



Free Microsoft DP-420 Practice Questions

Shared by Maddox on 13-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

DP-420 Mix Questions IN THIS CASE STUDY

You have an Azure Cosmos DB for NoSQL database named db1 that is monitored by using Azure Monitor. You need to identify the operations that have the highest throughput. Which metric should you use?

Options:

- A- Data Usage
- B- Autoscale Max Throughput
- C- Index Usage
- D- Total Request Units

Answer:

D

Question 2

Question Type: MultipleChoice

You have an Azure Cosmos DB Core (SQL) API account.

You run the following query against a container in the account.

```
SELECT
```

```
IS_NUMBER("1234") AS A,
```

```
IS_NUMBER(1234) AS B,
```

```
IS_NUMBER({prop: 1234}) AS C
```

What is the output of the query?

Options:

- A- [{ 'A': false, 'B': true, 'C': false }]
- B- [{ 'A': true, 'B': false, 'C': true }]
- C- [{ 'A': true, 'B': true, 'C': false }]
- D- [{ 'A': true, 'B': true, 'C': true }]

Answer:

A

Explanation:

IS_NUMBER returns a Boolean value indicating if the type of the specified expression is a number.

'1234' is a string, not a number.

Question 3

Question Type: MultipleChoice

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.

Options:

- 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

Explanation:

<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 4

Question Type: MultipleChoice

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.

Options:

- 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.

Answer:

C, E

Explanation:

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.

Question 5

Question Type: MultipleChoice

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?

Options:

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:



Question 6

Question Type: MultipleChoice

You have an Azure Cosmos DB for NoSQL account1 that is configured for automatic failover. The account1 account has a single read-write region in West US and a and a read region in East US.

You run the following PowerShell command.

```
Update-AzCosmosDBAccountFailoverPriority -ResourceGroupName "rg1" -Name "account1" -FailoverPolicy @("East US", "West US")
```

What is the effect of running the command?

Options:

- A- A manual failover will occur.
- B- The account will be unavailable to writes during the change.
- C- The provisioned throughput for account1 will increase.
- D- The account will be configured for multi-region writes.

Answer:

B

Explanation:

[You can use the Set-AzCosmosDBAccountRegion cmdlet to update the regions that an Azure Cosmos DB account uses. You can use this cmdlet to add a region or change the region failover order. The cmdlet requires a resource group name, an Azure Cosmos DB account name, and a list of regions in desired failover order1.](#)

For your scenario, based on the PowerShell command, you are using the Set-AzCosmosDBAccountRegion cmdlet to update the regions for an Azure Cosmos DB account named account1 that is configured for automatic failover. The command specifies two regions: West US and East US. The effect of running the command is that the account will be configured for multi-region writes.

[Multi-region writes is a feature of Azure Cosmos DB that allows you to write data to any region in your account and have it automatically replicated to all other regions. This feature provides high availability and low latency for write operations across multiple regions. To enable multi-region writes, you need to specify at least two regions in your account and set them as write regions2. In your command, you are setting both West US and East US as write regions by using](#)

[the -IsZoneRedundant parameter with a value of \\$true for both regions.](#)

Question 7

Question Type: MultipleChoice

You have a database named db1 in an Azure Cosmos DB. You have a third-party application that is exposed through an OData endpoint. You need to migrate data from the application to a container in Azure Cosmos DB for NoSQL. What should you use?

Options:

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

Answer:

B

Explanation:

[You can migrate data from various data sources to Azure Cosmos DB using different tools and methods. The choice of the migration tool depends on factors such as the data source, the Azure Cosmos DB API, the size of data, and the expected migration duration¹. Some of the common migration tools are:](#)

[Azure Cosmos DB Data Migration tool: This is an open source tool that can import data to Azure Cosmos DB from sources such as JSON files, MongoDB, SQL Server, CSV files, and Azure Cosmos DB collections. This tool supports the SQL API and the Table API of Azure Cosmos DB².](#)

[Azure Data Factory: This is a cloud-based data integration service that can copy data from various sources to Azure Cosmos DB using connectors. This tool supports the SQL API, MongoDB API, Cassandra API, Gremlin API, and Table API of Azure Cosmos DB³.](#)

[Azure Cosmos DB live data migrator: This is a command-line tool that can migrate data from one Azure Cosmos DB container to another container within the same or different account. This tool supports live migration with minimal downtime and works with any Azure Cosmos DB API⁴.](#)

[For your scenario, if you want to migrate data from a third-party application that is exposed through an OData endpoint to a container in Azure Cosmos DB for NoSQL, you should use Azure Data Factory. Azure Data Factory has an OData connector that can read data from an OData source and write it to an Azure Cosmos DB sink using the SQL API⁵. You can create a copy activity in Azure Data Factory that specifies the OData source and the Azure Cosmos DB sink.](#)

[and run it on demand or on a schedule.](#)

Question 8

Question Type: MultipleChoice

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

You need to create an alert based on a custom Log Analytics query.

Which signal type should you use?

Options:

- A- Metrics
- B- Log
- C- Activity Log
- D- Resource Health

Answer:

B



To Get Premium Files for DP-420 Visit

<https://www.p2pexams.com/products/dp-420>

For More Free Questions Visit

<https://www.p2pexams.com/microsoft/pdf/dp-420>

20%
DISCOUNT

P2P
exams