



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.pass4itsure.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**

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

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.

Correct Answer: C

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>

QUESTION 2

HOTSPOT

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.



Hot Area:

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"/>

Correct Answer:

Answer Area

Statements	Yes	No
The query performs a cross-partition query	<input type="radio"/>	<input checked="" type="radio"/>
The query uses an index	<input type="radio"/>	<input checked="" type="radio"/>
Recreating the container with the partition key set to /categoryId will improve the performance of the query	<input checked="" type="radio"/>	<input type="radio"/>

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>

QUESTION 3

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.

Correct Answer: C

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>

QUESTION 4

You have a database named db1 in an Azure Cosmos DB Core (SQL) API account.

You have a third-party application that is exposed through a REST API.

You need to migrate data from the application to a container in db1 on a weekly basis.

What should you use?

- A. Database Migration Assistant
- B. Azure Data Factory
- C. Azure Migrate

Correct Answer: B

You can use Copy Activity in Azure Data Factory to copy data from and to Azure Cosmos DB (SQL API).

The Azure Cosmos DB (SQL API) connector is supported for the following activities:

Copy activity with supported source/sink matrix

Mapping data flow

Lookup activity

Incorrect:

Not A: Azure Migrate provides a centralized hub to assess and migrate on-premises servers, infrastructure, applications, and data to Azure. It assesses on-premises databases and migrate them to Azure SQL Database or to SQL Managed

Instance.

Not C: Data Migration Assistant (DMA) enables you to upgrade to a modern data platform by detecting compatibility issues that can impact database functionality on your new version of SQL Server. It recommends performance and reliability



improvements for your target environment.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-cosmos-db>

QUESTION 5

You have a database in an Azure Cosmos DB for NoSQL account.

The database contains a container named container1. The indexing mode container1 is set to none.

You configure Azure Cognitive Search to extract data from container1 and make the data searchable.

You discover that the Cognitive Search index is missing all the data from the Azure Cosmos DB index.

What should you do to resolve the issue?

- A. Modify The index attributes in Cognitive Search to searchable.
- B. Modify the index attributes in Cognitive Search to Retrievable.
- C. Change the indexing mode of container 1 to consistent-
- D. Modify the indexing policy of container 1 to exclude the / * path

Correct Answer: C

QUESTION 6


HOTSPOT

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.



Total Requests by Status Code



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.

Hot Area:

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

Correct Answer:



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	

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>

QUESTION 7

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.



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
```

A. Option A

B. Option B



C. Option C

D. Option D

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 8

You are implementing an Azure Data Factory data flow that will use an Azure Cosmos DB (SQL API) sink to write a dataset. The data flow will use 2,000 Apache Spark partitions. You need to ensure that the ingestion from each Spark partition is balanced to optimize throughput.

Which sink setting should you configure?

A. Throughput

B. Write throughput budget

C. Batch size

D. Collection action

Correct Answer: C

Batch size: An integer that represents how many objects are being written to Cosmos DB collection in each batch. Usually, starting with the default batch size is sufficient. To further tune this value, note:

Cosmos DB limits single request's size to 2MB. The formula is "Request Size = Single Document Size * Batch Size". If you hit error saying "Request size is too large", reduce the batch size value.

The larger the batch size, the better throughput the service can achieve, while make sure you allocate enough RUs to empower your workload.

Incorrect Answers:

A: Throughput: Set an optional value for the number of RUs you'd like to apply to your CosmosDB collection for each execution of this data flow. Minimum is 400.

B: Write throughput budget: An integer that represents the RUs you want to allocate for this Data Flow write operation, out of the total throughput allocated to the collection.

D: Collection action: Determines whether to recreate the destination collection prior to writing.

None: No action will be done to the collection. Recreate: The collection will get dropped and recreated

Reference: <https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-cosmos-db>

**QUESTION 9**

You plan to create an operational system that will store data in an Azure Cosmos DB or NoSQL account. You need to configure the account to meet the following requirements:

1.

Support Spar* queries.

2.

Support the analysis of data from the last six months.

3.

Only pay for analytical compute when running queries.

Which three actions should you perform?

Each correct answer presents part of the solution. NOTE Each correct selection is worth one point.

A. Create an Azure Synapse linked service.

B. Create a container and set the time to live to six months.

C. Create a container and set the analytical property to six months.

D. Create an Azure Synapse pipeline.

E. Create an Azure Databricks notebook.

F. Enable Azure Synapse Link for the account

Correct Answer: CEF

QUESTION 10

You have a container in an Azure Cosmos DB for NoSQL account. The container stores data about families. Data about parents, children, and pets are stored as separate documents.

Each document contains the address of each family. Members of the same family share the same partition key named family Id

You need to update the address for each member of the same family that share the same address.

The solution must meet the following requirements:

1.

Be atomic consistent isolated, and durable (ACID).

2.



Provide the lowest latency. What should you do?

- A. Update the document of each family member by using a transactional batch operation.
- B. Update the document of each family member separately by using a patch operation.
- C. Update the document of each family member separately and set the consistency level to strong.

Correct Answer: A

QUESTION 11

You have a database named db1 in an Azure Cosmos DB for NoSQL account named account 1.

You need to write JSON data to db1 by using Azure Stream Analytics. The solution must minimize costs.

Which should you do before you can use db1 as an output of Stream Analytics?

- A. In account, add a private endpoint.
- B. In db1, create containers that have a custom indexing policy and analytical store disabled.
- C. In account, enable a dedicated gateway.
- D. In db1, create containers that have an automatic indexing policy and analytical store enabled.

Correct Answer: A

QUESTION 12

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" : []
}
```

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

Correct Answer:

Answer Area

Statements	Yes	No
If you run the following query, the query will run as a cross-partition query <pre>SELECT * FROM c where c.orderDate = "2019-05-03"</pre>	<input checked="" type="radio"/>	<input type="radio"/>
If you run the following query, the query will run as a cross-partition query <pre>SELECT * FROM c where c.customerId = "5425"</pre>	<input checked="" type="radio"/>	<input type="radio"/>
If you run the following query, the query will run as a cross-partition query <pre>SELECT * FROM c where c.orderDate = "2019-05-03" and c.orderId = "9d7816e6-f401-42ba-ad05-0e03de35c0b8"</pre>	<input type="radio"/>	<input checked="" type="radio"/>

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

QUESTION 13

You have operational data in an Azure Cosmos DB for NoSQL database.

Database users report that the performance of the database degrades significantly when a business analytics team runs large Apache Spark-based queries against the database.

You need to reduce the impact that running the Spark-based queries has on the database users.

What should you implement?

- A. Azure Synapse Link
- B. a default consistency level of Consistent Prefix
- C. a default consistency level of Strong
- D. the Spark connector

Correct Answer: A

QUESTION 14

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 leaseCollectionName
- C. Configure the trigger for each function to use a different leaseCollectionName
- D. Configure the trigger for each function to use the same leaseCollectionPrefix

Correct Answer: AC

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>

QUESTION 15

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

Correct Answer: B

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>

[DP-420 PDF Dumps](#)

[DP-420 Exam Questions](#)

[DP-420 Braindumps](#)