DP-420 Dumps

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

https://www.certleader.com/DP-420-dumps.html

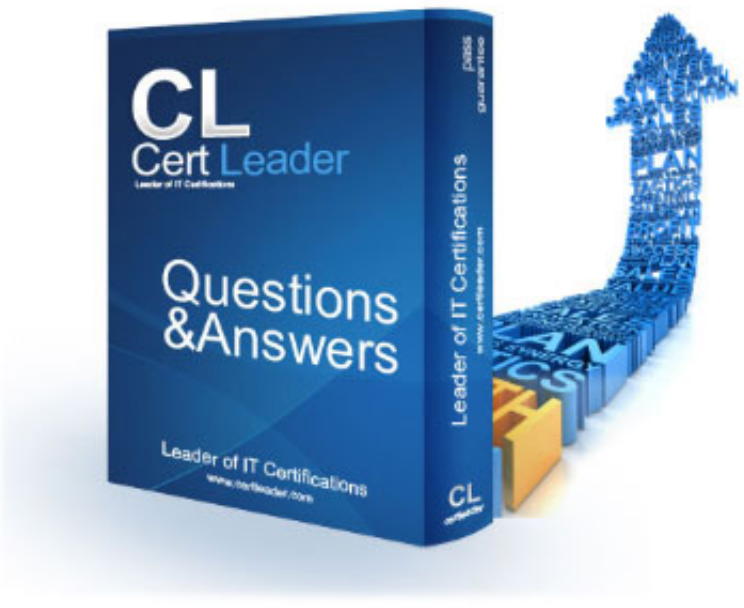

NEW QUESTION 1
- (Exam Topic 1)
You are troubleshooting the current issues caused by the application updates.
Which action can address the application updates issue without affecting the functionality of the application?

A. Enable time to live for the con-product container.
B. Set the default consistency level of account1 to strong.
C. Set the default consistency level of account1 to bounded staleness.
D. Add a custom indexing policy to the con-product container.

Answer: C

Explanation:
Bounded staleness is frequently chosen by globally distributed applications that expect low write latencies but require total global order guarantee. Bounded staleness is great for applications featuring group collaboration and sharing, stock ticker, publish-subscribe/queueing etc.
Scenario: Application updates in con-product frequently cause HTTP status code 429 "Too many requests". You discover that the 429 status code relates to excessive request unit (RU) consumption during the updates.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels

NEW QUESTION 2
- (Exam Topic 1)
You configure multi-region writes for account1.
You need to ensure that App1 supports the new configuration for account1. The solution must meet the business requirements and the product catalog requirements.
What should you do?

A. Set the default consistency level of accountl to bounded staleness.
B. Create a private endpoint connection.
C. Modify the connection policy of App1.
D. Increase the number of request units per second (RU/s) allocated to the con-product and con-productVendor containers.

Answer: D

Explanation:
App1 queries the con-product and con-productVendor containers.
Note: Request unit is a performance currency abstracting the system resources such as CPU, IOPS, and memory that are required to perform the database operations supported by Azure Cosmos DB.
Scenario:
Develop an app named App1 that will run from all locations and query the data in account1.
Once multi-region writes are configured, maximize the performance of App1 queries against the data in account1.
Whenever there are multiple solutions for a requirement, select the solution that provides the best performance, as long as there are no additional costs associated.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels

NEW QUESTION 3
- (Exam Topic 2)
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.

The databases contain the containers shown in the following table.

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
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/

NEW QUESTION 4
- (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.

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

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 5
- (Exam Topic 2)
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?

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.
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/sql/sql-query-is-number

NEW QUESTION 6
- (Exam Topic 2)
The following is a sample of a document in orders.

The orders container uses customerId as the partition key.
You need to provide a report of the total items ordered per month by item type. The solution must meet the following requirements:
Ensure that the report can run as quickly as possible. Minimize the consumption of request units (RUs). What should you do?

A. Configure the report to query orders by using a SQL query.
B. Configure the report to query a new aggregate containe
C. Populate the aggregates by using the change feed.
D. Configure the report to query orders by using a SQL query through a dedicated gateway.
E. Configure the report to query a new aggregate containe
F. Populate the aggregates by using SQL queries that run daily.

Answer: B

Explanation:
You can facilitate aggregate data by using Change Feed and Azure Functions, and then use it for reporting.
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed

NEW QUESTION 7
- (Exam Topic 2)
You have an Azure Cosmos DB Core (SQL) API account.
You configure the diagnostic settings to send all log information to a Log Analytics workspace.
You need to identify when the provisioned request units per second (RU/s) for resources within the account were modified.
You write the following query. AzureDiagnostics
| where Category == "ControlPlaneRequests" What should you include in the query?

A. | where OperationName startswith "AccountUpdateStart"
B. | where OperationName startswith "SqlContainersDelete"
C. | where OperationName startswith "MongoCollectionsThroughputUpdate"
D. | where OperationName startswith "SqlContainersThroughputUpdate"

Answer: A

Explanation:

The following are the operation names in diagnostic logs for different operations: RegionAddStart, RegionAddComplete RegionRemoveStart, RegionRemoveComplete AccountDeleteStart, AccountDeleteComplete RegionFailoverStart, RegionFailoverComplete AccountCreateStart, AccountCreateComplete
AccountUpdateStart, AccountUpdateComplete VirtualNetworkDeleteStart, VirtualNetworkDeleteComplete DiagnosticLogUpdateStart, DiagnosticLogUpdateComplete
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/audit-control-plane-logs

NEW QUESTION 8
- (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 9
- (Exam Topic 2)
You have the indexing policy shown in the following exhibit.

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.

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 10
- (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.

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 10
- (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 11
- (Exam Topic 2)
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

Answer: C

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

NEW QUESTION 13
- (Exam Topic 2)
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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
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

NEW QUESTION 17
- (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.

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

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 22
- (Exam Topic 2)
You have a database named telemetry in an Azure Cosmos DB Core (SQL) API account that stores IoT data. The database contains two containers named readings and devices.
Documents in readings have the following structure.
id
deviceid
timestamp
ownerid
measures (array)
- type
- value
- metricid
Documents in devices have the following structure.
id
deviceid
owner
- ownerid
- emailaddress
- name brand model
For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Yes
Need to join readings and devices.
Box 2: No
Only readings is required. All required fields are in readings.
Box 3: No
Only devices is required. All required fields are in devices.

NEW QUESTION 26
- (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.

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.

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 30
- (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.

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.

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

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 35
- (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 Azure Monitor alert to trigger the function. Does this meet the goal?

A. Yes
B. No

Answer: A

Explanation:
You can set up alerts from the Azure Cosmos DB pane or the Azure Monitor service in the Azure portal. Note: Alerts are used to set up recurring tests to monitor the availability and responsiveness of your Azure
Cosmos DB resources. Alerts can send you a notification in the form of an email, or execute an Azure Function when one of your metrics reaches the threshold or if a specific event is logged in the activity log.
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/create-alerts

NEW QUESTION 40
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your DP-420 Exam with Our Prep Materials Via below:

https://www.certleader.com/DP-420-dumps.html