Microsoft.DP-420.v2022-05-16.q18

Exam Code:	DP-420
Exam Name:	Designing and Implementing Cloud-Native Applications Using Microsoft Azure Cosmos DB
Certification Provider:	Microsoft
Free Question Number:	18
Version:	v2022-05-16
# of views:	2923
# of Questions views:	1156
https://www.dumpsfiles.com/files/Microsoft/DP-420/Microsoft.DP-420.v2022-05-16.q18	

NEW QUESTION: 1
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.

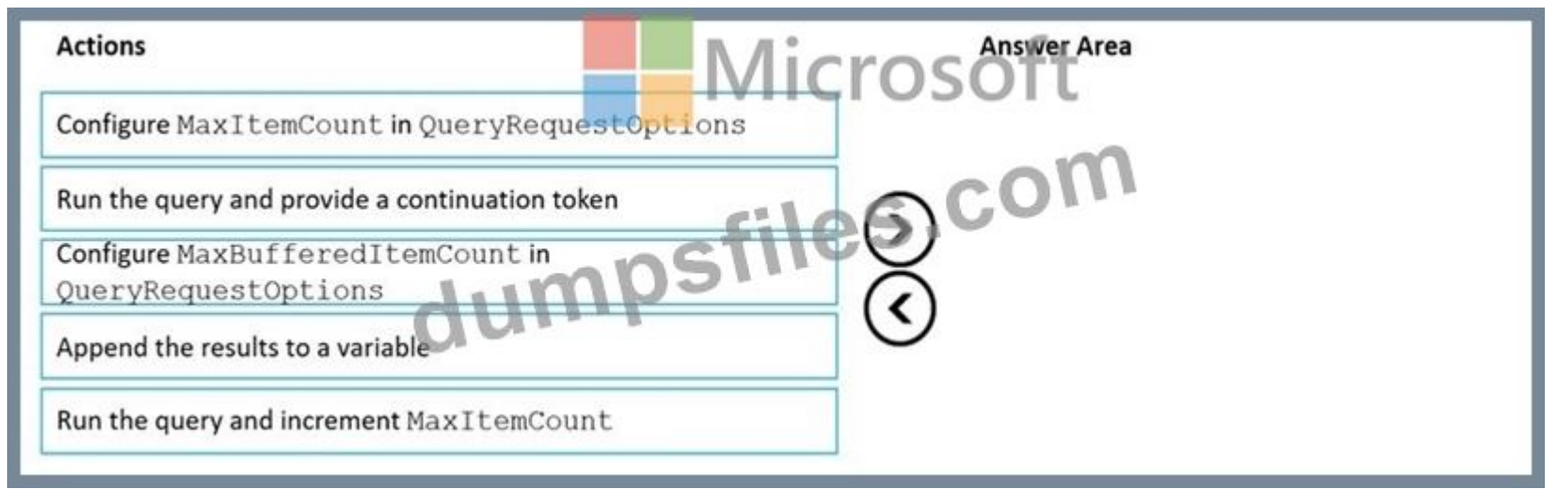

Answer:

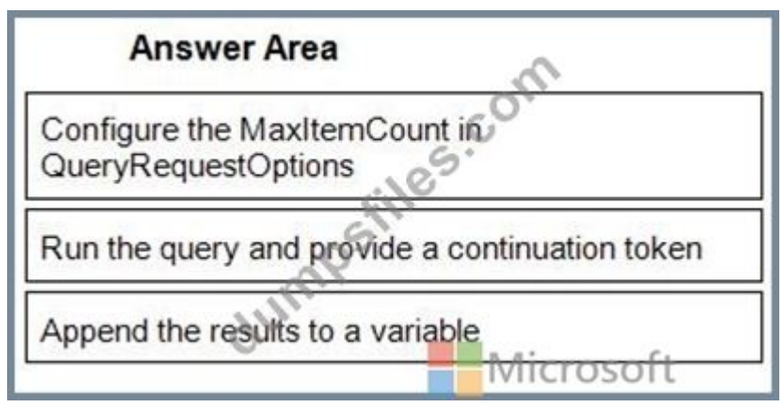

1 - Configure the MaxItemCount in QueryRequestOptions
2 - Run the query and provide a continuation token

3 - Append the results to a variable

NEW QUESTION: 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: (SHOW ANSWER)
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.

NEW QUESTION: 3
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	• Supports the development of new application features • Used intermittently as needed during development
shipments	• Captures over 100,000 updates per second generated at unpredictable rimes throughout the business day • Used with Azure Synapse Link for analytics

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.

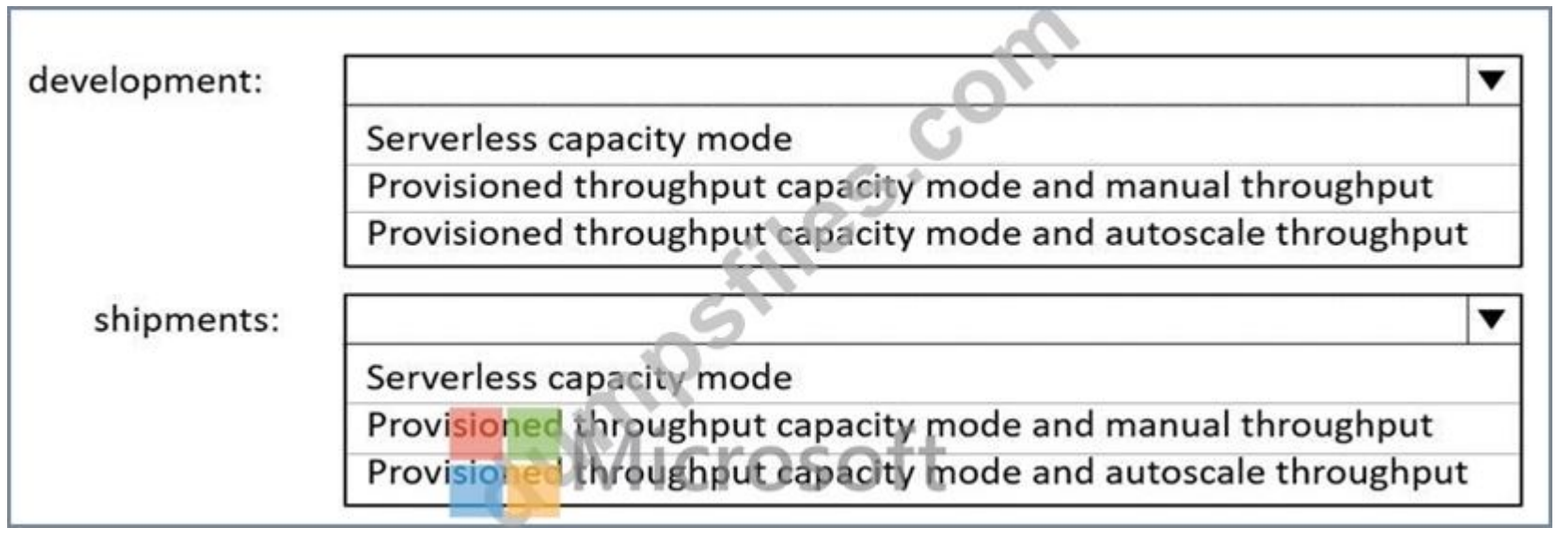

Answer:

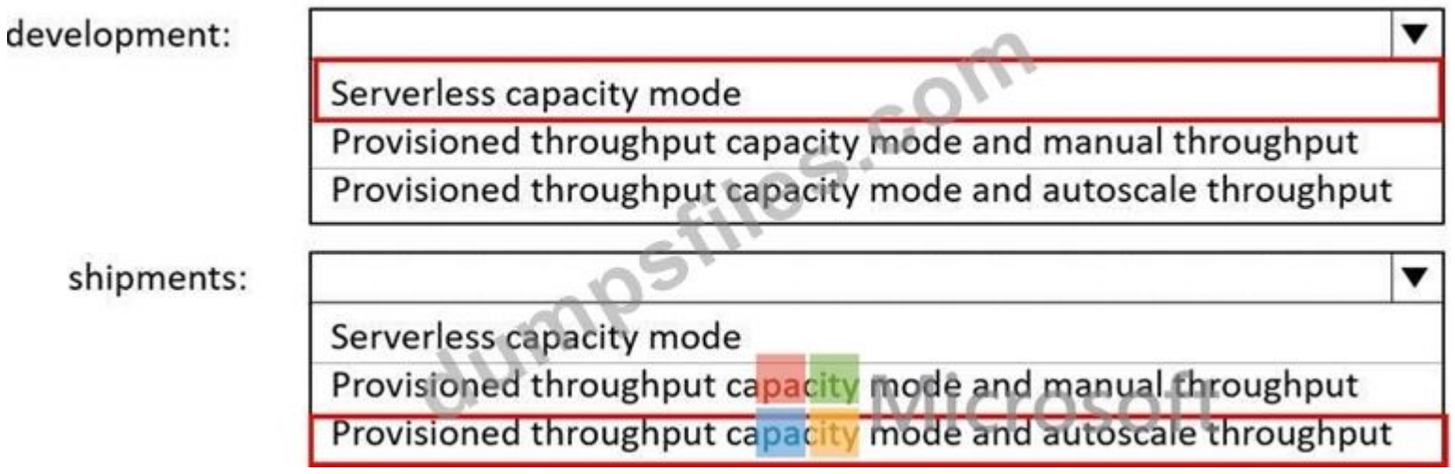

Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/serverless
https://docs.microsoft.com/en-us/azure/cosmos-db/provision-throughput-autoscale#use-cases-of-autoscale

NEW QUESTION: 4
You have a database named telemetry in an Azure Cosmos DB Core (SQL) API account that stores IoT dat a. 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.

Statements	Yes	No
To return for all devices owned by a specific `emailaddress`, multiple queries must be performed	○	○
To return `deviceid`, `ownerid`, `timestamp`, and `value` for a specific `metricid`, a join must be performed	○	○
To return `deviceid`, `ownerid`, `emailaddress`, and `model`, a join must be performed	○	○

Answer:

Statements	Yes	No
To return for all devices owned by a specific `emailaddress`, multiple queries must be performed	○	○
To return `deviceid`, `ownerid`, `timestamp`, and `value` for a specific `metricid`, a join must be performed	○	○
To return `deviceid`, `ownerid`, `emailaddress`, and `model`, a join must be performed	○	○

NEW QUESTION: 5
You have the indexing policy shown in the following exhibit.

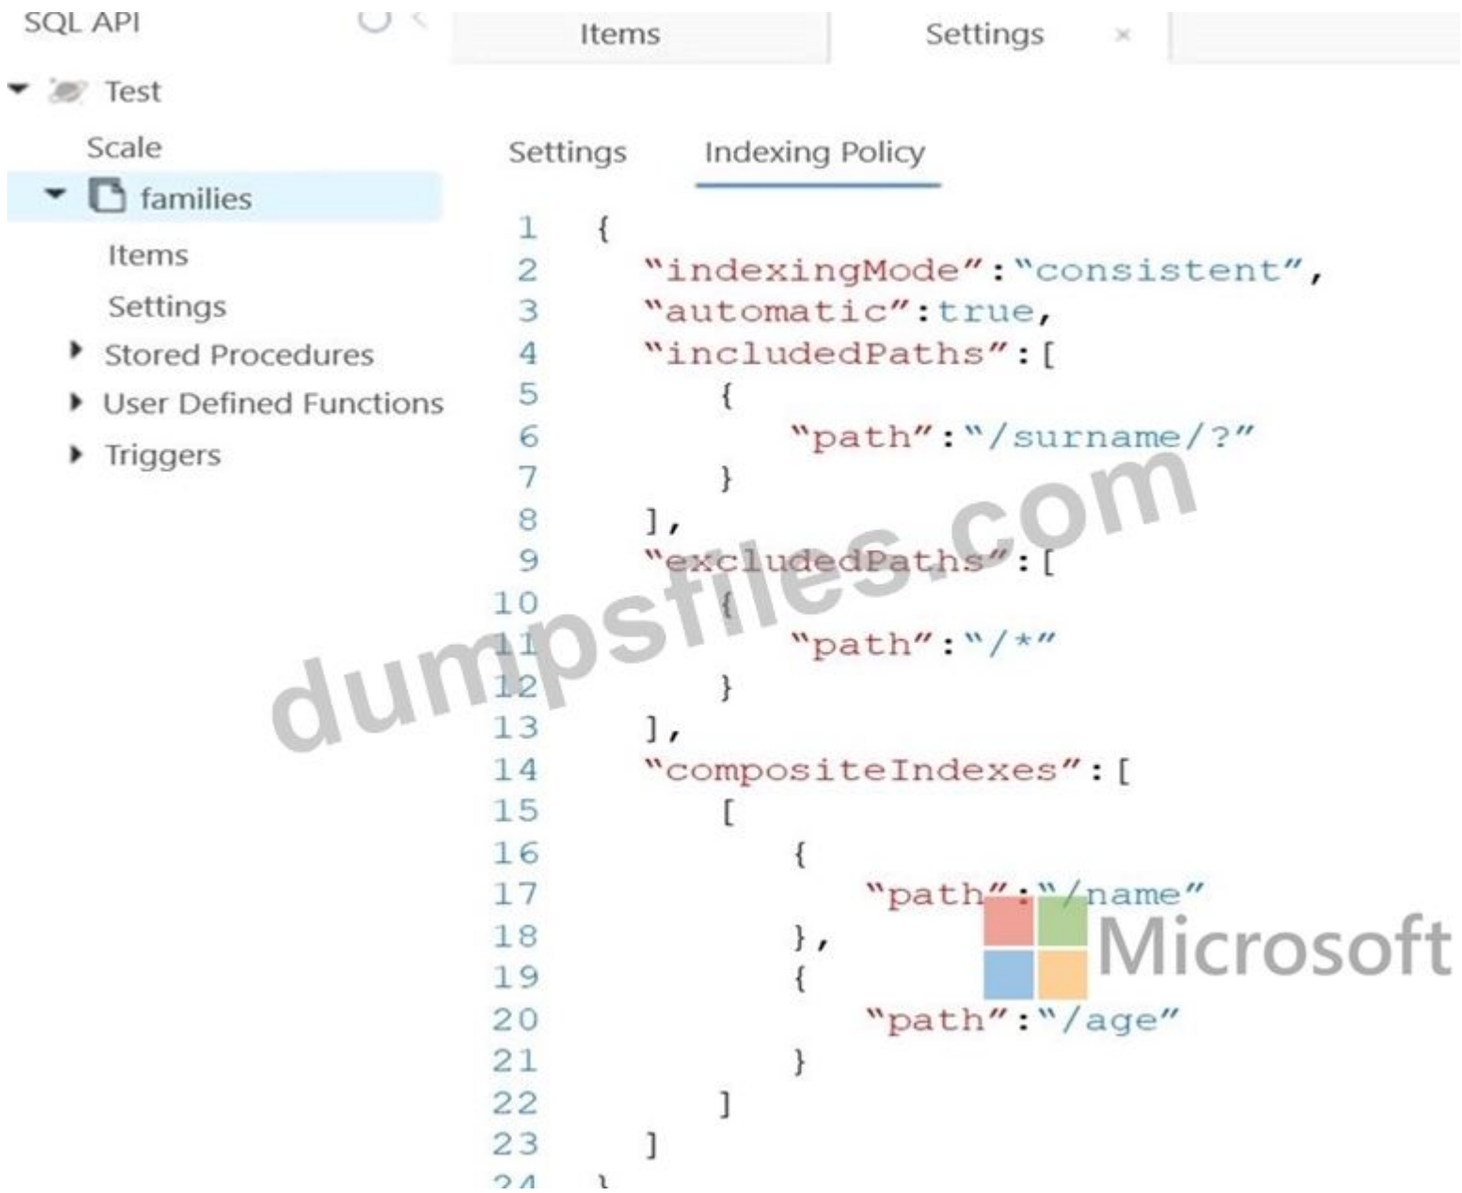

Use the drop-down menus to select the answer choice that answers each question based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

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

Answer:

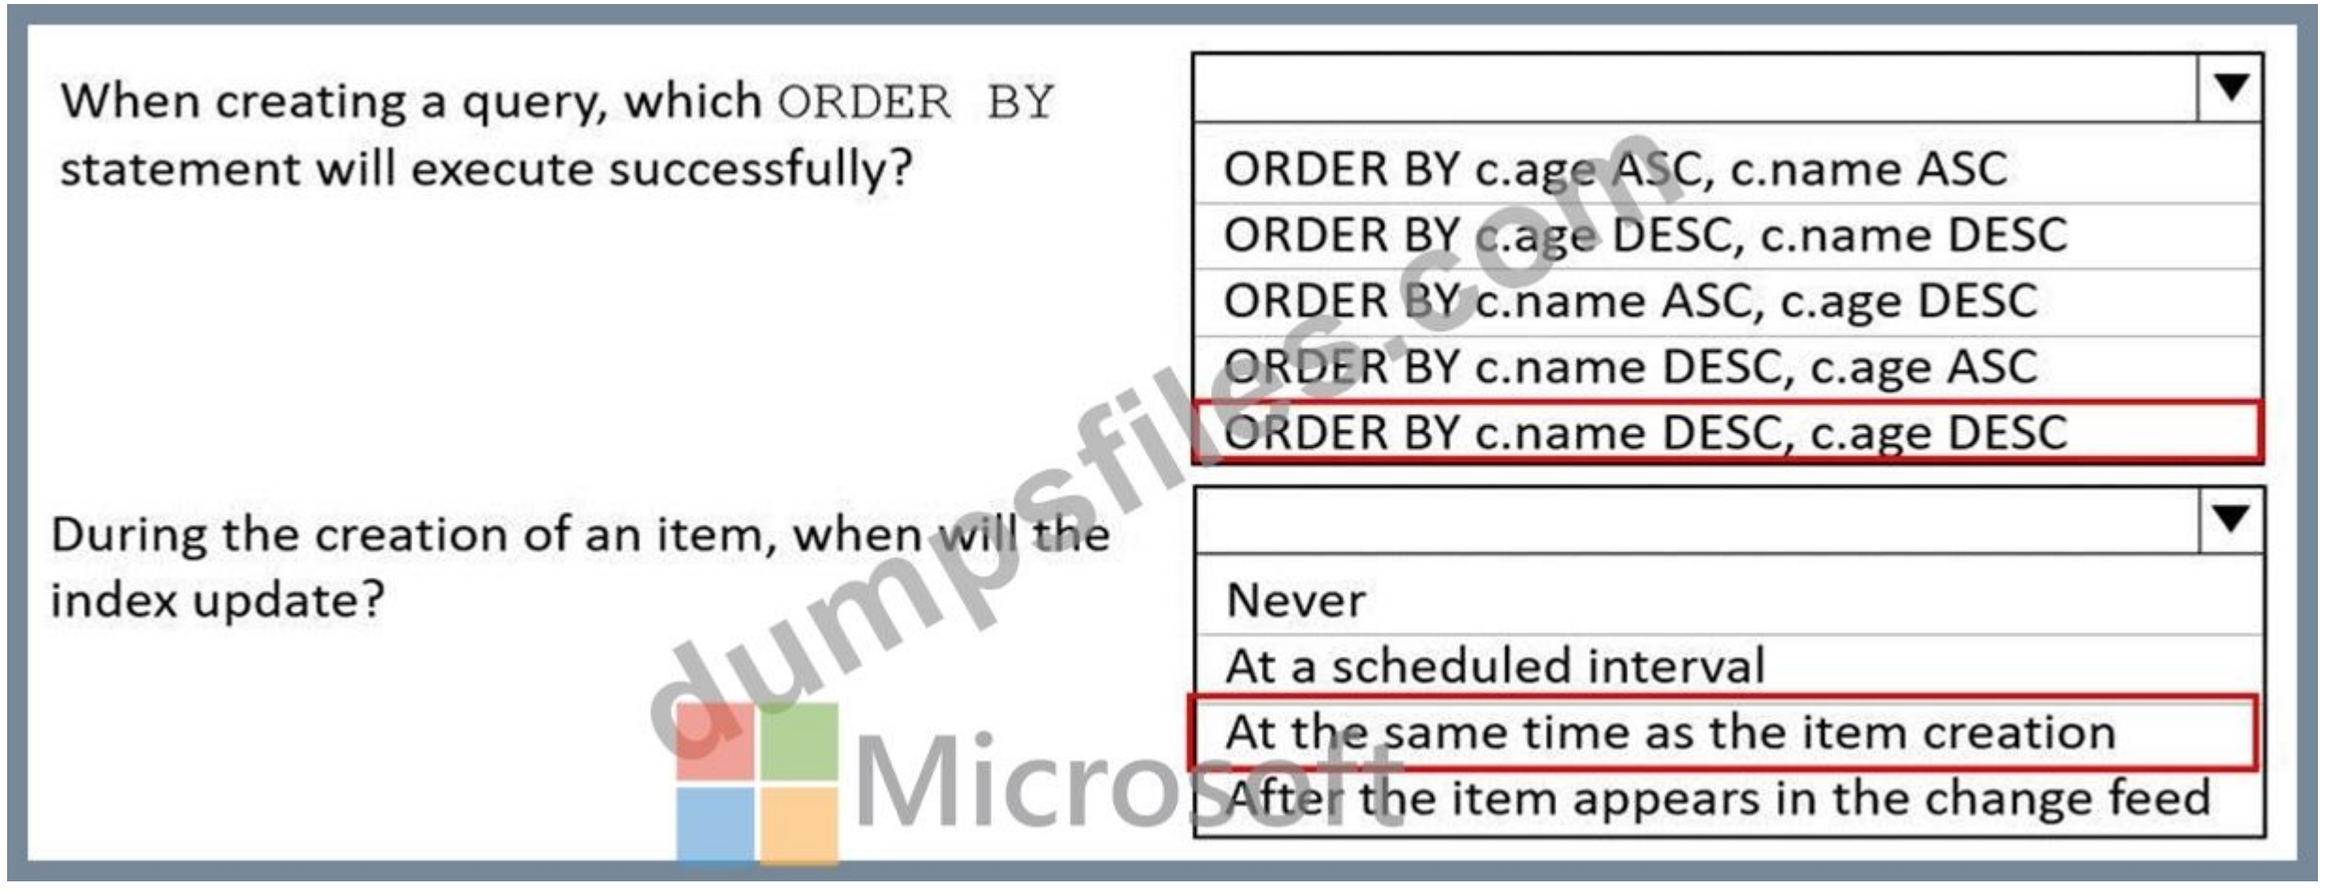

NEW QUESTION: 6

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 (LEAVE A REPLY)
Cosmos DB Operator: Can provision Azure Cosmos accounts, databases, and containers. Cannot access any data or use Data Explorer.
Incorrect Answers:
B: DocumentDB Account Contributor can manage Azure Cosmos DB accounts. Azure Cosmos DB is formerly known as DocumentDB.
C: DocumentDB Account Contributor: Can manage Azure Cosmos DB accounts.

NEW QUESTION: 7
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 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
Answer: B (LEAVE A REPLY)
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.

NEW QUESTION: 8
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.

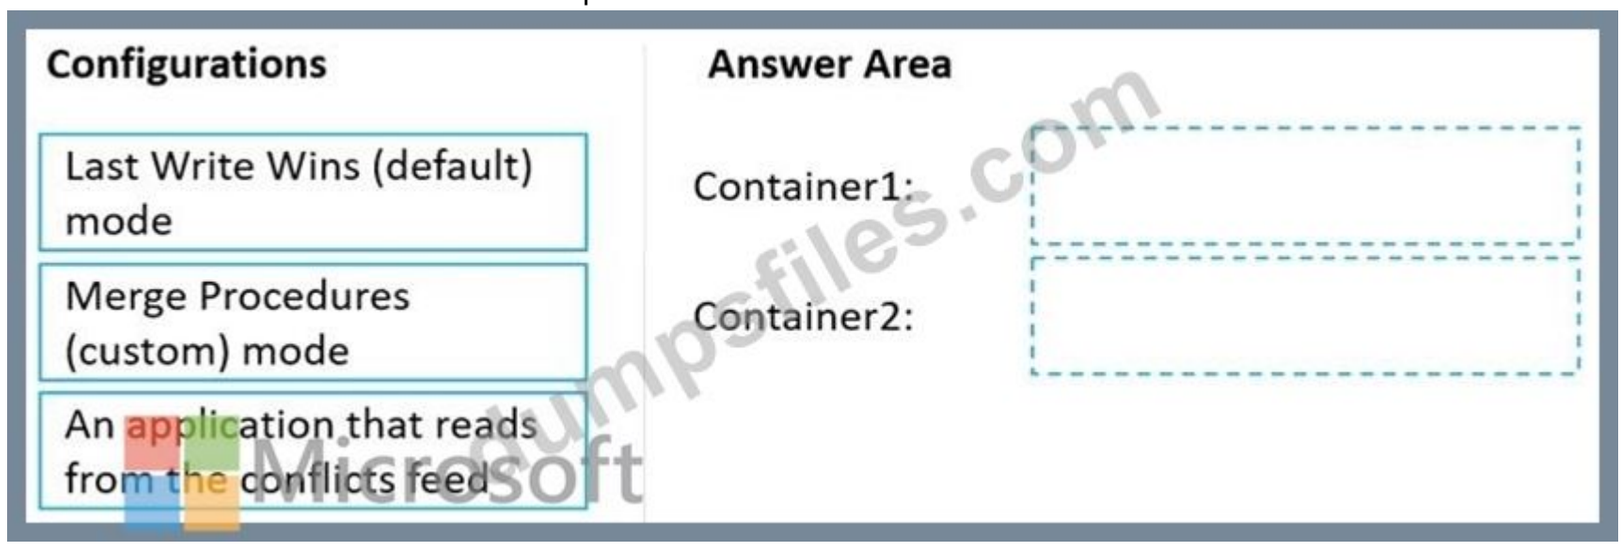

Answer:

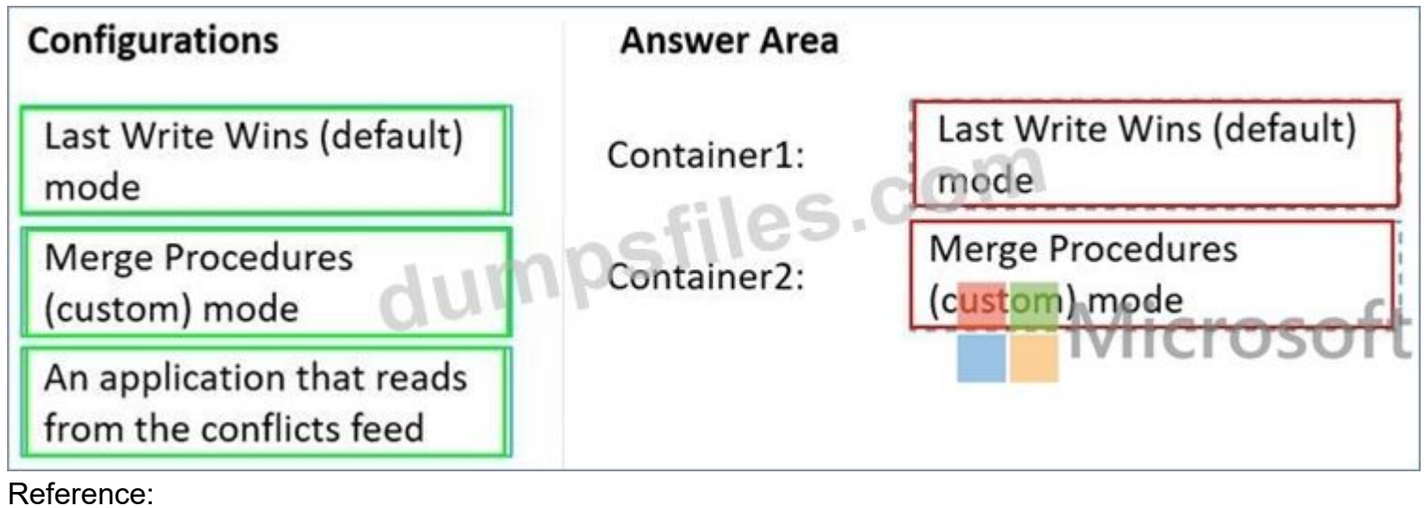

Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/conflict-resolution-policies
https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-manage-conflicts

NEW QUESTION: 9
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 accountl 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: (SHOW ANSWER)
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: 10
You need to provide a solution for the Azure Functions notifications following updates to con-product. The solution must meet the business requirements and the product catalog requirements.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.
A. Configure the trigger for each function to use a different leaseCollectionPrefix
B. Configure the trigger for each function to use the same leaseCollectionNair.e
C. Configure the trigger for each function to use a different leaseCollectionName
D. Configure the trigger for each function to use the same leaseCollectionPrefix
Answer: (SHOW ANSWER)
leaseCollectionPrefix: when set, the value is added as a prefix to the leases created in the Lease collection for this Function. Using a prefix allows two separate Azure Functions to share the same Lease collection by using different prefixes.
Scenario: Use Azure Functions to send notifications about product updates to different recipients.
Trigger the execution of two Azure functions following every update to any document in the con-product container.
Reference:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-cosmosdb-v2-trigger

NEW QUESTION: 11
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: A,D (LEAVE A REPLY)

Reference:

https://github.com/microsoft/MCW-Cosmos-DB-Real-Time-Advanced-Analytics/blob/main/Hands-on%20lab/HOL%20step-by%20step%20-%20Cosmos%20DB%20real-time%20advanced%20analytics.md

NEW QUESTION: 12

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.

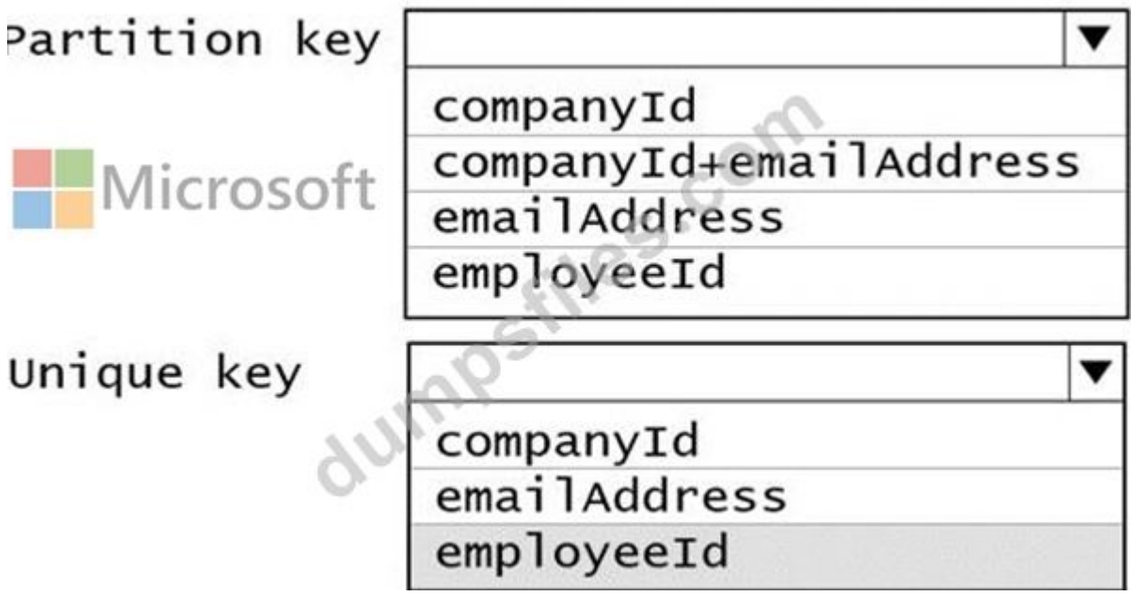

Answer:

NEW QUESTION: 13

You have a database in an Azure Cosmos DB Core (SQL) API account. The database is backed up every two hours.

You need to implement a solution that supports point-in-time restore.
What should you do first?
A. Configure the Point In Time Restore settings for the account.
B. Configure the Backup & Restore settings for the account.
C. Create a new account that has a periodic backup policy.
D. Enable Continuous Backup for the account.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 14
You are designing an Azure Cosmos DB Core (SQL) API solution to store data from IoT devices. Writes from the devices will be 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:
Minimizes the partition skew
Avoids capacity limits
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.
Answer: (SHOW ANSWER)
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: Senser1Value has only two values.
C: All the devices could have the same manufacturer.

NEW QUESTION: 15

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.

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:

NEW QUESTION: 16

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 (LEAVE A REPLY)

NEW QUESTION: 17

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 (LEAVE A REPLY)
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:

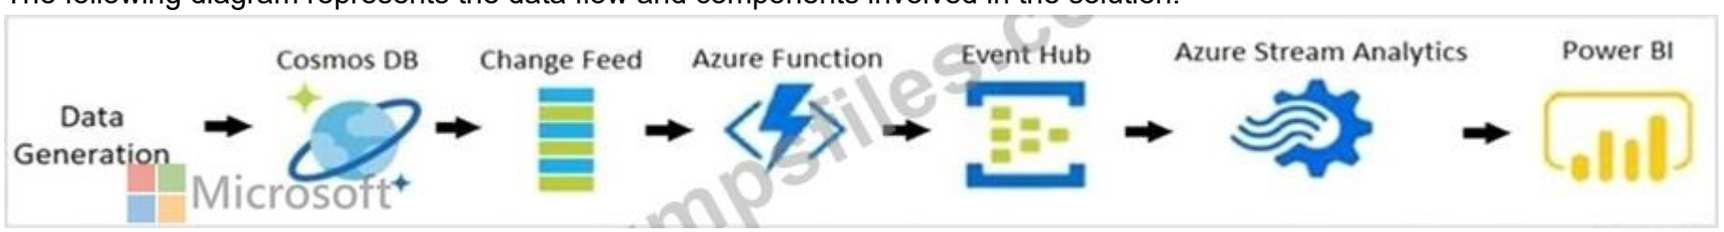

NEW QUESTION: 18
You have an Azure Cosmos DB Core (SQL) account that has a single write region in West Europe.
You run the following Azure CLI script.

```
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.

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	○	○

Answer:

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