Microsoft.DP-420.v2024-06-23.q124

Exam Code:	DP-420
Exam Name:	Designing and Implementing Cloud-Native Applications Using Microsoft Azure Cosmos DB
Certification Provider:	Microsoft
Free Question Number:	124
Version:	v2024-06-23
# of views:	691
# of Questions views:	1240
https://www.exam-tests.com/DP-420-exam/Microsoft.DP-420.v2024-06-23.q124.html	

NEW QUESTION: 1
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.
Answer:
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: 2

You have an Azure Cosmos DB for NoSQL account that has multiple write regions.
You need to receive an alert when requests that target the database exceed the available request units per second (RU/s).
Which Azure Monitor signal should you use?
A. Region Removed
B. Document Quota
C. Metadata Requests
D. Data Usage
Answer: B (LEAVE A REPLY)
Explanation
Azure Monitor is a service that provides comprehensive monitoring for Azure resources, including Azure Cosmos DB. You can use Azure Monitor to collect, analyze, and alert on metrics and logs from your Azure Cosmos DB account. You can create alerts for Azure Cosmos DB using Azure Monitor based on the metrics, activity log events, or Log Analytics logs on your account For your scenario, if you want to receive an alert when requests that target the database exceed the available request units per second (RU/s), you should use the Document Quota metric. This metric measures the percentage of RU/s consumed by your account or container. You can create an alert rule on this metric from the Azure portal by following these steps In the Azure portal, select the Azure Cosmos DB account you want to monitor.
Under the Monitoring section of the sidebar, select Alerts, and then select New alert rule.
In the Create alert rule pane, fill out the Scope section by selecting your subscription name and resource type (Azure Cosmos DB accounts).
In the Condition section, select Add condition and choose Document Quota from the list of signals.
In the Configure signal logic pane, specify the threshold value and operator for your alert condition. For example, you can choose Greater than or equal to 90 as the threshold value and operator to receive an alert when your RU/s consumption reaches 90% or more of your provisioned throughput.
In the Alert rule details section, specify a name and description for your alert rule.
In the Actions section, select Add action group and choose how you want to receive notifications for your alert. For example, you can choose Email/SMS/Push/Voice as an action type and enter your email address or phone number as a receiver.
Review your alert rule settings and select Create alert rule to save it.

NEW QUESTION: 3
You have an Azure subscription that contains a resource group named RG1. RG1 contains an Azure Cosmos DB for NoSQL account named account1.
You plan to audit changes to the control plane of account1.
You need to ensure that audit events include the details of the security principal that performs each control plane change.

How should you complete the az command? To answer select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Answer:
Explanation

NEW QUESTION: 4
You have an Azure Synapse Analytics workspace named workspace1 that contains a server less SQL pool.
You have an Azure Table Storage account that stores operational data.
You need to replace the Table storage account with Azure Cosmos DB for NoSQL The solution must meet the following requirements:
* Support Queries from the server less SQL pool.
* Only pay for analytical compute when running queries.
* Ensure that analytical processes do
NOTE: affect operational processes.
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.
Answer:
Explanation

NEW QUESTION: 5
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.
Answer:
Explanation
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

NEW QUESTION: 6
You have an Azure Cosmos DB for NoSQL account.
You plan 10 create a container named container1. The container1 container will store items that include two properties named nm and age The most commonly executed queries will query container1 for a specific name. The following is a sample of the query.
You need to define an opt-in Indexing policy for container1. The solution must meet the following requirements:
* Minimize the number of request units consumed by the queries.
* Ensure that the _etag property is excluded from indexing.
How should you define the indexing poky? To answer, select the appropriate options in the answer are a. NOTE: Each correct selection Is worth one point.
Answer:

NEW QUESTION: 7
You have a container named container1 in an Azure Cosmos DB Core (SQL) API account.
The following is a sample of a document in container1.

```
{
"studentId": "631282",
"firstName": "James",
"lastName": "Smith",
"enrollmentYear": 1990,
"isActivelyEnrolled": true,
"address": {
"street": "",
"city": "",
"stateProvince": "",
"postal": "",
}
}
```

The container1 container has the following indexing policy.

```
{
"indexingMode": "consistent",
"includePaths": [
{
"path": "/*"
},
{
"path": "/address/city/?"
}
```

],
"excludePaths": [
{
"path": "/address/*"
},
{
"path": "/firstName/?"
}
]
}
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.
Answer:

NEW QUESTION: 8
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.
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.
Answer:
Explanation
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

NEW QUESTION: 9
You configure a backup for an Azure Cosmos DB for NoSQL 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.
Answer:

NEW QUESTION: 10
You are designing a data model for an Azure Cosmos DB for NoSQL account.

What are the partition limits for request units per second (RU/s) and storage? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Answer:

NEW QUESTION: 11
You have a container named container1 in an Azure Cosmos DB Core (SQL) API account.
Upserts of items in container1 occur every three seconds.
You have an Azure Functions app named function1 that is supposed to run whenever items are inserted or replaced in container1.
You discover that function1 runs, but not on every upsert.
You need to ensure that function1 processes each upsert within one second of the upsert.
Which property should you change in the Function.json file of function1?
A. checkpointInterval
B. leaseCollectionsThroughput
C. maxItemsPerInvocation
D. feedPollDelay
Answer: D (LEAVE A REPLY)
Explanation
With an upsert operation we can either insert or update an existing record at the same time.
FeedPollDelay: The time (in milliseconds) for the delay between polling a partition for new changes on the feed, after all current changes are drained. Default is 5,000 milliseconds, or 5 seconds.
Reference: https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-cosmosdb-v2-trigger

NEW QUESTION: 12
You have operational data in an Azure Cosmos OB 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 10 reduce the impact that running the Spark-based queries has on the database users.
What should you implement?
A. a default consistency level of Consistent Prefix
B. Azure Synapse Link
C. the Spark connector
D. a default consistency level of Strong
Answer: B (LEAVE A REPLY)

NEW QUESTION: 13

You have an Azure subscription that contains an Azure Cosmos DB for NoSQL account named account1 and a Log Analytics workspace named Workspace1. Workspace 1 stores the logs of account1.
You need to identify which operations used the most request units per second (RU/s) during the last 24 hours.
How should you complete the query? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point
Answer:
Explanation

NEW QUESTION: 14
You have an Azure Synapse Analytics workspace named workspace1 that contains a server less SQL pool.
You have an Azure Table Storage account that stores operational data.
You need to replace the Table storage account with Azure Cosmos DB for NoSQL The solution must meet the following requirements:
* Support Queries from the server less SQL pool.
* Only pay for analytical compute when running queries.
* Ensure that analytical processes do
NOTE: affect operational processes.
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.
Answer:
1 - Create an Azure Cosmos DB for NoSQL account.
2 - Enable Azure Synapse Link.
3 - Create a database and a container that has Analytical store enabled.

NEW QUESTION: 15
You configure a backup for an Azure Cosmos DB for NoSQL 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.
Answer:

NEW QUESTION: 16
You provision Azure resources by using the following Azure Resource Manager (ARM) template.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer:

Explanation

Box 1: No

An alert is triggered when the DB key is regenerated, not when it is used.

Note: The az cosmosdb keys regenerate command regenerates an access key for a Azure Cosmos DB database account.

Box 2: No

Only an SMS action will be taken.

Emailreceivers is empty so no email action is taken.

Box 3: Yes

Yes, an alert is triggered when the DB key is regenerated.

Reference: https://docs.microsoft.com/en-us/cli/azure/cosmosdb/keys

NEW QUESTION: 17

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.

Answer: A,C (LEAVE A REPLY)

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.

NEW QUESTION: 18

You have an Azure Cosmos DB account named account1 that has a default consistency level of session.

You have an app named App1.
You need to ensure that the read operations of App1 can request either bounded staleness or consistent prefix consistency.
What should you modify for each consistency level? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Answer:

NEW QUESTION: 19
You have an Azure Cosmos DB account named account1.
You have several apps that connect to account1 by using the account's secondary key.
You then configure the apps to authenticate by using service principals.
You need to ensure that account1 will only allow apps to connect by using an Azure AD identity.
Which account property should you modify?
A. disableKeyBasedMetadataWriteAccess ,
B. disableLocalAuth
C. userAssignedIdentatxe
D. allowedOrxgins
Answer: B (LEAVE A REPLY)
The disableLocalAuth property is a boolean flag that indicates whether local authentication methods such as primary/secondary keys are disabled for the Azure Cosmos DB account. Setting this property to true improves security by ensuring that Azure Cosmos DB accounts exclusively require Azure Active Directory identities for authentication1.

NEW QUESTION: 20
You have a container in an Azure Cosmos DB for NoSQL account. The database that has a manual throughput of 30,000 request units per second (RU/s). The current consumption details are shewn in the following chart.
Normalized RU Consumption (%) By PartitionKeyRangeID
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.
Answer:

NEW QUESTION: 21
You need to identify which connectivity mode to use when implementing App2. The solution must support the planned changes and meet the business requirements.
Which connectivity mode should you identify?
A. Direct mode over HTTPS
B. Gateway mode (using HTTPS)
C. Direct mode over TCP

Answer: C (LEAVE A REPLY)
Scenario: Develop an app named App2 that will run from the retail stores and query the data in account2. App2 must be limited to a single DNS endpoint when accessing account2.
By using Azure Private Link, you can connect to an Azure Cosmos account via a private endpoint. The private endpoint is a set of private IP addresses in a subnet within your virtual network.
When you're using Private Link with an Azure Cosmos account through a direct mode connection, you can use only the TCP protocol. The HTTP protocol is not currently supported.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-configure-private-endpoints

NEW QUESTION: 22
You have a container m 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:
* Be atomic consistent isolated, and durable (ACID).
* 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.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 23
You have a database in an Azure Cosmos DB for NoSQL account that is configured for multi-region writes.
You need to use the Azure Cosmos DB SDK to implement the conflict resolution policy for a container. The solution must ensure that any conflict sent to the conflict feed.
Solution: You set ConfilictResolutionMode to Custom and you use the default settings for the policy.
Does this meet the goal?
A. Yes
B. No
Answer: B (LEAVE A REPLY)
Explanation
Setting ConflictResolutionMode to Custom and using the default settings for the policy will not ensure that conflicts are sent to the conflict feed. You need to define a custom stored procedure using the

"conflictingItems" parameter to handle conflicts properly.

NEW QUESTION: 24
You need to implement a solution to meet the product catalog requirements.
What should you do to implement the conflict resolution policy.
A. Remove frequently changed field from the index policy of the con-product container.
B. Disable indexing on all fields in the index policy of the con-product container.
C. Set the default consistency level for account1 to eventual.
D. Create a new container and migrate the product catalog data to the new container.
Answer: (SHOW ANSWER)
Topic 1, Litware, inc
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.
Overview
Litware, Inc. is a United States-based grocery retailer. Litware has a main office and a primary datacenter in Seattle. The company has 50 retail stores across the United States and an emerging online presence. Each store connects directly to the internet.
Existing environment. Cloud and Data Service Environments.
Litware has an Azure subscription that contains the resources shown in the following table.
Each container in productdb is configured for manual throughput.
The con-product container stores the company's product catalog data. Each document in con-product includes a con-productvendor value. Most queries targeting the data in con-product are in the following format.
SELECT * FROM con-product p WHERE p.con-productVendor - 'name'
Most queries targeting the data in the con-productVendor container are in the following format
SELECT * FROM con-productVendor pv ORDER BY pv.creditRating, pv.yearFounded Existing environment. Current Problems.
Litware identifies the following issues:
Updates to product categories in the con-productVendor container do not propagate automatically to documents in the con-product container.
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.
Requirements. Planned Changes

Litware plans to implement a new Azure Cosmos DB Core (SQL) API account named account2 that will contain a database named iotdb. The iotdb database will contain two containers named con-iot1 and con-iot2.
Litware plans to make the following changes:
Store the telemetry data in account2.
Configure account1 to support multiple read-write regions.
Implement referential integrity for the con-product container.
Use Azure Functions to send notifications about product updates to different recipients.
Develop an app named App1 that will run from all locations and query the data in account1.
Develop an app named App2 that will run from the retail stores and query the data in account2.
App2 must be limited to a single DNS endpoint when accessing account2.
Requirements. Business Requirements
Litware identifies the following business requirements:
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.
Ensure that Azure Cosmos DB costs for IoT-related processing are predictable.
Minimize the number of firewall changes in the retail stores.
Requirements. Product Catalog Requirements
Litware identifies the following requirements for the product catalog:
Implement a custom conflict resolution policy for the product catalog data.
Minimize the frequency of errors during updates of the con-product container.
Once multi-region writes are configured, maximize the performance of App1 queries against the data in account1.
Trigger the execution of two Azure functions following every update to any document in the con-product container.

NEW QUESTION: 25
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.
Answer: (SHOW ANSWER)
Explanation

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.

NEW QUESTION: 26
You have an Azure Cosmos DB for NoSQL account.
You plan 10 create a container named container1. The container1 container will store items that include two properties named nm and age The most commonly executed queries will query container1 for a specific name. The following is a sample of the query.
You need to define an opt-in Indexing policy for container1. The solution must meet the following requirements:
* Minimize the number of request units consumed by the queries.
* Ensure that the _etag property is excluded from indexing.
How should you define the indexing poky? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection Is worth one point.
Answer:
Explanation

NEW QUESTION: 27
You have an Azure Cosmos DB for NoSQL account that has multiple write regions.
You need to receive an alert when requests that target the database exceed the available request units per second (RU/s).
Which Azure Monitor signal should you use?
A. Region Removed
B. Provisioned Throughput
C. Metadata Requests
D. Data Usage
Answer: C (LEAVE A REPLY)
Explanation
Azure Monitor is a service that provides comprehensive monitoring for Azure resources, including Azure Cosmos DB. You can use Azure Monitor to collect, analyze, and alert on metrics and logs from your Azure Cosmos DB account. You can create alerts for Azure Cosmos DB using Azure Monitor based on the metrics, activity log events, or Log Analytics logs on your account1.

For your scenario, if you want to receive an alert when requests that target the database exceed the available request units per second (RU/s), you should use the Document Quota metric. This metric measures the percentage of RU/s consumed by your account or container. You can create an alert rule on this metric from the Azure portal by following these steps2:
In the Azure portal, select the Azure Cosmos DB account you want to monitor.
Under the Monitoring section of the sidebar, select Alerts, and then select New alert rule.
In the Create alert rule pane, fill out the Scope section by selecting your subscription name and resource type (Azure Cosmos DB accounts).
In the Condition section, select Add condition and choose Document Quota from the list of signals.
In the Configure signal logic pane, specify the threshold value and operator for your alert condition. For example, you can choose Greater than or equal to 90 as the threshold value and operator to receive an alert when your RU/s consumption reaches 90% or more of your provisioned throughput.
In the Alert rule details section, specify a name and description for your alert rule.
In the Actions section, select Add action group and choose how you want to receive notifications for your alert. For example, you can choose Email/SMS/Push/Voice as an action type and enter your email address or phone number as a receiver.
Review your alert rule settings and select Create alert rule to save it.

NEW QUESTION: 28
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.
Answer:
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: 29

You plan to create an Azure Cosmos DB database named db1 that will contain two containers. One of the containers will contain blog posts, and the other will contain users. Each item in the blog post container will include:

* A single blog post
* All the comments associated to the blog post
* The names of the users who created the blog post and added the comments.

You need to design a solution to update usernames m the user container without causing data integrity issues.

The solution must minimize administrative and development effort. What should you include in the solution?

To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer:

Explanation

NEW QUESTION: 30

You have an Apache Spark pool in Azure Synapse Analytics that runs the following Python code in a notebook.

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer:

Explanation

New and updated orders will be added to contoso-erp.orders: Yes

The code performs bulk data ingestion from contoso-app: No

Both contoso-app and contoso-erp have Analytics store enabled: Yes

The code uses the spark.readStream method to read data from a container named orders in a database named contoso-app. The data is then filtered by a condition and written to another container named orders in a database named contoso-erp using the spark.writeStream method. The write mode is set to "append", which means that new and updated orders will be added to the destination container1.

The code does not perform bulk data ingestion from contoso-app, but rather stream processing. Bulk data ingestion is a process of loading large amounts of data into a data store in batches. Stream processing is a process of continuously processing data as it arrives in real-time2.

Both contoso-app and contoso-erp have Analytics store enabled, because they are both accessed by Spark pools using the spark.cosmos.oltp method. This method requires that the containers have Analytics store enabled, which is a feature that allows Spark pools to query data stored in Azure Cosmos DB containers using SQL APIs3.

NEW QUESTION: 31

You plan to implement con-iot1 and con-iot2.

You need to configure the default Time to Live setting for each container. The solution must meet the IoT telemetry requirements.
What should you configure? To answer, select the appropriate options in the answer NOTE: Each correct selection is worth one point.

Answer:

Explanation

Box 1 = On (no default)For con-iot1, you need to configure the default TTL setting to On (no default), which means that items in this container do not expire by default, but you can override the TTL value on a per-item basis. This meets the requirement of retaining all telemetry data unless overridden1.

Box 2 = On (3600 seconds)For con-iot2, you need to configure the default TTL setting to On (3600 seconds), which means that items in this container will expire 3600 seconds (one hour) after their last modified time. This meets the requirement of deleting all telemetry data older than one hour1.

NEW QUESTION: 32

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.
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.

Answer:

Explanation

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

NEW QUESTION: 33
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.
Which path will be indexed?
A. /product/brand
B. /product/category
C. /product/[ ]/category
D. /product/brand/tailspin
Answer: (SHOW ANSWER)
Explanation
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.

NEW QUESTION: 34
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.
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.
Answer:
Explanation
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

NEW QUESTION: 35

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.
Answer:

NEW QUESTION: 36
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. Cosmos DB Built-in Data Contributor only
C. DocumentDB Account Contributor only
D. DocumentDB Account Contributor and Cosmos DB Built-in Data Contributor
Answer: A (LEAVE A REPLY)
Explanation
Cosmos DB Operator: Can provision Azure Cosmos accounts, databases, and containers.
Cannot access any data or use Data Explorer.
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/role-based-access-control

NEW QUESTION: 37
You have an Azure Cosmos DB container named container! that has a provisioned throughput and two physical partitions. You monitor the following metrics for container1
* Normalized RU consumption
* The percentage of requests that have an HTTP status code of 429
You need to confirm that container1 is configured to maximize resource utilization.
What are the optimal values for each metric? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Answer:
Explanation

NEW QUESTION: 38
You plan to store order data in Azure Cosmos DB for NoSQL account. The data contains information about orders and their associated items.
You need to develop a model that supports order read operations. The solution must minimize the number or requests.

A. Create a single database that contains one container. Store orders and order items in separate documents in the container.
B. Create a single database that contains one container. Create a separate document for each order and embed the order items into the order documents.
C. Create a database for orders and a database for order items.
D. Create a single database that contains a container for order and a container for order items.
Answer: B (LEAVE A REPLY)
Explanation
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 develop a model that supports order read operations and minimizes the number of requests, 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 model that you could implement is B. Create a single database that contains one container. Create a separate document for each order and embed the order items into the order documents.
This model has the following advantages:
It stores orders and order items as self-contained documents that can be easily retrieved by order ID1.
It avoids storing redundant data or creating additional containers for order items1.
It allows you to view the order history of a customer with simple queries1.
It leverages the benefits of embedding data, such as reducing the number of requests, improving query performance, and simplifying data consistency2.
This model also has some limitations, such as:
It may not be suitable for some order items that 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 order items need to be queried independently from orders or aggregated by other criteria2.
It may not support transactions across multiple orders or customers, as transactions are scoped to a single logical partition2.
Depending on your specific use case and requirements, you may need to adjust this model 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.

NEW QUESTION: 39
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.
Answer:
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: 40
You have an Azure Cosmos DB for NoSQL account named account that has the disablekeyBasedletadatwriteAccess property enabled.
You are developing an app named App1 that will be used by a use1 named DevUser1 to create containers in account1. DevUser1 has a non-privileged user account in the 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.
Answer:
Explanation

NEW QUESTION: 41
You have a container that stores data about families. The following is a sample document.
For each of the following statements, select Yes if the statement is true. otherwise, select No.
NOTE: Each correct selection is worth one point.
Answer:

NEW QUESTION: 42
You plan to use a multi-region Azure Cosmos DB for NoSQL account to store data for a new application suite. The suite contains the applications shown in the following table.
Each application should use the weakest consistency level possible.
Which consistency level should you configure for each application? To answer, select the appropriate options in the answer are a. NOTE: Each correct selection is worth one point.
Answer:

NEW QUESTION: 43
You have a container named container1 in an Azure Cosmos DB Core (SQL) API account.
The following is a sample of a document in container1.

```
{
"studentId": "631282",
"firstName": "James",
"lastName": "Smith",
"enrollmentYear": 1990,
"isActivelyEnrolled": true,
"address": {
"street": "",
"city": "",
"stateProvince": "",
"postal": "",
}
}
```

The container1 container has the following indexing policy.

```
{
"indexingMode": "consistent",
"includePaths": [
{
"path": "/*"
},
{
"path": "/address/city/?"
}
],
"excludePaths": [
{
"path": "/address/*"
},
{
```

"path": "/firstName/?"
}
]
}
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer:

Explanation
Box 1: Yes
"path": "/*" is in includePaths.
Include the root path to selectively exclude paths that don't need to be indexed. This is the recommended approach as it lets Azure Cosmos DB proactively index any new property that may be added to your model.
Box 2: No
"path": "/firstName/?" is in excludePaths.
Box 3: Yes
"path": "/address/city/?" is in includePaths
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/index-policy

NEW QUESTION: 44

You have a container in an Azure Cosmos DB Core (SQL) API account.
You need to use the Azure Cosmos DB SDK to replace a document by using optimistic concurrency.
What should you include in the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer:

Explanation
Box 1: ConsistencyLevel
The ItemRequestOptions Class ConsistencyLevel property gets or sets the consistency level required for the request in the Azure Cosmos DB service.
Azure Cosmos DB offers 5 different consistency levels. Strong, Bounded Staleness, Session, Consistent Prefix and Eventual - in order of strongest to weakest consistency.
Box 2: _etag
The ItemRequestOptions class helped us implement optimistic concurrency by specifying that we wanted the SDK to use the If-Match header to allow the server to decide whether a resource should be updated. The If-Match value is the ETag value to be checked against. If the ETag value matches the server ETag value, the resource is updated.
Reference:
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.itemrequestoptions

https://cosmosdb.github.io/labs/dotnet/labs/10-concurrency-control.html

NEW QUESTION: 45
You have a multi-region Azure Cosmos DB account named account1 that has a default consistency level of strong.
You have an app named App1 that is configured to request a consistency level of session.
How will the read and write operations of App1 be handled? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Answer:

NEW QUESTION: 46
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 container. Populate the aggregates by using the change feed.
C. Configure the report to query orders by using a SQL query through a dedicated gateway.
D. Configure the report to query a new aggregate container. Populate the aggregates by using SQL queries that run daily.
Answer: (SHOW ANSWER)
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: 47

You have three containers in an Azure Cosmos DB Core (SQL) API account as shown in the following table.
You have the following Azure functions:
A function named Fn1 that reads the change feed of cn1
A function named Fn2 that reads the change feed of cn2
A function named Fn3 that reads the change feed of cn3
You perform the following actions:
Delete an item named item1 from cn1.
Update an item named item2 in cn2.
For an item named item3 in cn3, update the item time to live to 3,600 seconds.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.
Answer:
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/sql/change-feed-design-patterns
https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed

NEW QUESTION: 48
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.
Answer:
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: 49
You plan to implement con-iot1 and con-iot2.
You need to configure the default Time to Live setting for each container. The solution must meet the loT telemetry requirements.
What should you configure? To answer, select the appropriate options in the answer NOTE: Each correct selection is worth one point.
Answer:

NEW QUESTION: 50

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 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
Answer: B (LEAVE A REPLY)
Explanation
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.
The following diagram represents the data flow and components involved in the solution:
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/sql/changefeed-ecommerce-solution

NEW QUESTION: 51
You have a multi-region Azure Cosmos DB account named account1 that has a default consistency level of strong.
You have an app named App1 that is configured to request a consistency level of session.
How will the read and write operations of App1 be handled? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Answer:
Explanation
Box 1 = Write and replicate data to every region synchronously
This is because the write concern is mapped to the default consistency level configured on your Azure Cosmos DB account2, which is strong in this case. Strong consistency ensures that every write operation is synchronously committed to every region associated with your Azure Cosmos DB account1. The request level consistency level of session only applies to the read operations of App11.
Box 2: That has the lowest estimated latency to the client

This is because the read operations of App1 will use the session consistency level that is specified in the request options. Session consistency is a client-centric consistency model that guarantees monotonic reads, monotonic writes, and read-your-own-writes within a session. A session is scoped to a client connection or a stored procedure execution. Session consistency allows clients to read from any region that has the lowest latency to the client.

NEW QUESTION: 52

You are developing an application that will connect to an Azure Cosmos DB for NoSQL account. The account has a single readme region and one agonal read region. The regions are configured for automatic failover.

The account has the following connect strings. (Line numbers are included for reference only.)

For each of the following statements, select Yes if the statement is true. otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer:

Explanation

If the primary write region fails, applications that write to the database must use a different connection string to continue to use the service. = NOYou do not need to use a different connection string to continue to use the service if the primary write region fails. This is because Azure Cosmos DB supports automatic failover, which means that it will automatically switch the primary write region to another region in case of a regional outage2

. The application does not need to change the connection string or specify the failover priority3. The connection string contains a list of all the regions associated with your account, and Azure Cosmos DB will route the requests to the appropriate region based on the availability and latency1.

The primary Read-Only SQL Connection String and the Secondary Read-Only SQL Connection String will connect to different regions from an application running in the East US Azure region = YesThe primary read-only SQL connection string and the secondary read-only SQL connection string will connect to different regions from an application running in the East US Azure region. This is because the primary read-only SQL connection string contains the endpoint for the East US region, which is the same as the primary write region.

The secondary read-only SQL connection string contains the endpoint for the West US region, which is the additional read region. Therefore, if an application running in the East US Azure region uses these connection strings, it will connect to different regions depending on which one it chooses.

Applications can choose from which region by setting the PreferredLocations property within their connection properties = Yes Applications can choose from which region by setting the PreferredLocations property within their connection properties. This property allows you to specify a list of regions that you prefer to read from based on their proximity to your application2. Azure Cosmos DB will route the requests to the appropriate region based on the availability and latency1. You can also set the ApplicationRegion property to the region where your application is

deployed, and Azure Cosmos DB will automatically populate the PreferredLocations property based on the geo-proximity from that location1.

NEW QUESTION: 53
You plan to create an Azure Cosmos DB Core (SQL) API account that will use customer-managed keys stored in Azure Key Vault.
You need to configure an access policy in Key Vault to allow Azure Cosmos DB access to the keys.
Which three permissions should you enable in the access policy? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.
A. Wrap Key
B. Get
C. List
D. Update
E. Sign
F. Verify
G. Unwrap Key
Answer: (SHOW ANSWER)
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-setup-cmk

NEW QUESTION: 54
You configure a backup for an Azure Cosmos DB for NoSQL 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.
Answer:
Explanation
Box 1 = The current backup policy provides protection for: 2 Hours Azure Cosmos DB automatically takes backups of your data at regular intervals. The backup interval and the retention period can be configured from the Azure portal. You can also choose between two backup modes: periodic backup mode and continuous backup mode. Periodic backup mode is the default mode for all existing accounts and it takes a full backup of your database every 4 hours by default. Continuous backup mode is a new mode that allows you to restore to any point of time within either 7 or 30 days1.
For your scenario, based on the exhibit, you have configured a backup for an Azure Cosmos DB for NoSQL account using the periodic backup mode with a backup interval of 1 hour and a retention period of 2 hours.
This means that Azure Cosmos DB will take a full backup of your database every hour and keep only the latest two backups. Therefore, the current backup policy provides protection for 2 hours.

Box 2: In case of emergency, you must (answer choice) to restore the backup = create a support ticket Azure Cosmos DB automatically takes backups of your data at regular intervals. You can configure the backup interval and the retention period from the Azure portal. You can also choose between two backup modes:
periodic backup mode and continuous backup mode. Periodic backup mode is the default mode for all existing accounts and it takes a full backup of your database every 4 hours by default. Continuous backup mode is a new mode that allows you to restore to any point of time within either 7 or 30 days1.
For your scenario, based on the exhibit, you have configured a backup for an Azure Cosmos DB for NoSQL account using the periodic backup mode with a backup interval of 1 hour and a retention period of 2 hours.
This means that Azure Cosmos DB will take a full backup of your database every hour and keep only the latest two backups. In case of emergency, you must create a support ticket to restore the backup. This is the answer to your question.
To restore data from a periodic backup, you need to create a support request with Azure Cosmos DB team and provide the following information:
* The name of your Azure Cosmos DB account
* The name of the database or container that you want to restore
* The date and time (in UTC) that you want to restore from
* The name of the target Azure Cosmos DB account where you want to restore the data
* The name of the target resource group where you want to restore the data The Azure Cosmos DB team will then initiate the restore process and notify you when it is completed2.

NEW QUESTION: 55
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.
Answer:
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: 56
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. CDBControlPlaneRequests
D. CDBDataPlaneRequests
Answer: C (LEAVE A REPLY)

NEW QUESTION: 57
You need to recommend indexes for con-product and con-productVendor. The solution must meet the product catalog requirements and the business requirements.
Which type of index should you recommend for each container? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Answer:
Explanation

NEW QUESTION: 58
You have a container in an Azure Cosmos DB for NoSQL account. The database that has a manual throughput of 30,000 request units per second (RU/s). The current consumption details are shewn in the following chart.
Normalized RU Consumption (%) By PartitionKeyRangeID
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.
Answer:

NEW QUESTION: 59
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 container! 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 that uses the Azure Cosmos 08 for NoSQL change feed as a trigger and an Azure event hub as the output.
Does this meet the goal?
A. No
B. Yes
Answer: B (LEAVE A REPLY)

NEW QUESTION: 60
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.
Answer:
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: 61
You have a container that stores data about families. The following is a sample document.
For each of the following statements, select Yes if the statement is true. otherwise, select No.
NOTE: Each correct selection is worth one point.
Answer:
Explanation
Children who do no have parents defined will appear on the list = NO
Children who do not have parents defined will not appear on the list. This is because the document schema defines the children property as an array of objects that contain the firstName and gender properties of each child, as well as a parents property that references the id values of

the parents. If a child does not have parents defined, it means that the parents property is either missing or empty for that child. Therefore, such a child will not be included in the list of children who have parents defined.
Children who have more than one pet will appear on the list multiple times. = Yes Children who have more than one pet will appear on the list multiple times. This is because the document schema defines the pets property as an array of objects that contain the givenName and type properties of each pet, as well as a children property that references the id values of the children who own the pet. If a child has more than one pet, it means that the child's id value will appear in the children property of multiple pet objects. Therefore, such a child will be included in the list of children who have pets multiple times.
Children who do no have pets defined will appear on the list = No
Children who do not have pets defined will not appear on the list. This is because the document schema defines the pets property as an array of objects that contain the givenName and type properties of each pet, as well as a children property that references the id values of the children who own the pet. If a child does not have pets defined, it means that the child's id value does not appear in the children property of any pet object.
Therefore, such a child will not be included in the list of children who have pets defined.

NEW QUESTION: 62
You have a database named db1 in an Azure Cosmos DB for NoSQL
You are designing an application that will use dbl.
In db1, you are creating a new container named coll1 that will store in coll1.
The following is a sample of a document that will be stored in coll1.
The application will have the following characteristics:
* New orders will be created frequently by different customers.
* Customers will often view their past order history.
You need to select the partition key value for coll1 to support the application. The solution must minimize costs.
To what should you set the partition key?
A. id
B. customerId
C. orderDate

D. orderId

Answer: C (LEAVE A REPLY)

Based on the characteristics of the application and the provided document structure, the most suitable partition key value for coll1 in the given scenario would be the customerId, Option B. The application frequently creates new orders by different customers and customers often view their past order history. Using customerId as the partition key would ensure that all orders associated with a particular customer are stored in the same partition. This enables efficient querying of past order history for a specific customer and reduces cross-partition queries, resulting in lower costs and improved performance.

a partition key is a JSON property (or path) within your documents that is used by Azure Cosmos DB to distribute data among multiple partitions3. A partition key should have a high cardinality, which means it should have many distinct values, such as hundreds or thousands1. A partition key should also align with the most common query patterns of your application, so that you can efficiently retrieve data by using the partition key value1.

Based on these criteria, one possible partition key that you could use for coll1 is B. customerId. This partition key has the following advantages:

It has a high cardinality, as each customer will have a unique ID3.

It aligns with the query patterns of the application, as customers will often view their past order history3.

It minimizes costs, as it reduces the number of cross-partition queries and optimizes the storage and throughput utilization1.

This partition key also has some limitations, such as:

It may not be optimal for scenarios where orders need to be queried independently from customers or aggregated by date or other criteria3.

It may result in hot partitions or throttling if some customers create orders more frequently than others or have more data than others1.

It may not support transactions across multiple customers, as transactions are scoped to a single logical partition2.

Depending on your specific use case and requirements, you may need to adjust this partition key or choose a different one. For example, you could use a synthetic partition key that concatenates multiple properties of an item2, or you could use a partition key with a random or pre-calculated suffix to distribute the workload more evenly2.

NEW QUESTION: 63

You plan to store order data in Azure Cosmos DB for NoSQL account. The data contains information about orders and their associated items.

You need to develop a model that supports order read operations. The solution must minimize the number or requests.

A. Create a single database that contains one container. Store orders and order items in separate documents in the container.

B. Create a single database that contains one container. Create a separate document for each order and embed the order items into the order documents.
C. Create a database for orders and a database for order items.
D. Create a single database that contains a container for order and a container for order items.
Answer: B (LEAVE A REPLY)
Explanation
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 develop a model that supports order read operations and minimizes the number of requests, 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 model that you could implement is B. Create a single database that contains one container. Create a separate document for each order and embed the order items into the order documents.
This model has the following advantages:
* It stores orders and order items as self-contained documents that can be easily retrieved by order ID1.
* It avoids storing redundant data or creating additional containers for order items1.
* It allows you to view the order history of a customer with simple queries1.
* It leverages the benefits of embedding data, such as reducing the number of requests, improving query performance, and simplifying data consistency2.
This model also has some limitations, such as:
* It may not be suitable for some order items that 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 order items need to be queried independently from orders or aggregated by other criteria
* It may not support transactions across multiple orders or customers, as transactions are scoped to a single logical partition2.
Depending on your specific use case and requirements, you may need to adjust this model 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.

NEW QUESTION: 64

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 (LEAVE A REPLY)
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: 65
You have three containers in an Azure Cosmos DB Core (SQL) API account as shown in the following table.
You have the following Azure functions:
A function named Fn1 that reads the change feed of cn1
A function named Fn2 that reads the change feed of cn2
A function named Fn3 that reads the change feed of cn3
You perform the following actions:
Delete an item named item1 from cn1.
Update an item named item2 in cn2.
For an item named item3 in cn3, update the item time to live to 3,600 seconds.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.
Answer:
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/sql/change-feed-design-patterns
https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed

NEW QUESTION: 66
You have a database named db1 in an Azure Cosmos DB for NoSQL
You are designing an application that will use dbl.

In db1, you are creating a new container named coll1 that will store in coll1.
The following is a sample of a document that will be stored in coll1.
The application will have the following characteristics:
* New orders will be created frequently by different customers.
* Customers will often view their past order history.
You need to select the partition key value for coll1 to support the application. The solution must minimize costs.
To what should you set the partition key?
A. id
B. customerId
C. orderDate
D. orderId
Answer: B (LEAVE A REPLY)
Explanation
Based on the characteristics of the application and the provided document structure, the most suitable partition key value for coll1 in the given scenario would be the customerId, Option B.
The application frequently creates new orders by different customers and customers often view their past order history. Using customerId as the partition key would ensure that all orders associated with a particular customer are stored in the same partition. This enables efficient querying of past order history for a specific customer and reduces cross-partition queries, resulting in lower costs and improved performance.
a partition key is a JSON property (or path) within your documents that is used by Azure Cosmos DB to distribute data among multiple partitions3. A partition key should have a high cardinality, which means it should have many distinct values, such as hundreds or thousands1. A partition key should also align with the most common query patterns of your application, so that you can efficiently retrieve data by using the partition key value1.
Based on these criteria, one possible partition key that you could use for coll1 is B. customerId. This partition key has the following advantages:
* It has a high cardinality, as each customer will have a unique ID
* It aligns with the query patterns of the application, as customers will often view their past order history3.
* It minimizes costs, as it reduces the number of cross-partition queries and optimizes the storage and throughput utilization1.
This partition key also has some limitations, such as:
* It may not be optimal for scenarios where orders need to be queried independently from customers or aggregated by date or other criteria
* It may result in hot partitions or throttling if some customers create orders more frequently than others or have more data than others
* It may not support transactions across multiple customers, as transactions are scoped to a single logical partition2.

Depending on your specific use case and requirements, you may need to adjust this partition key or choose a different one. For example, you could use a synthetic partition key that concatenates multiple properties of an item2, or you could use a partition key with a random or pre-calculated suffix to distribute the workload more evenly2.

NEW QUESTION: 67
You have a database named telemetry in an Azure Cosmos DB Core (SQL) API account that stores IoT dat a. 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.
Answer:

NEW QUESTION: 68
You have a container in an Azure Cosmos DB Core (SQL) API account.
You need to use the Azure Cosmos DB SDK to replace a document by using optimistic concurrency.
What should you include in the code? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Answer:
Reference:
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.itemrequestoptions
https://cosmosdb.github.io/labs/dotnet/labs/10-concurrency-control.html

NEW QUESTION: 69
You have a container in an Azure Cosmos DB Core (SQL) API account. The container stores telemetry data from IoT devices. The container uses telemetryId as the partition key and has a throughput of 1,000 request units per second (RU/s). Approximately 5,000 IoT devices submit data every five minutes by using the same telemetryId value.
You have an application that performs analytics on the data and frequently reads telemetry data for a single IoT device to perform trend analysis.
The following is a sample of a document in the container.
You need to reduce the amount of request units (RUs) consumed by the analytics application.
What should you do?
A. Decrease the offerThroughput value for the container.
B. Increase the offerThroughput value for the container.
C. Move the data to a new container that has a partition key of deviceId.
D. Move the data to a new container that uses a partition key of date.
Answer: C (LEAVE A REPLY)
The partition key is what will determine how data is routed in the various partitions by Cosmos DB and needs to make sense in the context of your specific scenario. The IoT Device ID is generally the "natural" partition key for IoT applications.

NEW QUESTION: 70
You have an Azure Cosmos DB for NoSQL account named account1 that has a single read-write region and one additional read region. Account1 uses the strong default consistency level.
You have an application that uses the eventual consistency level when submitting requests to account1.
How will writes from the application be handled?
A. Writes will use the strong consistency level.
B. Azure Cosmos DB will reject writes from the application.
C. The write order is not guaranteed during replication.
D. Writes will use the eventual consistency level.
Answer: A (LEAVE A REPLY)
Explanation
This is because the write concern is mapped to the default consistency level configured on your Azure Cosmos DB account, which is strong in this case. Strong consistency ensures that every write operation is synchronously committed to every region associated with your Azure Cosmos DB account. The eventual consistency level that the application uses only applies to the read operations. Eventual consistency offers higher availability and better performance, but it does not guarantee the order or latency of the reads.

NEW QUESTION: 71
You have a database in an Azure Cosmos DB Core (SQL) API account.

You need to create an Azure function that will access the database to retrieve records based on a variable named accountnumber. The solution must protect against SQL injection attacks.
How should you define the command statement in the function?
A. cmd = "SELECT * FROM Persons p
WHERE p.accountnumber = 'accountnumber'"
B. cmd = "SELECT * FROM Persons p
WHERE p.accountnumber = LIKE @accountnumber"
C. cmd = "SELECT * FROM Persons p
WHERE p.accountnumber = @accountnumber"
D. cmd = "SELECT * FROM Persons p
WHERE p.accountnumber = '" + accountnumber + "'"
Answer: C (LEAVE A REPLY)
Azure Cosmos DB supports queries with parameters expressed by the familiar @ notation. Parameterized SQL provides robust handling and escaping of user input, and prevents accidental exposure of data through SQL injection.
For example, you can write a query that takes lastName and address.state as parameters, and execute it for various values of lastName and address.state based on user input.
SELECT *
FROM Families f
WHERE f.lastName = @lastName AND f.address.state = @addressState

NEW QUESTION: 72
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.
Answer:
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: 73
You need to select the partition key for con-iot1. The solution must meet the IoT telemetry requirements.
What should you select?
A. the timestamp
B. the humidity
C. the temperature
D. the device ID
Answer: D (LEAVE A REPLY)
The partition key is what will determine how data is routed in the various partitions by Cosmos DB and needs to make sense in the context of your specific scenario. The IoT Device ID is generally the "natural" partition key for IoT applications.
Scenario: The iotdb database will contain two containers named con-iot1 and con-iot2.
Ensure that Azure Cosmos DB costs for IoT-related processing are predictable.
Reference:
https://docs.microsoft.com/en-us/azure/architecture/solution-ideas/articles/iot-using-cosmos-db

NEW QUESTION: 74
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?
A. Yes
B. No
Answer: A (LEAVE A REPLY)
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:

NEW QUESTION: 75
You have a container named container! in an Azure Cosmos DB for NoSQL account.
You need to provide a user named User1 with the ability to insert items into container1 by using role-based access The solution must use the principle of least privilege.
Which roles should you assign to User1?
A. Cosmos DB Built-in Data Contributor only
B. Cosmos DB Operator only
C. DocumentDB Account Contribute only
D. DocumentDB Account Contributor and Cosmos DB Built-in Data Contributor
Answer: A (LEAVE A REPLY)
The Cosmos DB Built-in Data Contributor role provides the necessary permissions to insert items into a container in an Azure Cosmos DB for NoSQL account. This role grants the minimum required privileges for the described task, adhering to the principle of least privilege.

NEW QUESTION: 76
You have three containers in an Azure Cosmos DB Core (SQL) API account as shown in the following table.
You have the following Azure functions:
A function named Fn1 that reads the change feed of cn1
A function named Fn2 that reads the change feed of cn2
A function named Fn3 that reads the change feed of cn3
You perform the following actions:
Delete an item named item1 from cn1.
Update an item named item2 in cn2.
For an item named item3 in cn3, update the item time to live to 3,600 seconds.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.
Answer:
Explanation
Box 1: No
Azure Cosmos DB's change feed is a great choice as a central data store in event sourcing architectures where all data ingestion is modeled as writes (no updates or deletes).
Note: The change feed does not capture deletes. If you delete an item from your container, it is also removed from the change feed. The most common method of handling this is adding a soft marker on the items that are being deleted. You can add a property called "deleted" and set it to "true" at the time of deletion. This document update will show up in the change feed. You can set a TTL on this item so that it can be automatically deleted later.
Box 2: No
The _etag format is internal and you should not take dependency on it, because it can change anytime.
Box 3: Yes

Change feed support in Azure Cosmos DB works by listening to an Azure Cosmos container for any changes.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/sql/change-feed-design-patterns
https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed

NEW QUESTION: 77
You plan to create an operational system that will store data in an Azure Cosmos OB 'or NoSQL account. You need to configure the account to meet the following requirements:
* Support Spar* queries.
* Support the analysis of data from the last six months.
* 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 analytical property to six months.
C. Enable Azure Synapse Link for the account
D. Create a container and set the time to live to six months.
E. Create an Azure Databanks notebook.
F. Create an Azure Synapse pipeline.
Answer: (SHOW ANSWER)

NEW QUESTION: 78
You plan to create an Azure Cosmos DB database named db1 that will contain two containers. One of the containers will contain blog posts, and the other will contain users. Each item in the blog post container will include:
* A single blog post
* All the comments associated to the blog post
* The names of the users who created the blog post and added the comments.
You need to design a solution to update usernames m the user container without causing data integrity issues. The solution must minimize administrative and development effort. What should you include in the solution? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer:

NEW QUESTION: 79

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

You need to use the Azure Cosmos DB SDK to replace a document by using optimistic concurrency.

What should you include in the code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer:

Explanation

Box 1: ConsistencyLevel

The ItemRequestOptions Class ConsistencyLevel property gets or sets the consistency level required for the request in the Azure Cosmos DB service.

Azure Cosmos DB offers 5 different consistency levels. Strong, Bounded Staleness, Session, Consistent Prefix and Eventual - in order of strongest to weakest consistency.

Box 2: _etag

The ItemRequestOptions class helped us implement optimistic concurrency by specifying that we wanted the SDK to use the If-Match header to allow the server to decide whether a resource should be updated. The If-Match value is the ETag value to be checked against. If the ETag value matches the server ETag value, the resource is updated.

Reference:

https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.itemrequestoptions

https://cosmosdb.github.io/labs/dotnet/labs/10-concurrency-control.html

NEW QUESTION: 80

You have a container in an Azure Cosmos DB for NoSQL account. The database that has a manual throughput of 30,000 request units per second (RU/s). The current consumption details are shewn in the following chart.

Normalized RU Consumption (%) By PartitionKeyRangeID

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.

Answer:

Explanation

NEW QUESTION: 81

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.

What is the effect of running the command?

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 (LEAVE A REPLY)

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.

NEW QUESTION: 82

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

You need to provide a user named User1 with the ability to insert items into container1 by using role-based access The solution must use the principle of least privilege.

Which roles should you assign to User1?

A. Cosmos DB Built-in Data Contributor only

B. Cosmos DB Operator only

C. DocumentDB Account Contribute only

D. DocumentDB Account Contributor and Cosmos DB Built-in Data Contributor

Answer: A (LEAVE A REPLY)

Explanation

The Cosmos DB Built-in Data Contributor role provides the necessary permissions to insert items into a container in an Azure Cosmos DB for NoSQL account. This role grants the minimum required privileges for the described task, adhering to the principle of least privilege.

NEW QUESTION: 83

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 leaseCollectionNair.e
C. Configure the trigger for each function to use a different leaseCollectionName
D. Configure the trigger for each function to use the same leaseCollectionPrefix
Answer: (SHOW ANSWER)
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

NEW QUESTION: 84
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:
* Store companies and the users employed by them. Each company will have less than 1,000 users.
* Some users have data that is greater than 2 KB.
* Associate each user to only one company.
* Provide the ability to browse by company.
* Provide the ability to browse the users by company.
* Whenever a company or user profile is selected, show a details page for the company and all the related users.
* 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.
Answer: B (LEAVE A REPLY)
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.

NEW QUESTION: 85
You plan to use a multi-region Azure Cosmos DB for NoSQL account to store data for a new application suite. The suite contains the applications shown in the following table.
Each application should use the weakest consistency level possible.
Which consistency level should you configure for each application? To answer, select the appropriate options in the answer are a. NOTE: Each correct selection is worth one point.
Answer:

NEW QUESTION: 86
You have a container named container1 in an Azure Cosmos DB Core (SQL) API account.
The following is a sample of a document in container1.

{
"studentId": "631282",
"firstName": "James",
"lastName": "Smith",
"enrollmentYear": 1990,
"isActivelyEnrolled": true,
"address": {
"street": "",
"city": "",
"stateProvince": "",
"postal": "",
}
}
The container1 container has the following indexing policy.
{
"indexingMode": "consistent",
"includePaths": [
{
"path": "/*"
},
{
"path": "/address/city/?"
}
],
"excludePaths": [
{
"path": "/address/*"
},
{
"path": "/firstName/?"
}
]
}
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer:

Explanation

Box 1: Yes

"path": "/*" is in includePaths.

Include the root path to selectively exclude paths that don't need to be indexed. This is the recommended approach as it lets Azure Cosmos DB proactively index any new property that may be added to your model.
Box 2: No
"path": "/firstName/?" is in excludePaths.
Box 3: Yes
"path": "/address/city/?" is in includePaths
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/index-policy

NEW QUESTION: 87
You have a multi-region Azure Cosmos DB account named account1 that has a default consistency level of strong.
You have an app named App1 that is configured to request a consistency level of session.
How will the read and write operations of App1 be handled? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Answer:
Explanation
Box 1 = Write and replicate data to every region synchronously
This is because the write concern is mapped to the default consistency level configured on your Azure Cosmos DB account2, which is strong in this case. Strong consistency ensures that every write operation is synchronously committed to every region associated with your Azure Cosmos DB account1. The request level consistency level of session only applies to the read operations of App11.
Box 2: That has the lowest estimated latency to the client
This is because the read operations of App1 will use the session consistency level that is specified in the request options. Session consistency is a client-centric consistency model that guarantees monotonic reads, monotonic writes, and read-your-own-writes within a session. A session is scoped to a client connection or a stored procedure execution. Session consistency allows clients to read from any region that has the lowest latency to the client.

NEW QUESTION: 88
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?
A. Metrics
B. Log
C. Resource Health
D. Activity Log
Answer: B (LEAVE A REPLY)

NEW QUESTION: 89
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.
Answer:

NEW QUESTION: 90
You configure a backup for an Azure Cosmos DB for NoSQL 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.
Answer:
Explanation
Box 1 = The current backup policy provides protection for: 2 Hours Azure Cosmos DB automatically takes backups of your data at regular intervals. The backup interval and the retention period can be configured from the Azure portal. You can also choose between two backup modes: periodic backup mode and continuous backup mode. Periodic backup mode is the default mode for all existing accounts and it takes a full backup of your database every 4 hours by default. Continuous backup mode is a new mode that allows you to restore to any point of time within either 7 or 30 days1.
For your scenario, based on the exhibit, you have configured a backup for an Azure Cosmos DB for NoSQL account using the periodic backup mode with a backup interval of 1 hour and a retention period of 2 hours.
This means that Azure Cosmos DB will take a full backup of your database every hour and keep only the latest two backups. Therefore, the current backup policy provides protection for 2 hours.
Box 2: In case of emergency, you must (answer choice) to restore the backup = create a support ticket Azure Cosmos DB automatically takes backups of your data at regular intervals. You can configure the backup interval and the retention period from the Azure portal. You can also choose between two backup modes:
periodic backup mode and continuous backup mode. Periodic backup mode is the default mode for all existing accounts and it takes a full backup of your database every 4 hours by default. Continuous backup mode is a new mode that allows you to restore to any point of time within either 7 or 30 days1.
For your scenario, based on the exhibit, you have configured a backup for an Azure Cosmos DB for NoSQL account using the periodic backup mode with a backup interval of 1 hour and a retention period of 2 hours.

This means that Azure Cosmos DB will take a full backup of your database every hour and keep only the latest two backups. In case of emergency, you must create a support ticket to restore the backup. This is the answer to your question.
To restore data from a periodic backup, you need to create a support request with Azure Cosmos DB team and provide the following information:
The name of your Azure Cosmos DB account
The name of the database or container that you want to restore
The date and time (in UTC) that you want to restore from
The name of the target Azure Cosmos DB account where you want to restore the data The name of the target resource group where you want to restore the data The Azure Cosmos DB team will then initiate the restore process and notify you when it is completed2.

NEW QUESTION: 91
You have a container named contained in an Azure Cosmos DB for NoSQL account.
The following is a sample of a document in container1.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.
Answer:
Explanation

NEW QUESTION: 92
You have an Azure Cosmos DB for NoSQL account.
You plan 10 create a container named container1. The container1 container will store items that include two properties named nm and age The most commonly executed queries will query container1 for a specific name. The following is a sample of the query.
You need to define an opt-in Indexing policy for container1. The solution must meet the following requirements:
* Minimize the number of request units consumed by the queries.
* Ensure that the _etag property is excluded from indexing.
How should you define the indexing poky? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection Is worth one point.

Answer:
Explanation

NEW QUESTION: 93
You have an Azure Cosmos DB for NoSQL account that has multiple write regions.
You need to receive an alert when requests that target the database exceed the available request units per second (RU/s).
Which Azure Monitor signal should you use?
A. Region Removed
B. Document Quota
C. Metadata Requests
D. Data Usage
Answer: B (LEAVE A REPLY)
Explanation
Azure Monitor is a service that provides comprehensive monitoring for Azure resources, including Azure Cosmos DB. You can use Azure Monitor to collect, analyze, and alert on metrics and logs from your Azure Cosmos DB account. You can create alerts for Azure Cosmos DB using Azure Monitor based on the metrics, activity log events, or Log Analytics logs on your account1.
For your scenario, if you want to receive an alert when requests that target the database exceed the available request units per second (RU/s), you should use the Document Quota metric. This metric measures the percentage of RU/s consumed by your account or container. You can create an alert rule on this metric from the Azure portal by following these steps2:
* In the Azure portal, select the Azure Cosmos DB account you want to monitor.
* Under the Monitoring section of the sidebar, select Alerts, and then select New alert rule.
* In the Create alert rule pane, fill out the Scope section by selecting your subscription name and resource type (Azure Cosmos DB accounts).
* In the Condition section, select Add condition and choose Document Quota from the list of signals.
* In the Configure signal logic pane, specify the threshold value and operator for your alert condition. For example, you can choose Greater than or equal to 90 as the threshold value and operator to receive an alert when your RU/s consumption reaches 90% or more of your provisioned throughput.
* In the Alert rule details section, specify a name and description for your alert rule.
* In the Actions section, select Add action group and choose how you want to receive notifications for your alert. For example, you can choose Email/SMS/Push/Voice as an action type and enter your email address or phone number as a receiver.
* Review your alert rule settings and select Create alert rule to save it.

NEW QUESTION: 94
You have an Azure Cosmos DB Core (SQL) API account that is configured for multi-region writes. The account contains a database that has two containers named container1 and container2.

The following is a sample of a document in container1:

```
{
"customerId": 1234,
"firstName": "John",
"lastName": "Smith",
"policyYear": 2021
}
```

The following is a sample of a document in container2:

```
{
"gpsId": 1234,
"latitude": 38.8951,
"longitude": -77.0364
}
```

You need to configure conflict resolution to meet the following requirements:
For container1 you must resolve conflicts by using the highest value for policyYear.
For container2 you must resolve conflicts by accepting the distance closest to latitude: 40.730610 and longitude: -73.935242.
Administrative effort must be minimized to implement the solution.
What should you configure for each container? To answer, drag the appropriate configurations to the correct containers. Each configuration may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

Answer:

Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/conflict-resolution-policies
https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-manage-conflicts

NEW QUESTION: 95

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.

Answer:

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: 96
You need to identify which connectivity mode to use when implementing App2. The solution must support the planned changes and meet the business requirements.
Which connectivity mode should you identify?
A. Direct mode over HTTPS
B. Gateway mode (using HTTPS)
C. Direct mode over TCP
Answer: (SHOW ANSWER)
Explanation
Scenario: Develop an app named App2 that will run from the retail stores and query the data in account2.
App2 must be limited to a single DNS endpoint when accessing account2.
By using Azure Private Link, you can connect to an Azure Cosmos account via a private endpoint. The private endpoint is a set of private IP addresses in a subnet within your virtual network.
When you're using Private Link with an Azure Cosmos account through a direct mode connection, you can use only the TCP protocol. The HTTP protocol is not currently supported.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-configure-private-endpoints

NEW QUESTION: 97
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.
Answer:
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: 98
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 (LEAVE A REPLY)
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: 99
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 container. Populate the aggregates by using the change feed.
C. Configure the report to query orders by using a SQL query through a dedicated gateway.

D. Configure the report to query a new aggregate container. Populate the aggregates by using SQL queries that run daily.

Answer: (SHOW ANSWER)

You can facilitate aggregate data by using Change Feed and Azure Functions, and then use it for reporting.

NEW QUESTION: 100

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.

Answer:

NEW QUESTION: 101

You have a database in an Azure Cosmos DB for NoSQL account that is configured for multi-region writes.

You need to use the Azure Cosmos DB SDK to implement the conflict resolution policy for a container. The solution must ensure that any conflict sent to the conflict feed.

Solution: You set ConfilictResolutionMode to Custom. You Set ResolutionProcedures to a custom stored procedure. You configure the custom stored procedure to use the isTomstone parameter to resolve conflict.

Does this meet the goal?

A. Yes

B. No

Answer: (SHOW ANSWER)

The solution is incorrect because there is no "isTom" parameter in the Azure Cosmos DB SDK. The correct parameter is "isTombstone".

NEW QUESTION: 102

You have an Azure Cosmos DB Core (SQL) API account that uses a custom conflict resolution policy. The account has a registered merge procedure that throws a runtime exception.

The runtime exception prevents conflicts from being resolved.

You need to use an Azure function to resolve the conflicts.

What should you use?

A. a function that pulls items from the conflicts feed and is triggered by a timer trigger

B. a function that receives items pushed from the change feed and is triggered by an Azure Cosmos DB trigger

C. a function that pulls items from the change feed and is triggered by a timer trigger

D. a function that receives items pushed from the conflicts feed and is triggered by an Azure Cosmos DB trigger

Answer: D (LEAVE A REPLY)

Explanation

The Azure Cosmos DB Trigger uses the Azure Cosmos DB Change Feed to listen for inserts and updates across partitions. The change feed publishes inserts and updates, not deletions.

Reference: https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-cosmosdb

NEW QUESTION: 103

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.

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.

Answer:

NEW QUESTION: 104

You have an Azure Cosmos DB for NoSQL container. The container contains items that have the following properties.

You need to protect the data stored in the container by using Always Encrypted. For each property, you must use the strongest type of encryption and ensure that queries execute properly.

What is the strongest type of encryption that you can apply to each property? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer:

NEW QUESTION: 105

You have an Azure Synapse Analytics workspace named workspace1 that contains a server less SQL pool.

You have an Azure Table Storage account that stores operational data.

You need to replace the Table storage account with Azure Cosmos DB for NoSQL The solution must meet the following requirements:

* Support Queries from the server less SQL pool.
* Only pay for analytical compute when running queries.
* Ensure that analytical processes do

NOTE: affect operational processes.

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.

Answer:

Explanation

NEW QUESTION: 106
You have an Azure Cosmos DB for NoSQL account.
The change feed is enabled on a container named invoice.
You create an Azure function that has a trigger on the change feed.
What is received by the Azure function?
A. all the properties of the updated items
B. only the partition key and the changed properties of the updated items
C. all the properties of the original items and the updated items
D. only the changed properties and the system-defined properties of the updated items
Answer: D (LEAVE A REPLY)
According to the Azure Cosmos DB documentation12, the change feed is a persistent record of changes to a container in the order they occur. The change feed outputs the sorted list of documents that were changed in the order in which they were modified.
The Azure function that has a trigger on the change feed receives all the properties of the updated items2. The change feed does not include the original items or only the changed properties. The change feed also includes some system-defined properties such as _ts (the last modified timestamp) and _lsn (the logical sequence number)3.
Therefore, the correct answer is:

NEW QUESTION: 107
You have a database in an Azure Cosmos DB for NoSQL account that is configured for multi-region writes.
You need to use the Azure Cosmos DB SDK to implement the conflict resolution policy for a container. The solution must ensure that any conflict sent to the conflict feed.
Solution: You set ConfilictResolutionMode to Custom and you use the default settings for the policy.
Does this meet the goal?
A. Yes
B. No
Answer: A (LEAVE A REPLY)

Setting ConflictResolutionMode to Custom and using the default settings for the policy will not ensure that conflicts are sent to the conflict feed. You need to define a custom stored procedure using the "conflictingItems" parameter to handle conflicts properly.

NEW QUESTION: 108
You have an Azure Cosmos DB Core (SQL) account that has a single write region in West Europe.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.
Answer:
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: 109
You have an Azure Cosmos DB for NoSQL account named account that has the disablekeyBasedletadatwriteAccess property enabled.
You are developing an app named App1 that will be used by a use1 named DevUser1 to create containers in account1. DevUser1 has a non-privileged user account in the 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.
Answer:
Explanation

NEW QUESTION: 110
You have an Azure Cosmos DB for NoSQL account that frequently receives the same three queries.
You need to configure indexing to minimize RUs consumed by the queries.
Which type of index should you use for each query? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Answer:

NEW QUESTION: 111
You have a database named telemetry in an Azure Cosmos DB Core (SQL) API account that stores IoT dat a. 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.
Answer:

NEW QUESTION: 112
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.
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.
Answer:
Explanation
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

NEW QUESTION: 113
You have a container named container1 in an Azure Cosmos DB Core (SQL) API account.
The following is a sample of a document in container1.

```
{
"studentId": "631282",
"firstName": "James",
"lastName": "Smith",
"enrollmentYear": 1990,
"isActivelyEnrolled": true,
"address": {
"street": "",
"city": "",
"stateProvince": "",
"postal": "",
}
}
```

The container1 container has the following indexing policy.

```
{
"indexingMode": "consistent",
"includePaths": [
{
"path": "/*"
},
{
"path": "/address/city/?"
}
],
"excludePaths": [
{
"path": "/address/*"
},
{
"path": "/firstName/?"
}
]
}
```

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

NEW QUESTION: 114
You have three containers in an Azure Cosmos DB Core (SQL) API account as shown in the following table.
You have the following Azure functions:
A function named Fn1 that reads the change feed of cn1
A function named Fn2 that reads the change feed of cn2
A function named Fn3 that reads the change feed of cn3
You perform the following actions:
Delete an item named item1 from cn1.
Update an item named item2 in cn2.
For an item named item3 in cn3, update the item time to live to 3,600 seconds.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.
Answer:
Explanation
Box 1: No
Azure Cosmos DB's change feed is a great choice as a central data store in event sourcing architectures where all data ingestion is modeled as writes (no updates or deletes).
Note: The change feed does not capture deletes. If you delete an item from your container, it is also removed from the change feed. The most common method of handling this is adding a soft marker on the items that are being deleted. You can add a property called "deleted" and set it to "true" at the time of deletion. This document update will show up in the change feed. You can set a TTL on this item so that it can be automatically deleted later.
Box 2: No
The _etag format is internal and you should not take dependency on it, because it can change anytime.
Box 3: Yes
Change feed support in Azure Cosmos DB works by listening to an Azure Cosmos container for any changes.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/sql/change-feed-design-patterns
https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed

NEW QUESTION: 115
You have a database named db1 in an Azure Cosmos DB for NoSQL account named account1.
The db1 database has a manual throughput of 4,000 request units per second (RU/s).
You need to move db1 from manual throughput to autoscale throughput by using the Azure CLI.
The solution must provide a minimum of 4,000 RU/s and a maximum of 40,000 RU/s.
How should you complete the CLI statements? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer:

Explanation

Migrate

40000

According to the Azure CLI reference , you need to use the az cosmosdb sql database throughput migrate command to migrate the throughput of the SQL database between autoscale and manually provisioned. You also need to use the --throughput-type parameter to specify the type of throughput to migrate to, and the --max-throughput parameter to specify the maximum throughput resource can scale to (RU/s).

To complete the CLI statements, you should replace the missing values with:

--throughput-type autoscale

--max-throughput 40000

The final command should look like this:

```
az cosmosdb sql database throughput migrate \
--account-name account1 \
--name db1 \
--resource-group rg1 \
--throughput-type autoscale \
--max-throughput 40000
```

NEW QUESTION: 116

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.

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.

Answer:

NEW QUESTION: 117

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 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
Answer: B (LEAVE A REPLY)

NEW QUESTION: 118
You have a database in an Azure Cosmos DB Core (SQL) API account.
You plan to create a container that will store employee data for 5,000 small businesses. Each business will have up to 25 employees. Each employee item will have an emailAddress value.
You need to ensure that the emailAddress value for each employee within the same company is unique.
To what should you set the partition key and the unique key? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Answer:

NEW QUESTION: 119
You have a multi-region Azure Cosmos DB account named account1 that has a default consistency level of strong.
You have an app named App1 that is configured to request a consistency level of session.
How will the read and write operations of App1 be handled? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Answer:

NEW QUESTION: 120
You need to configure an Apache Kafka instance to ingest data from an Azure Cosmos DB Core (SQL) API account. The data from a container named telemetry must be added to a Kafka topic named iot. The solution must store the data in a compact binary format.
Which three configuration items should you include in the solution? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.
A. "connector.class": "com.azure.cosmos.kafka.connect.source.CosmosDBSourceConnector"
B. "key.converter": "org.apache.kafka.connect.json.JsonConverter"
C. "key.converter": "io.confluent.connect.avro.AvroConverter"
D. "connect.cosmos.containers.topicmap": "iot#telemetry"
E. "connect.cosmos.containers.topicmap": "iot"
F. "connector.class": "com.azure.cosmos.kafka.connect.source.CosmosDBSinkConnector"
Answer: (SHOW ANSWER)
Explanation
C: Avro is binary format, while JSON is text.

F: Kafka Connect for Azure Cosmos DB is a connector to read from and write data to Azure Cosmos DB. The Azure Cosmos DB sink connector allows you to export data from Apache Kafka topics to an Azure Cosmos DB database. The connector polls data from Kafka to write to containers in the database based on the topics subscription.
D: Create the Azure Cosmos DB sink connector in Kafka Connect. The following JSON body defines config for the sink connector.
Extract:
"connector.class": "com.azure.cosmos.kafka.connect.sink.CosmosDBSinkConnector",
"key.converter": "org.apache.kafka.connect.json.AvroConverter"
"connect.cosmos.containers.topicmap": "hotels#kafka"
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/sql/kafka-connector-sink
https://www.confluent.io/blog/kafka-connect-deep-dive-converters-serialization-explained/

NEW QUESTION: 121
You have an Azure Cosmos DB Core (SQL) account that has a single write region in West Europe.
database named db
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.
Answer:
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: 122
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.
Answer:
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: 123
You have the following Azure Resource Manager (ARM) template.
You plan to deploy the template in incremental mode.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.

Answer:
Explanation

NEW QUESTION: 124
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.
Answer: C,E (LEAVE A REPLY)
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.
Reference: https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-use-stored-procedures-triggers-udfs