



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.geekcert.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

You have a container named container1 in an Azure Cosmos DB for NoSQL account named account1 that is set to the session default consistency level. The average size of an item in container1 is 20 KB.

You have an application named App1 that uses the Azure Cosmos DB SDK and performs a point read on the same set of items in container1 every minute.

You need to minimize the consumption of the request units (RUs) associated to the reads by App1.

What should you do?

- A. In account1, change the default consistency level to bounded staleness.
- B. In App1, change the consistency level of read requests to consistent prefix.
- C. In account1, provision a dedicated gateway and integrated cache
- D. In App1, modify the connection policy settings.

Correct Answer: B

The cost of a point read for a 1 KB item is 1 RU. The cost of other operations depends on factors such as item size, indexing policy, consistency level, and query complexity¹. To minimize the consumption of RUs, you can optimize these

factors according to your application needs.

For your scenario, one possible way to minimize the consumption of RUs associated to the reads by App1 is to change the consistency level of read requests to consistent prefix. Consistent prefix is a lower consistency level than session,

which is the default consistency level for Azure Cosmos DB. Lower consistency levels consume fewer RUs than higher consistency levels². Consistent prefix guarantees that reads never see out-of-order writes and that monotonic reads are

preserved¹. This may be suitable for your application if you can tolerate some eventual consistency.

QUESTION 2

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

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>

QUESTION 3

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 4

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: AB

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 5

You need to create a data store for a directory of small and medium-sized businesses (SMBs). The data store must meet the following requirements:

1.

Store companies and the users employed by them. Each company will have less than 1,000 users.

2.

Some users have data that is greater than 2 KB.

3.

Associate each user to only one company.

4.

Provide the ability to browse by company.

5.



Provide the ability to browse the users by company.

6.

Whenever a company or user profile is selected, show a details page for the company and all the related users.

7.

Be optimized for reading data.

Which design should you implement to optimize the data store for reading data?

A. In a directory container, create a document for each company and a document for each user. Use company ID as the partition key.

B. In a company container, create a document for each company. Embed the users into company documents. Use the company ID as the partition key.

C. Create a user container that uses the user ID as the partition key and a company container that container that uses the company ID as the partition key. Add the company ID to each user documents.

D. In a user container, create a document for each user. Embed the company into each user document. Use the user ID as the partition key.

Correct Answer: B

Azure Cosmos DB is a multi-model database that supports various data models, such as documents, key-value, graph, and column-family³. The core content-model of Cosmos DB's database engine is based on atom-record-sequence

(ARS), which allows it to store and query different types of data in a flexible and efficient way³. To optimize the data store for reading data, you should consider the following factors:

The size and shape of your data

The frequency and complexity of your queries

The latency and throughput requirements of your application The trade-offs between storage efficiency and query performance

Based on these factors, one possible design that you could implement is B. In a company container, create a document for each company. Embed the users into company documents. Use the company ID as the partition key.

This design has the following advantages:

It stores companies and users as self-contained documents that can be easily retrieved by company ID¹.

It avoids storing redundant data or creating additional containers for users¹. It allows you to browse by company and browse the users by company with simple queries¹.

It shows a details page for the company and all the related users by fetching a single document¹.

It leverages the benefits of embedding data, such as reducing the number of requests, improving query performance, and simplifying data consistency².

This design also has some limitations, such as:

It may not be suitable for some users who have data that is greater than 2 KB, as it could exceed the maximum



document size limit of 2 MB². It may not be optimal for scenarios where users need to be associated with more than one

company or queried independently from companies². It may not be scalable for companies that have more than 1,000 users, as it could result in hot partitions or throttling². Depending on your specific use case and requirements, you may

need to adjust this design or choose a different one. For example, you could use a hybrid data model that combines embedding and referencing data², or you could use a graph data model that expresses entities and relationships as vertices

and edges.

[DP-420 PDF Dumps](#)

[DP-420 Study Guide](#)

[DP-420 Braindumps](#)