Exam Questions DP-420

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

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

NEW QUESTION 1
- (Exam Topic 1)
You need to select the partition key for con-iot1. The solution must meet the IoT telemetry requirements. What should you select?

A. the timestamp
B. the humidity
C. the temperature
D. the device ID

Answer: D

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.
Scenario: The iotdb database will contain two containers named con-iot1 and con-iot2. Ensure that Azure Cosmos DB costs for IoT-related processing are predictable. Reference:
https://docs.microsoft.com/en-us/azure/architecture/solution-ideas/articles/iot-using-cosmos-db

NEW QUESTION 2
- (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	O	O
The query uses an index	O	O
Recreating the container with the partition key set to /categoryId will improve the performance of the query	O	O

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 3
- (Exam Topic 2)
You have a container named container1 in an Azure Cosmos DB Core (SQL) API account. The following is a sample of a document in container1.

{
"studentId": "631282", "firstName": "James", "lastName": "Smith", "enrollmentYear": 1990, "isActivelyEnrolled": true, "address": {
"street": "",
"city": "",
"stateProvince": "",
"postal": "",
}
}
The container1 container has the following indexing policy.
{
"indexingMode": "consistent", "includePaths": [
{
"path": "/*"
},
{
"path": "/address/city/?"
}
],
"excludePaths": [
{
"path": "/address/*"
},
{
"path": "/firstName/?"
}
]
}
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 /isActivelyEnrolled property is included in the index	○	○
The /fisrtname property is included in the index	○	○
The /address/city property is included in the index	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Yes
"path": "/*" is in includePaths.
Include the root path to selectively exclude paths that don't need to be indexed. This is the recommended approach as it lets Azure Cosmos DB proactively index any new property that may be added to your model.
Box 2: No
"path": "/firstName/?" is in excludePaths.
Box 3: Yes
"path": "/address/city/?" is in includePaths
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/index-policy

NEW QUESTION 4
- (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 Synapse pipeline that uses Azure Cosmos DB Core (SQL) API as the input and Azure Blob Storage as the output.
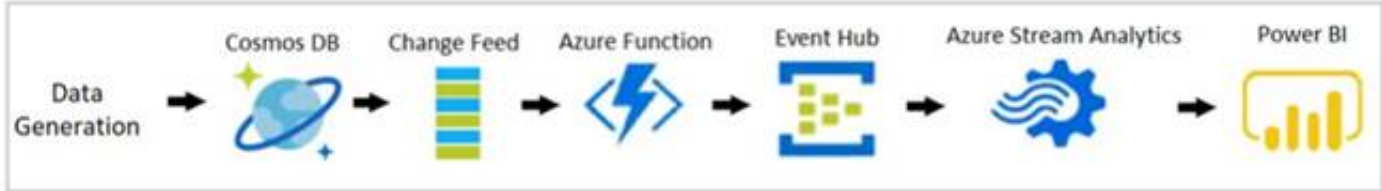Does this meet the goal?

A. Yes
B. No

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 5
- (Exam Topic 2)
You have a container in an Azure Cosmos DB Core (SQL) API account.
You need to use the Azure Cosmos DB SDK to replace a document by using optimistic concurrency. What should you include in the code? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: ConsistencyLevel
The ItemRequestOptions Class ConsistencyLevel property gets or sets the consistency level required for the request in the Azure Cosmos DB service.
Azure Cosmos DB offers 5 different consistency levels. Strong, Bounded Staleness, Session, Consistent Prefix and Eventual - in order of strongest to weakest consistency.
Box 2: _etag
The ItemRequestOptions class helped us implement optimistic concurrency by specifying that we wanted the SDK to use the If-Match header to allow the server to decide whether a resource should be updated. The If-Match value is the ETag value to be checked against. If the ETag value matches the server ETag value, the resource is updated.
Reference:
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.itemrequestoptions https://cosmosdb.github.io/labs/dotnet/labs/10-concurrency-control.html

NEW QUESTION 6
- (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.
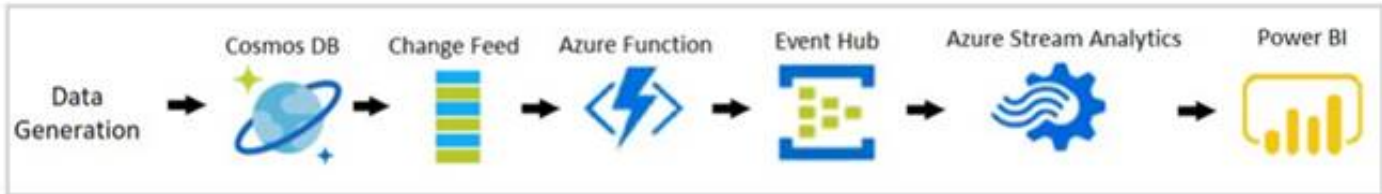Does this meet the goal?

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 7
- (Exam Topic 2)
You need to configure an Apache Kafka instance to ingest data from an Azure Cosmos DB Core (SQL) API account. The data from a container named telemetry must be added to a Kafka topic named iot. The solution must store the data in a compact binary format.
Which three configuration items should you include in the solution? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

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

Answer: CDF

Explanation:
C: Avro is binary format, while JSON is text.
F: Kafka Connect for Azure Cosmos DB is a connector to read from and write data to Azure Cosmos DB. The Azure Cosmos DB sink connector allows you to export data from Apache Kafka topics to an Azure Cosmos DB database. The connector polls data from Kafka to write to containers in the database based on the topics subscription.
D: Create the Azure Cosmos DB sink connector in Kafka Connect. The following JSON body defines config for the sink connector.
Extract:
"connector.class": "com.azure.cosmos.kafka.connect.sink.CosmosDBSinkConnector", "key.converter": "org.apache.kafka.connect.json.AvroConverter"
"connect.cosmos.containers.topicmap": "hotels#kafka"
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/sql/kafka-connector-sink https://www.confluent.io/blog/kafka-connect-deep-dive-converters-serialization-explained/

NEW QUESTION 8
- (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.
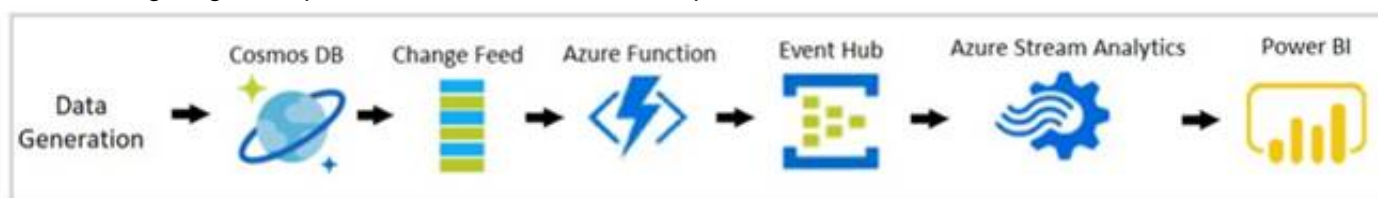Does this meet the goal?

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 9
- (Exam Topic 2)
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

Answer: D

Explanation:
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

NEW QUESTION 10
- (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.

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 10
- (Exam Topic 2)
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure Cosmos DB Core (SQL) API account named account 1 that uses autoscale throughput. You need to run an Azure function when the normalized request units per second for a container in account1 exceeds a specific value.
Solution: You configure 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

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

NEW QUESTION 11
- (Exam Topic 2)
You have the indexing policy shown in the following exhibit.

```
SQL API          ↻ <           Items          Settings      ×

  ▼ ◉ Test
      Scale              Settings     Indexing Policy
  ▼ 🗋 families          1   {
      Items             2       "indexingMode":"consistent",
      Settings          3       "automatic":true,
    ▸ Stored Procedures 4       "includedPaths":[
    ▸ User Defined Functions 5       {
    ▸ Triggers          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.

Answer Area

When creating a query, which ORDER BY statement will execute successfully?

▼
ORDER BY c.age ASC, c.name ASC
ORDER BY c.age DESC, c.name DESC
ORDER BY c.name ASC, c.age DESC
ORDER BY c.name DESC, c.age ASC
ORDER BY c.name DESC, c.age DESC

During the creation of an item, when will the index update?

▼
Never
At a scheduled interval
At the same time as the item creation
After the item appears in the change feed

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: ORDER BY c.name DESC, c.age DESC
Queries that have an ORDER BY clause with two or more properties require a composite index.
The following considerations are used when using composite indexes for queries with an ORDER BY clause with two or more properties:
If the composite index paths do not match the sequence of the properties in the ORDER BY clause, then the composite index can't support the query.
The order of composite index paths (ascending or descending) should also match the order in the ORDER BY clause.
The composite index also supports an ORDER BY clause with the opposite order on all paths. Box 2: At the same time as the item creation
Azure Cosmos DB supports two indexing modes:
Consistent: The index is updated synchronously as you create, update or delete items. This means that the consistency of your read queries will be the consistency configured for the account.
None: Indexing is disabled on the container.
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/index-policy

NEW QUESTION 16
- (Exam Topic 2)
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.

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

A. Mastered
B. Not Mastered

Answer: A

Explanation:
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.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/secure-access-to-data https://docs.microsoft.com/en-us/rest/api/resources/

NEW QUESTION 20
- (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

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

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 23
- (Exam Topic 2)
You have a container named container1 in an Azure Cosmos DB Core (SQL) API account.
You need to provide a user named User1 with the ability to insert items into container1 by using role-based access control (RBAC). The solution must use the principle of least privilege.
Which roles should you assign to User1?

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

Answer: A

Explanation:

Cosmos DB Operator: Can provision Azure Cosmos accounts, databases, and containers. Cannot access any data or use Data Explorer.
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/role-based-access-control

NEW QUESTION 24
- (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 27
- (Exam Topic 2)
You have an application named App1 that reads the data in an Azure Cosmos DB Core (SQL) API account. App1 runs the same read queries every minute. The default consistency level for the account is set to eventual.
You discover that every query consumes request units (RUs) instead of using the cache.
You verify the IntegratedCacheiteItemHitRate metric and the IntegratedCacheQueryHitRate metric. Both metrics have values of 0.
You verify that the dedicated gateway cluster is provisioned and used in the connection string. You need to ensure that App1 uses the Azure Cosmos DB integrated cache.
What should you configure?

A. the indexing policy of the Azure Cosmos DB container
B. the consistency level of the requests from App1
C. the connectivity mode of the App1 CosmosClient
D. the default consistency level of the Azure Cosmos DB account

Answer: C

Explanation:
Because the integrated cache is specific to your Azure Cosmos DB account and requires significant CPU and memory, it requires a dedicated gateway node.
Connect to Azure Cosmos DB using gateway mode.
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/integrated-cache-faq

NEW QUESTION 29
- (Exam Topic 2)
You plan to deploy two Azure Cosmos DB Core (SQL) API accounts that will each contain a single database. The accounts will be configured as shown in the following table.

Name	Description
development	• 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.

development:
▼
Serverless capacity mode
Provisioned throughput capacity mode and manual throughput
Provisioned throughput capacity mode and autoscale throughput

shipments:
▼
Serverless capacity mode
Provisioned throughput capacity mode and manual throughput
Provisioned throughput capacity mode and autoscale throughput

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Serverless capacity mode
Azure Cosmos DB serverless best fits scenarios where you expect intermittent and unpredictable traffic with long idle times. Because provisioning capacity in such situations isn't required and may be cost-prohibitive, Azure Cosmos DB serverless should be considered in the following use-cases:
Getting started with Azure Cosmos DB
Running applications with bursty, intermittent traffic that is hard to forecast, or low (<10%) average-to-peak traffic ratio
Developing, testing, prototyping and running in production new applications where the traffic pattern is unknown
Integrating with serverless compute services like Azure Functions
Box 2: Provisioned throughput capacity mode and autoscale throughput The use cases of autoscale include:
Variable or unpredictable workloads: When your workloads have variable or unpredictable spikes in usage, autoscale helps by automatically scaling up and down based on usage. Examples include retail websites that have different traffic patterns depending on seasonality; IOT workloads that have spikes at various times during the day; line of business applications that see peak usage a few times a month or year, and more. With autoscale, you no longer need to manually provision for peak or average capacity.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/serverless
https://docs.microsoft.com/en-us/azure/cosmos-db/provision-throughput-autoscale#use-cases-of-autoscale

NEW QUESTION 33
- (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	○	○
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	○	○

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 37
- (Exam Topic 2)
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. Enable Continuous Backup for the account.
B. Configure the Backup & Restore settings for the account.
C. Create a new account that has a periodic backup policy.
D. Configure the Point In Time Restore settings for the account.

Answer: A

Explanation:
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/provision-account-continuous-backup

NEW QUESTION 41
- (Exam Topic 2)
You are creating a database in an Azure Cosmos DB Core (SQL) API account. The database will be used by an application that will provide users with the ability to share online posts. Users will also be able to submit comments on other users' posts.
You need to store the data shown in the following table.

Type	Description
Users	Information about a user who will use the application
Posts	Text of up to 1,000 characters that a user will share with other users
Comments	Text of up to 280 characters that users will submit as a comment on a post
Interests	Information about a user's interests

The application has the following characteristics: Users can submit an unlimited number of posts.
The average number of posts submitted by a user will be more than 1,000. Posts can have an unlimited number of comments from different users.
The average number of comments per post will be 100, but many posts will exceed 1,000 comments. Users will be limited to having a maximum of 20 interests.
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 embed the posts data into the users data instead of creating a separate document for each post, you will increase the write operation costs for new posts	○	○
If you embed the comments data into the posts data instead of creating a separate document for each comment you will increase the write operation costs for new comments	○	○
If you embed the interests data into the users data instead of creating a separate document for each interest, you will increase the read operation costs for displaying the users and their associated interests	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Yes
Non-relational data increases write costs, but can decrease read costs.
Box 2: Yes
Non-relational data increases write costs, but can decrease read costs.
Box 3: No
Non-relational data increases write costs, but can decrease read costs.

NEW QUESTION 46
- (Exam Topic 2)
You have a database in an Azure Cosmos DB Core (SQL) API account.
You need to create an Azure function that will access the database to retrieve records based on a variable named accountnumber. The solution must protect against SQL injection attacks.

How should you define the command statement in the function?

A. cmd = "SELECT * FROM Persons pWHERE p.accountnumber = 'accountnumber'"
B. cmd = "SELECT * FROM Persons pWHERE p.accountnumber = LIKE @accountnumber"
C. cmd = "SELECT * FROM Persons pWHERE p.accountnumber = @accountnumber"
D. cmd = "SELECT * FROM Persons pWHERE p.accountnumber = '" + accountnumber + "'"

Answer: C

Explanation:
Azure Cosmos DB supports queries with parameters expressed by the familiar @ notation. Parameterized SQL provides robust handling and escaping of user input, and prevents accidental exposure of data through SQL injection.
For example, you can write a query that takes lastName and address.state as parameters, and execute it for various values of lastName and address.state based on user input.
SELECT *
FROM Families f
WHERE f.lastName = @lastName AND f.address.state = @addressState
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/sql/sql-query-parameterized-queries

NEW QUESTION 49
- (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	◯	◯
VM2 can access account 1	◯	◯
VM3 can access account 1	◯	◯

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 50
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

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

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

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

Money Back Guarantee

DP-420 Practice Exam Features:

* DP-420 Questions and Answers Updated Frequently

* DP-420 Practice Questions Verified by Expert Senior Certified Staff

* DP-420 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* DP-420 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year