit'})
ORDER BY i.orders DESC
```

https://dbcloud.academy/

Correct Answer: *BD*
ARRAY_CONTAINS returns a Boolean indicating whether the array contains the specified value. You can check for a partial or full match of an object by using a boolean expression within the command.
Incorrect Answers:
A: Default sorting ordering is Ascending. Must use Descending order.
C: Order on Orders not on Type.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/sql/sql-query-array-contains

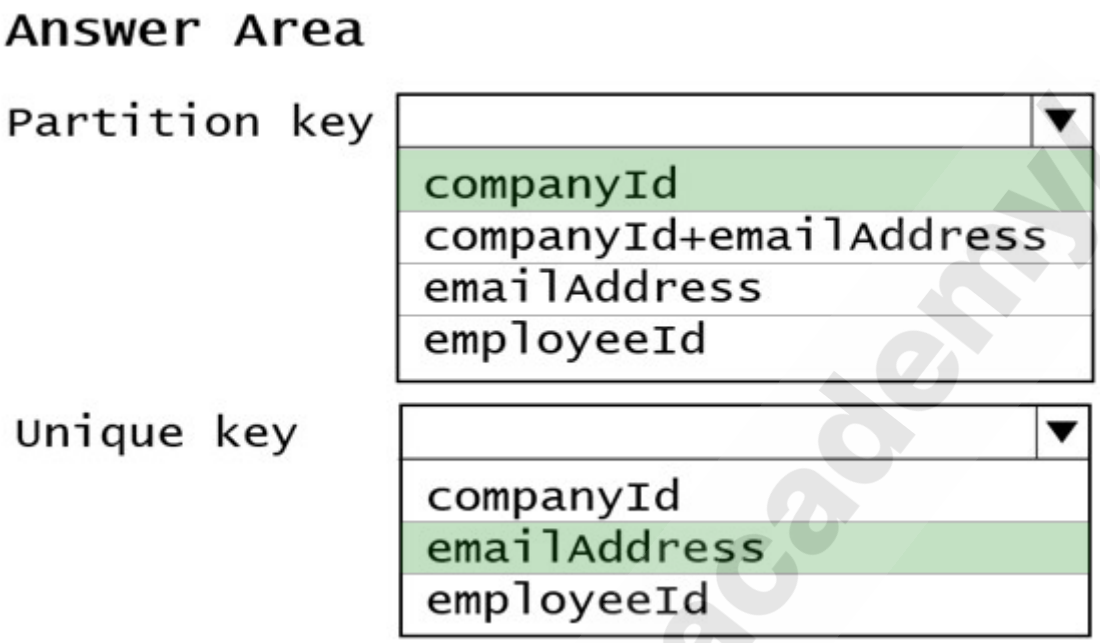
Question #4*Topic 1*

HOTSPOT -
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.
Hot Area:



Correct Answer:

https://dbcloud.academy/



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

Question #5*Topic 1*

HOTSPOT -
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" : []
}
```

https://dbcloud.academy/

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.
Hot Area:

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"`
	○	○

If you run the following query, the query will run as a cross-partition query
`SELECT * FROM c`
`where c.customerId = "5425"`
	○	○

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"`
	○	○

Correct Answer:

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"`
	○	○

If you run the following query, the query will run as a cross-partition query
`SELECT * FROM c`
`where c.customerId = "5425"`
	○	○

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"`
	○	○

Box 1: Yes -
Records with different OrderIDs will match.

Box 2: Yes -
Records with different OrderIDs will match.

https://dbcloud.academy/