DP-420 Q&As

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

Pass Microsoft DP-420 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.certbus.com/dp-420.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

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 sett 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 conlainer1 in an Azure Cosmos DB for NoSQL 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 to copy data to another Azure Cosmos DB for NoSQL container.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

QUESTION 2

HOTSPOT

You have an Azure Cosmos DB Core (SQL) API account named storage1 that uses provisioned throughput capacity mode.

The storage1 account contains the databases shown in the following table.

Name	Throughput	Max request units per second (RU/s)	Geo-redundancy	Multi-region writes	Number of regions
db1	Autoscale	5,000	Disabled	Disabled	1
db2	Autoscale	8,000	Enabled	Enabled	3

The databases contain the containers shown in the following table.

Name	Database	Throughput
cn01	db1	Container - autoscale maximum RU/s of 10,000
cn02	db1	Database
cn03	db1	Database
cn04	db1	Database
cn05	db1	Database
cn11	db2	Database
cn12	db2	Database
cn13	db2	Database
cn14	db2	Database
cn15	db2	Database
cn16	db2	Database
cn17	db2	Database
cn18	db2	Database

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
At a minimum, you will be billed for 4,000 RU/s per hour for db1	○	○
The maximum throughput that can be consumed by cn11 is 400 RU/s	○	○
To db2, you can add a new container that uses database throughput	○	○

Correct Answer:

Answer Area

Statements	Yes	No
At a minimum, you will be billed for 4,000 RU/s per hour for db1	○	○
The maximum throughput that can be consumed by cn11 is 400 RU/s	○	○
To db2, you can add a new container that uses database throughput	○	○

Box 1: No

Four containers with 1000 RU/s each.

Box 2: No

Max 8000 RU/s for db2. 8 containers, so 1000 RU/s for each container.

Box 3: Yes

Max 8000 RU/s for db2. 8 containers, so 1000 RU/s for each container. Can very well add an additional container.

Reference:

https://docs.microsoft.com/en-us/azure/cosmos-db/plan-manage-costs

https://azure.microsoft.com/en-us/pricing/details/cosmos-db/

QUESTION 3

You have a container in an Azure Cosmos DB for NoSQL account that stores data about orders. The following is a sample of an order document.

```
{
    "orderId" : "d4a9179b-5ead-43a3-b851-add9a71ac4b6",
    "customerId" : "f6e39103-bdc7-4346-9cfb-45daa4b2becf",
    "orderDate" : "2021-09-29",
    "orderItems" : [...],
    "total" : 12345
}
```

Documents are up to 2 KB.

You plan to receive one million orders daily.

Customers will frequently view then past order history.

You are the evaluating whether to use orderDate as the partition key.

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

NOTE: Each correct selection is worth one point.

A. You will exceed the maximum number of partition key values.

B. Queries will run cross-partition.

C. You will exceed the maximum storage per partition.

D. There will always be a hot partition.

Correct Answer: BD

QUESTION 4

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

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

Which path will be indexed?

A. /product/brand

B. /product/category

C. /product/[ ]/category

D. /product/brand/tailspin

Correct Answer: A

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

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

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

The excludedPaths array contains a single path: /* .

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

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

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

QUESTION 5

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.

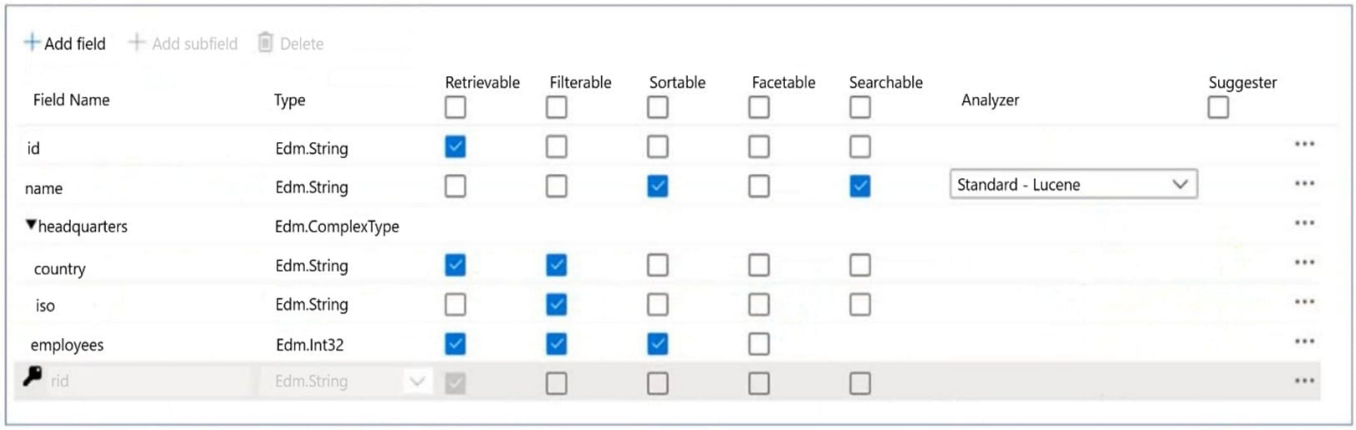

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

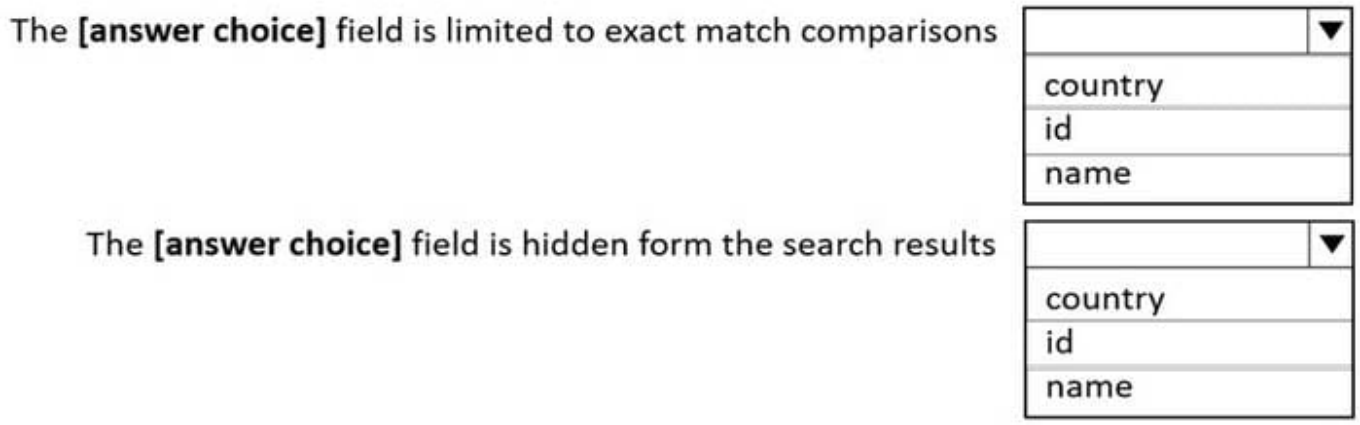

Correct Answer:

Answer Area

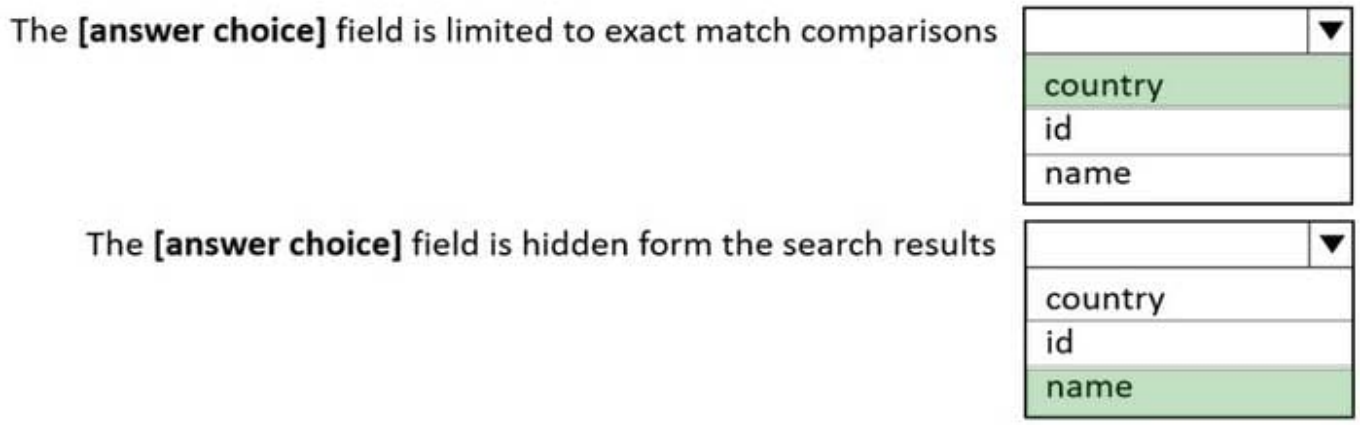

Box 1: country

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

QUESTION 6

DRAG DROP

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

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

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

Select and Place:

Actions

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

Answer Area

Correct Answer:

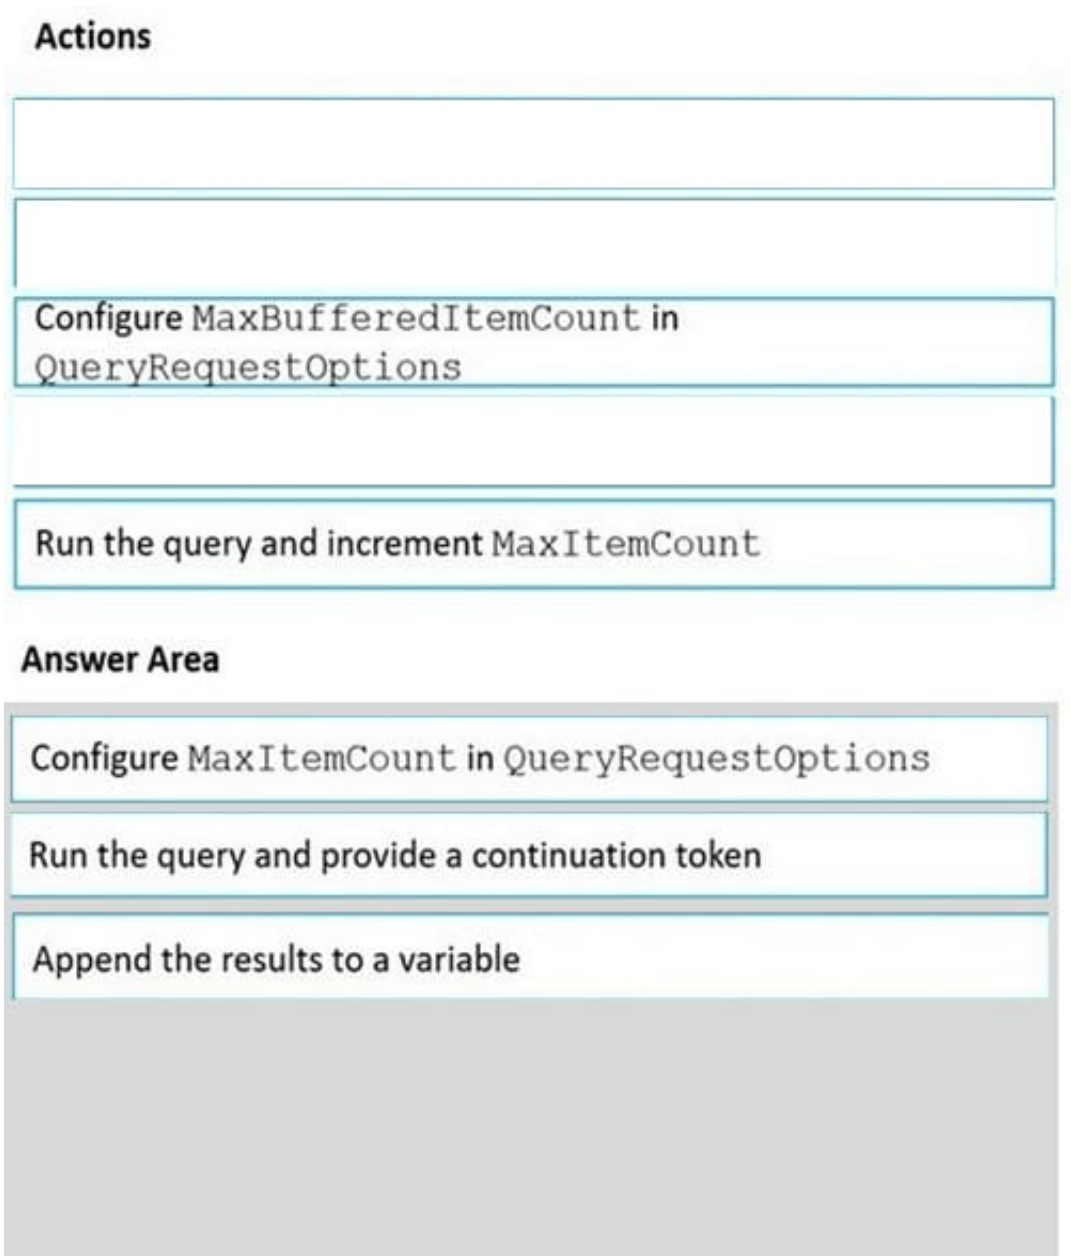

When DefaultTimeToLive is -1 then your Time to Live setting is On (No default)

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

Time to Live on an item:

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

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

Reference:

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

Step 1: Configure the MaxItemCount in QueryRequestOptions

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

Box 2: Run the query and provide a continuation token

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

token.

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

Step 3: Append the results to a variable

Reference:

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

---

QUESTION 7

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

Correct Answer: B

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

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

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

---

QUESTION 8

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

Correct Answer: A

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.

Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/role-based-access-control

---

QUESTION 9

HOTSPOT

You have an Azure Cosmos DB Core (SQL) API account named account1 that has the disableKeyBasedMetadataWriteAccessproperty 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

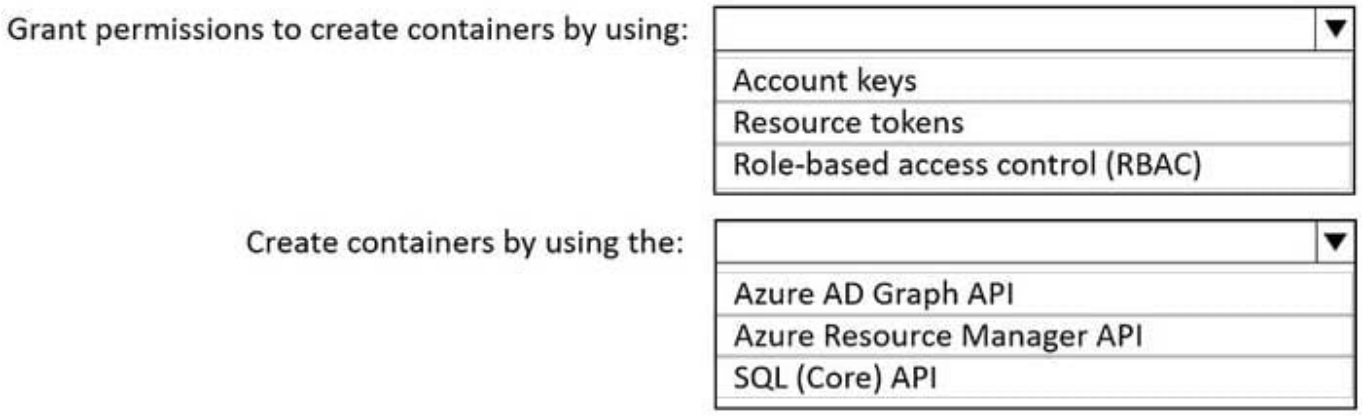

Correct Answer:

Answer Area

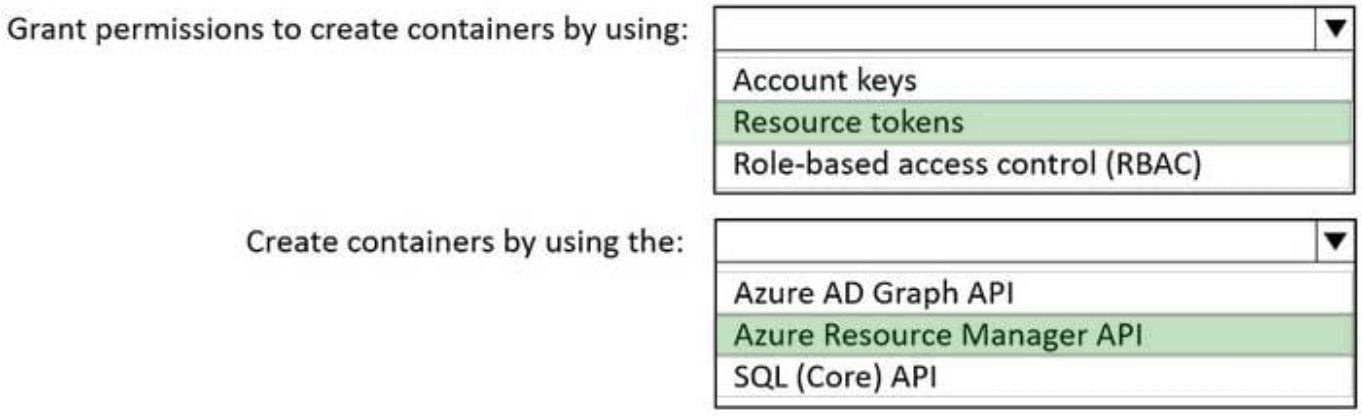

Box 1: Resource tokens

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 10

You have an Azure Cosmos DB for NoSQL account.

You need to create an Azure Monitor query that lists recent modifications to the regional failover policy.

Which Azure Monitor table should you query?

A. CDBPartitionKeyStatistics

B. CDBQueryRunTimeStatistics

C. CDBDataPlaneRequests

D. CDBControlPlaneRequests

Correct Answer: D

QUESTION 11

You plan to create an Azure Cosmos DB account that will use the NoSQL API.

You need to create a grouping strategy for items that will be stored in the account. The solution must ensure that write and read operations on the items can be performed within the same transact!

What should you use to group the items?

A. logical partitions

B. physical partitions

C. databases

D. containers

Correct Answer: A

QUESTION 12

You have a container m an Azure Cosmos DB for NoSQL account.

Data update volumes are unpredictable.

You need to process the change teed of the container by using a web app that has multiple instances. The change feed will be processed by using the change feed processor from the Azure Cosmos DB SDK. The multiple instances must share the workload.

Which three actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Configure the same processor name for all the instances.

B. Configure a different processor name for each instance.

C. Configure a different lease container configuration for each instance.

D. Configure the same instance name for all the instances. 13

E. Configure a different instance name for each instance.

F. Configure the same lease container configuration for all the instances.

Correct Answer: AEF

QUESTION 13

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.olapas 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 OPENROWSETand 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.oltpas the data source.

Correct Answer: AD

Explore analytical store with Apache Spark

1.

Navigate to the Data hub.

2.

Select the Linked tab (1), expand the Azure Cosmos DB group (if you don\\'t see this, select the Refresh button above), then expand the WoodgroveCosmosDb account (2). Right-click on the transactions container (3), select New notebook (4), then select Load to DataFrame (5).

3.

In the generated code within Cell 1 (3), notice that the spark.read format is set to cosmos.olap. This instructs Synapse Link to use the container\\'s analytical store. If we wanted to connect to the transactional store, like to read from the change feed or write to the container, we\\'d use cosmos.oltp instead.

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

QUESTION 14

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 and 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.

Correct Answer: A

When creating a new Azure Cosmos DB account, in the Backup policy tab, choose continuous mode to enable the point in time restore functionality for the new account. With the point-in-time restore, data is restored to a new account, currently you can\\'t restore to an existing account.

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

QUESTION 15

You have an Azure Cosmos DB for NoSQL account named account1 that supports an application named App1. App1 uses the consistent prefix consistency level.

You configure account1 to use a dedicated gateway and integrated cache.

You need to ensure that App1 can use the integrated cache.

Which two actions should you perform for APP1? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Change the connection mode to direct

B. Change the account endpoint to https://account1.sqlx.cosmos.azure.com.

C. Change the consistency level of requests to strong.

D. Change the consistency level of requests to session.

E. Change the account endpoint to https://account1.documents.azure.com

Correct Answer: BD

the Azure Cosmos DB integrated cache is an in-memory cache that is built- in to the Azure Cosmos DB dedicated

gateway. The dedicated gateway is a front-end compute that stores cached data and routes requests to the backend database.

You can choose from a variety of dedicated gateway sizes based on the number of cores and memory needed for your workload1. The integrated cache can reduce the RU consumption and latency of read operations by serving them from

the cache instead of the backend containers2.

For your scenario, to ensure that App1 can use the integrated cache, you should perform these two actions:

Change the account endpoint to https://account1.sqlx.cosmos.azure.com. This is the dedicated gateway endpoint that you need to use to connect to your Azure Cosmos DB account and leverage the integrated cache. The standard gateway

endpoint (https://account1.documents.azure.com) will not use the integrated cache2.

Change the consistency level of requests to session. This is the highest consistency level that is supported by the integrated cache. If you use a higher consistency level (such as strong or bounded staleness), your requests will bypass the

integrated cache and go directly to the backend containers

DP-420 PDF Dumps DP-420 VCE Dumps DP-420 Study Guide