

AZ-204 Dumps

Developing Solutions for Microsoft Azure (beta)

<https://www.certleader.com/AZ-204-dumps.html>



NEW QUESTION 1

- (Exam Topic 1)

You need to migrate on-premises shipping data to Azure. What should you use?

- A. Azure Migrate
- B. Azure Cosmos DB Data Migration tool (dt.exe)
- C. AzCopy
- D. Azure Database Migration service

Answer: D

Explanation:

Migrate from on-premises or cloud implementations of MongoDB to Azure Cosmos DB with minimal downtime by using Azure Database Migration Service. Perform resilient migrations of MongoDB data at scale and with high reliability.

Scenario: Data migration from on-premises to Azure must minimize costs and downtime.

The application uses MongoDB JSON document storage database for all container and transport information. References:

<https://azure.microsoft.com/en-us/updates/mongodb-to-azure-cosmos-db-online-and-offline-migrations-are-now>

NEW QUESTION 2

- (Exam Topic 3)

You use Azure Table storage to store customer information for an application. The data contains customer details and is partitioned by last name. You need to create a query that returns all customers with the last name Smith. Which code segment should you use?

- A. `TableQuery.GenerateFilterCondition("PartitionKey", Equals, "Smith")`
- B. `TableQuery.GenerateFilterCondition("LastName", Equals, "Smith")`
- C. `TableQuery.GenerateFilterCondition("PartitionKey", QueryComparisons.Equal, "Smith")`
- D. `TableQuery.GenerateFilterCondition("LastName", QueryComparisons.Equal, "Smith")`

Answer: C

Explanation:

Retrieve all entities in a partition. The following code example specifies a filter for entities where 'Smith' is the partition key. This example prints the fields of each entity in the query results to the console.

Construct the query operation for all customer entities where `PartitionKey="Smith"`. `TableQuery<CustomerEntity> query = new`

`TableQuery<CustomerEntity>().Where(TableQuery.GenerateFilterCondition("PartitionKey",`

`QueryComparisons.Equal, "Smith"));`

References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-dotnet>

NEW QUESTION 3

- (Exam Topic 3)

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 question, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are developing a website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.

You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level. You need to configure authorization.

Solution: Configure the Azure Web App for the website to allow only authenticated requests and require Azure AD log on.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Instead in the Azure AD application's manifest, set value of the `groupMembershipClaims` option to All. References:

<https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/>

NEW QUESTION 4

- (Exam Topic 3)

You must ensure that the external party cannot access the data in the SSN column of the Person table.

Will each protection method meet the requirement? To answer, drag the appropriate responses to the correct protection methods. Each response 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.

| Responses | Protection method | Response |
|------------------------------|--|--------------------------|
| <input type="checkbox"/> Yes | Enable AlwaysOn encryption. | <input type="checkbox"/> |
| <input type="checkbox"/> No | Set the column encryption setting to disabled. | <input type="checkbox"/> |
| | Assign users to the Public fixed database role. | <input type="checkbox"/> |
| | Store column encryption keys in the system catalog view in the database. | <input type="checkbox"/> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

You can configure Always Encrypted for individual database columns containing your sensitive data. When setting up encryption for a column, you specify the information about the encryption algorithm and cryptographic keys used to protect the data in the column.

Box 2: No

Box 3: Yes

In SQL Database, the VIEW permissions are not granted by default to the public fixed database role. This enables certain existing, legacy tools (using older versions of DacFx) to work properly. Consequently, to work with encrypted columns (even if not decrypting them) a database administrator must explicitly grant the two VIEW permissions.

Box 4: No

All cryptographic keys are stored in an Azure Key Vault. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

NEW QUESTION 5

- (Exam Topic 3)

A company is developing a Java web app. The web app code is hosted in a GitHub repository located at <https://github.com/Contoso/webapp>.

The web app must be evaluated before it is moved to production. You must deploy the initial code release to a deployment slot named staging.

You need to create the web app and deploy the code.

How should you complete the commands? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

```
gitrepo=https://github.com/Contoso/webapp
webappname=businesswebapp
resourcegroupname=BusinessAppResourceGroup
```

az ▼

- group
- webapp
- appservice plan
- webapp deployment slot
- webapp deployment source

```
create --location centralus - --name $resourcegroupname
create --name $webappname - --resource-group $resourcegroupname
- --sku S3
create --name $webappname - --resource-group $resourcegroupname
\ - --plan $webappname
create --name $webappname - --resource-group $resourcegroupname
\ - --slot staging
```

az ▼

- group
- webapp
- appservice plan
- webapp deployment slot
- webapp deployment source

```
config - --name $webappname - --resource-group $resourcegroupname
\ - --slot staging - --repo-url
$gitrepo - --branch master - --manual-integration
```

az ▼

- group
- webapp
- appservice plan
- webapp deployment slot
- webapp deployment source

az ▼

- group
- webapp
- appservice plan
- webapp deployment slot
- webapp deployment source

az ▼

- group
- webapp
- appservice plan
- webapp deployment slot
- webapp deployment source

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: group

Create a resource group.

```
az group create --location westeurope --name myResourceGroup
```

Box 2: appservice plan

```
# Create an App Service plan in STANDARD tier (minimum required by deployment slots). az appservice plan create --name $webappname --resource-group myResourceGroup --sku S1
```

Box 3: webapp

Create a web app.

```
az webapp create --name $webappname --resource-group myResourceGroup \ --plan $webappname
```

Box 4: webapp deployment slot

#Create a deployment slot with the name "staging".

```
az webapp deployment slot create --name $webappname --resource-group myResourceGroup \ --slot staging
```

Box 5: webapp deployment source

Deploy sample code to "staging" slot from GitHub.

```
az webapp deployment source config --name $webappname --resource-group myResourceGroup \ --slot staging --repo-url $gitrepo --branch master --manual-integration
```

References:
<https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-deploy-staging-environment>

NEW QUESTION 6

- (Exam Topic 3)

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 question, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

Margie's Travel is an international travel and bookings management service. The company is expanding into restaurant bookings. You are tasked with implementing Azure Search for the restaurants listed in their solution.

You create the index in Azure Search.

You need to import the restaurant data into the Azure Search service by using the Azure Search .NET SDK. Solution:

- * 1. Create a SearchIndexClient object to connect to the search index.
- * 2. Create a DataContainer that contains the documents which must be added.
- * 3. Create a DataSource instance and set its Container property to the DataContainer.
- * 4. Call the Documents.Suggest method of the SearchIndexClient and pass the DataSource. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Use the following method:

- * 1. - Create a SearchIndexClient object to connect to the search index
- * 2. - Create an IndexBatch that contains the documents which must be added.
- * 3. - Call the Documents.Index method of the SearchIndexClient and pass the IndexBatch. References:
<https://docs.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk>

NEW QUESTION 7

- (Exam Topic 3)

You develop a web app that uses tier D1 app service plan by using the Web Apps feature of Microsoft Azure App Service.

Spikes in traffic have caused increases in page load times.

You need to ensure that the web app automatically scales when CPU load is about 85 percent and minimize costs.

Which four 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.

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Actions

Answer Area

Configure the web app to the Premium App Service tier.

Configure the web app to the Standard App Service tier.

Enable autoscaling on the web-app.

Add a Scale rule.

Switch to an Azure App Services consumption plan.

Configure a Scale condition.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Configure the web app to the Standard App Service Tier
The Standard tier supports auto-scaling, and we should minimize the cost. Step 2: Enable autoscaling on the web app
First enable autoscale Step 3: Add a scale rule
Step 4: Add a Scale condition Reference:
<https://docs.microsoft.com/en-us/azure/monitoring-and-diagnostics/monitoring-autoscale-get-started>

NEW QUESTION 8

- (Exam Topic 3)

You are developing a data storage solution for a social networking app. The solution requires a mobile app that stores user information using Azure Table Storage. You need to develop code that can insert multiple sets of user information. How should you complete the code? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

```
CloudStorageAccount storageAccount = CloudStorageAccount.Parse(
    CloudConfigurationManager.GetSetting("StorageConnectionString"));
CloudTableClient tableClient = storageAccount.CreateCloudTableClient();
CloudTable table = tableClient.GetTableReference("clients");
Table.CreateIfNotExists();
```

op = new ();

| |
|--------------------|
| ▼ |
| TableOperation |
| TableBatchOperaton |
| TableEntity |
| TableQuery |

| |
|--------------------|
| ▼ |
| TableOperation |
| TableBatchOperaton |
| TableEntity |
| TableQuery |

table. (op);

| |
|---------------|
| ▼ |
| ExecuteBatch |
| Execute |
| Insert |
| InsertOrMerge |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1, Box 2: TableBatchOperation Create the batch operation.

TableBatchOperation op = new TableBatchOperation(); Box 3: ExecuteBatch

/ Execute the batch operation. table.ExecuteBatch(op);

Note: You can insert a batch of entities into a table in one write operation. Some other notes on batch operations:

You can perform updates, deletes, and inserts in the same single batch operation. A single batch operation can include up to 100 entities.

All entities in a single batch operation must have the same partition key.

While it is possible to perform a query as a batch operation, it must be the only operation in the batch. References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-dotnet>

NEW QUESTION 9

- (Exam Topic 3)

Your company is migrating applications to Azure. The IT department must allow internal developers to communicate with Microsoft support.

The service agents of the IT department must only have view resources and create support ticket permissions to all subscriptions. A new custom role must be created by reusing a default role definition and changing the permissions.

You need to create the custom role.

To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

| Item | Value |
|--------------------|---|
| Powershell command | <div style="border: 1px solid black; padding: 5px;"> <pre>Get-AzureRmRoleDefinition-Name"Reader" ConvertTo-Json Out-File C:\SupportRole.json Get-AzureRmRoleDefinition-Name"Operator" ConvertTo-Json Out-File C:\SupportRole.json Set-AzureRmRoleDefinition-Name"Reader" Input-File C:\SupportRole.json Set-AzureRmRoleDefinition Input-File C:\SupportRole.json</pre> </div> |
| Actions section | <div style="border: 1px solid black; padding: 5px;"> <pre>"*/read*", "Microsoft.Support/*" "*/read*" "*/read*", "Microsoft.Support/*" "*/read*"</pre> </div> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Set-AzureRmRoleDefinition Input-File C:\SupportRole.json

The Set-AzureRmRoleDefinition cmdlet updates an existing custom role in Azure Role-Based Access Control. Provide the updated role definition as an input to the command as a JSON file or a PSRoleDefinition object.

The role definition for the updated custom role MUST contain the Id and all other required properties of the role even if they are not updated: DisplayName, Description, Actions, AssignableScope


Box 2: "*/read*", "Microsoft.Support/*" Microsoft.Support/* Create and manage support tickets "Microsoft.Support" role definition azure

NEW QUESTION 10

- (Exam Topic 3)


You are validating the configuration of an Azure Search indexer.

The service has been configured with an indexer that uses the Import Data option. The index is configured using options as shown in the Index Configuration exhibit. (Click the Index Configuration tab.)



You use an Azure table as the data source for the import operation. The table contains three records with item inventory data that matches the fields in the Storage data exhibit. These records were imported when the index was created. (Click the Storage Data tab.) When users search with no filter, all three records are displayed.

| PartitionKey | RowKey | Timestamp | InStockCount | ItemDescription | ItemName | LocationRow | LocationShelf | SKU |
|--------------|--------|---------------------------|--------------|-------------------------------|-----------|-------------|---------------|--------|
| Food | 3 | 2018-08-25T 15:47:29.135Z | 32 | A box of chocolate candy bars | Choco-bar | 5 | 3 | 123421 |
| Hardware | 2 | 2018-08-25T 15:46:08.402Z | 2 | A bag of bolts | Bolts | 1 | 4 | 67894 |
| Hardware | 1 | 2018-08-25T 15:46:41.402Z | 23 | A box of nails | Nails | 2 | 1 | 654365 |



When users search for items by description, Search explorer returns no records. The Search Explorer exhibit shows the query and results for a test. In the test, a user is trying to search for all items in the table that have a description that contains the word bag. (Click the Search Explorer tab.)

You need to resolve the issue.

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

- | | Yes | No |
|--|-----------------------|-----------------------|
| You can resolve the issue by recreating the search index with the same settings for all fields except ItemDescription. Select the SEARCHABLE option for this field | <input type="radio"/> | <input type="radio"/> |
| You can resolve the issue by selecting the index, editing the ItemDescription field, and selecting the SEARCHABLE option for the field. | <input type="radio"/> | <input type="radio"/> |
| You can resolve the issue by running the indexer. | <input type="radio"/> | <input type="radio"/> |
| You can resolve the issue by changing the query string in Search explorer to <code>bag of</code> to return the correct results | <input type="radio"/> | <input type="radio"/> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

The ItemDescription field in not searchable. Box 2: No

The ItemDescription field in not searchable, but we would need to recreate the index. Box 3: Yes

An indexer in Azure Search is a crawler that extracts searchable data and metadata from an external Azure data source and populates an index based on field-to-field mappings between the index and your data source. This approach is sometimes referred to as a 'pull model' because the service pulls data in without you having to write any code that adds data to an index.

Box 4: No References:

<https://docs.microsoft.com/en-us/azure/search/search-what-is-an-index> <https://docs.microsoft.com/en-us/azure/search/search-indexer-overview>

NEW QUESTION 10

- (Exam Topic 3)

You are a developer for a software as a service (SaaS) company that uses an Azure Function to process orders. The Azure Function currently runs on an Azure Function app that is triggered by an Azure Storage queue.

You are preparing to migrate the Azure Function to Kubernetes using Kubernetes-based Event Driven Autoscaling (KEDA).

You need to configure Kubernetes Custom Resource Definitions (CRD) for the Azure Function.

Which CRDs should you configure? To answer, drag the appropriate CRD types to the correct locations. Each CRD type 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 Area

| CRD types | Setting | CRD type |
|-----------------------|---------------------------------|----------------------|
| Secret | Azure Function code | <input type="text"/> |
| Deployment | Polling interval | <input type="text"/> |
| ScaledObject | Azure Storage connection string | <input type="text"/> |
| TriggerAuthentication | | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Deployment

To deploy Azure Functions to Kubernetes use the func kubernetes deploy command has several attributes that directly control how our app scales, once it is deployed to Kubernetes.

Box 2: ScaledObject

With --polling-interval, we can control the interval used by KEDA to check Azure Service Bus Queue for messages.

Example of ScaledObject with polling interval apiVersion: keda.k8s.io/v1alpha1

kind: ScaledObject metadata:

name: transformer-fn namespace: tt

labels:

deploymentName: transformer-fn spec:

scaleTargetRef: deploymentName: transformer-fn pollingInterval: 5

minReplicaCount: 0

maxReplicaCount: 100

Box 3: Secret

Store connection strings in Kubernetes Secrets. Example: to create the Secret in our demo Namespace:

create the k8s demo namespace kubectl create namespace tt

grab connection string from Azure Service Bus KEDA_SCALER_CONNECTION_STRING=\$(az servicebus queue authorization-rule keys list \

-g \$RG_NAME \

--namespace-name \$SBN_NAME \

--queue-name inbound \

-n keda-scaler \

--query "primaryConnectionString" \

-o tsv)

create the kubernetes secret

kubectl create secret generic tt-keda-auth \

--from-literal KedaScaler=\$KEDA_SCALER_CONNECTION_STRING \

--namespace tt Reference:

<https://www.thinktecture.com/en/kubernetes/serverless-workloads-with-keda/>

NEW QUESTION 14

- (Exam Topic 3)

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 develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure Storage Blob storage. The storage account type is

General-purpose V2.

When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.

You need to design the process that starts the photo processing.

Solution: Move photo processing to an Azure Function triggered from the blob upload. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

Azure Storage events allow applications to react to events. Common Blob storage event scenarios include image or video processing, search indexing, or any file-oriented workflow.

Events are pushed using Azure Event Grid to subscribers such as Azure Functions, Azure Logic Apps, or even to your own http listener.

Note: Only storage accounts of kind StorageV2 (general purpose v2) and BlobStorage support event integration. Storage (general purpose v1) does not support integration with Event Grid.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview>

NEW QUESTION 16

- (Exam Topic 3)

You have a web service that is used to pay for food deliveries. The web service uses Azure Cosmos DB as the data store.

You plan to add a new feature that allows users to set a tip amount. The new feature requires that a property named tip on the document in Cosmos DB must be present and contain a numeric value.

There are many existing websites and mobile apps that use the web service that will not be updated to set the tip property for some time.

How should you complete the trigger?

NOTE: Each correct selection is worth one point.

```
function ensureTip() {
  var r = 
  
  var i = r.getBody();
  
  
  
}
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: getContext().getRequest(); Box 2: if(isNaN(i) ["tip"] ..

In JavaScript, there are two ways to check if a variable is a number :

isNaN() – Stands for “is Not a Number”, if variable is not a number, it return true, else return false. typeof – If variable is a number, it will returns a string named “number”.

Box 3:r.setBody(i);

// update the item that will be created References:

<https://docs.microsoft.com/bs-latn-ba/azure/cosmos-db/how-to-write-stored-procedures-triggers-udfs>

<https://mkyong.com/javascript/check-if-variable-is-a-number-in-javascript/>

NEW QUESTION 19

- (Exam Topic 3)

You have an app that stores player scores for an online game. The app stores data in Azure tables using a class named PlayerScore as the table entity. The table is populated with 100,000 records.

You are reviewing the following section of code that is intended to retrieve 20 records where the player score exceeds 15,000. (Line numbers are included for reference only.)

```

1 public void GetScore(string playerId, int score, string gameName)
2 {
3     TableQuery<DynamicTableEntity> query = new TableQuery<DynamicTableEntity>().Select(new string[] { "Score" })
        .Where(TableQuery.GenerateFilterConditionForInt("Score", QueryComparisons.GreaterThanOrEqual, 15000)).Take
(20);
4     EntityResolver<KeyValuePair<string, int?>> resolver =
        (partitionKey, rowKey, ts, props, etag) => new KeyValuePair<string, int?>(rowKey, props["Score"].Int32Value);
5     foreach (var scoreItem in scoreTable.ExecuteQuery(query, resolver, null, null))
6     {
7         Console.WriteLine($"{scoreItem.Key} {scoreItem.Value}");
8     }

9     public class PlayerScore : TableEntity
10 {
11     public PlayerScore(string gameId, string playerId, int score, long timePlayed)
12     {
13         PartitionKey = gameId;
14         RowKey = playerId;
15         Score = score;
16         TimePlayed = timePlayed;
17     }
18     public int Score { get; set; }
19     public long TimePlayed { get; set; }
20 }

```

You have the following code. (Line numbers are included for reference only.)

You store customer information in an Azure Cosmos database. The following data already exists in the database:

```

01 CloudTableClient tableClient = account.CreateCloudTableClient();
02 CloudTable table = tableClient.GetTableReference("people");
03 TableQuery<CustomerEntity> query = new TableQuery<CustomerEntity>()
04     .Where(TableQuery.CombineFilters(
05     TableQuery.Generate.And, TableQuery.GenerateFilterCondition(Email, QueryComparisons.Equal, "Smith")
06     TableOperators.And, TableQuery.GenerateFilterCondition(Email, QueryComparisons.Equal,
"ssmith@contoso.com")
07 ));
08 await table.ExecuteQuerySegmentedAsync<CustomerEntity>(query, null);

```

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

| | Yes | No |
|---|-----------------------|-----------------------|
| The code queries the Azure table and retrieves the TimePlayed property from the table | <input type="radio"/> | <input type="radio"/> |
| The code will display a maximum of twenty records. | <input type="radio"/> | <input type="radio"/> |
| All records will be sent to the client. The client will display records for scores greater than or equal to 15,000. | <input type="radio"/> | <input type="radio"/> |
| The scoreItem.Key property of the KeyValuePair that ExecuteQuery returns will contain a value for PlayerID. | <input type="radio"/> | <input type="radio"/> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: No

Box 2: Yes

The TableQuery.Take method defines the upper bound for the number of entities the query returns. Example:

query.Take(10);

Box 3: Yes

Box 4: Yes

References:
<https://www.vkinfotek.com/azureqa/how-do-i-query-azure-table-storage-using-tablequery-class.html>

NEW QUESTION 22

- (Exam Topic 3)

A company develops a series of mobile games. All games use a single leaderboard service. You have the following requirements:

- Code should be scalable and allow for growth.
- Each record must consist of a playerId, gameId, score, and time played.

- When users reach a new high score, the system will save the new score using the SaveScore function below.
- Each game is assigned and Id based on the series title.

You have the following code. (Line numbers are included for reference only.)

You store customer information in an Azure Cosmos database. The following data already exists in the database:

```
01 CloudTableClient tableClient = account.CreateCloudTableClient();
02 CloudTable table = tableClient.GetTableReference("people");
03 TableQuery<CustomerEntity> query = new TableQuery<CustomerEntity>()
04     .Where(TableQuery.CombineFilters(
05         TableQuery.Generate.And, TableQuery.GenerateFilterCondition(Email, QueryComparisons.Equal, "Smith")
06         TableOperstors.And, TableQuery.GenerateFilterCondition(Email, QueryComparisons.Equal,
07         "ssmith@contoso.com")
08     ));
08 await table.ExecuteQuerySegmentedAsync<CustomerEntity>(query, null);
```

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

| | Yes | No |
|---|-----------------------|-----------------------|
| The code will work with Cosmos DB. | <input type="radio"/> | <input type="radio"/> |
| The save score function will update and replace a record if one already exists with the same playerId and gameId. | <input type="radio"/> | <input type="radio"/> |
| The data for the game will be automatically partitioned. | <input type="radio"/> | <input type="radio"/> |
| This code will store the values for the gameId and playerId parameters in the database. | <input type="radio"/> | <input type="radio"/> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

Code for CosmosDB, example:

```
// Parse the connection string and return a reference to the storage account. CloudStorageAccount storageAccount = CloudStorageAccount.Parse(
CloudConfigurationManager.GetSetting("StorageConnectionString"));
```

```
// Create the table client.
```

```
CloudTableClient tableClient = storageAccount.CreateCloudTableClient();
```

```
// Retrieve a reference to the table.
```

```
CloudTable table = tableClient.GetTableReference("people");
```

```
// Create the TableOperation object that inserts the customer entity. TableOperation insertOperation = TableOperation.Insert(customer1);
```

Box 2: No
A new record will always be added as TableOperation.Insert is used, instead of TableOperation.InsertOrReplace.

Box 3: No

No partition key is used. Box 4: Yes

References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-dotnet>

NEW QUESTION 23

- (Exam Topic 3)

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 develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure Storage Blob storage. The storage account type is General-purpose V2.

When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.

You need to design the process that starts the photo processing.

Solution: Convert the Azure Storage account to a BlockBlobStorage storage account. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Not necessary to convert the account, instead move photo processing to an Azure Function triggered from the blob upload..

Azure Storage events allow applications to react to events. Common Blob storage event scenarios include image or video processing, search indexing, or any file-oriented workflow.

Note: Only storage accounts of kind StorageV2 (general purpose v2) and BlobStorage support event integration. Storage (general purpose v1) does not support integration with Event Grid.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview>

NEW QUESTION 28

- (Exam Topic 3)

You are using Azure Front Door Service.

You are expecting inbound files to be compressed by using Brotli compression. You discover that inbound XML files are not compressed. The files are 9 megabytes (MB) in size.

You need to determine the root cause for the issue.

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

NOTE: Each correct selection is worth one point.

Answer Area

| Statement | Yes | No |
|---|-----------------------|-----------------------|
| The file MIME type is supported by the service. | <input type="radio"/> | <input type="radio"/> |
| Edge nodes must be purged of all cache assets. | <input type="radio"/> | <input type="radio"/> |
| The compression type is supported. | <input type="radio"/> | <input type="radio"/> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: No

Front Door can dynamically compress content on the edge, resulting in a smaller and faster response to your clients. All files are eligible for compression. However, a file must be of a MIME type that is eligible for compression list.

Box 2: No

Sometimes you may wish to purge cached content from all edge nodes and force them all to retrieve new updated assets. This might be due to updates to your web application, or to quickly update assets that contain incorrect information.

Box 3: Yes

These profiles support the following compression encodings: Gzip (GNU zip), Brotli Reference:

<https://docs.microsoft.com/en-us/azure/frontdoor/front-door-caching>

NEW QUESTION 32

- (Exam Topic 3)

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.

Margie's Travel is an international travel and bookings management service. The company is expanding into restaurant bookings. You are tasked with implementing Azure Search for the restaurants listed in their solution.

You create the index in Azure Search.

You need to import the restaurant data into the Azure Search service by using the Azure Search .NET SDK. Solution:

- * 1. Create a SearchIndexClient object to connect to the search index.
- * 2. Create a DataContainer that contains the documents which must be added.
- * 3. Create a DataSource instance and set its Container property to the DataContainer
- * 4 Call the Documents.Suggest method of the SearchIndexClient and pass the DataSource.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 36

- (Exam Topic 3)

You are developing a web app that is protected by Azure Web Application Firewall (WAF). All traffic to the web app is routed through an Azure Application Gateway instance that is used by multiple web apps. The web app address is contoso.azurewebsites.net.

All traffic must be secured with SSL. The Azure Application Gateway instance is used by multiple web apps. You need to configure the Azure Application Gateway for the app.

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. In the Azure Application Gateway's HTTP setting, enable the Use for App service setting.
- B. Convert the web app to run in an Azure App service environment (ASE).
- C. Add an authentication certificate for contoso.azurewebsites.net to the Azure Application gateway.
- D. In the Azure Application Gateway's HTTP setting, set the value of the Override backend path option to contoso22.azurewebsites.net.

Answer: AD

Explanation:

D: The ability to specify a host override is defined in the HTTP settings and can be applied to any back-end pool during rule creation. The ability to derive the host name from the IP or FQDN of the back-end pool members. HTTP settings also provide an option to dynamically pick the host name from a back-end pool member's FQDN if configured with the option to derive host name from an individual back-end pool member.
A (not C): SSL termination and end to end SSL with multi-tenant services.
In case of end to end SSL, trusted Azure services such as Azure App service web apps do not require whitelisting the backends in the application gateway. Therefore, there is no need to add any authentication certificates.

Add HTTP setting
saiappgw-appgw

* Protocol

Authentication certificates are not required for trusted Azure certificates for end to end ssl to work

* Port ⓘ

* Request timeout (seconds)

Override backend path ⓘ

Use for App service

Use custom probe ⓘ

Reference:

<https://docs.microsoft.com/en-us/azure/application-gateway/application-gateway-web-app-overview>

NEW QUESTION 41

- (Exam Topic 3)

You provide an Azure API Management managed web service to clients. The back end web service implements HTTP Strict Transport Security (HSTS). Every request to the backend service must include a valid HTTP authorization header. You need to configure the Azure API Management instance with an authentication policy. Which two policies can you use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. Certificate Authentication
- B. Basic Authentication
- C. OAuth Client Credential Grant
- D. Digest Authentication

Answer: AC

NEW QUESTION 44

- (Exam Topic 3)

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 question, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are developing a website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.

You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level. You need to configure authorization.

Solution:

- Create a new Azure AD application's manifest, set value of the groupMembershipClaims option to All.
- In the website, use the value of the groups claim from the JWT for the user to determine permissions. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

To configure Manifest to include Group Claims in Auth Token

* 1. Go to Azure Active Directory to configure the Manifest. Click on Azure Active Directory, and go to App registrations to find your application:

* 2. Click on your application (or search for it if you have a lot of apps) and edit the Manifest by clicking on it.

* 3. Locate the "groupMembershipClaims" setting. Set its value to either "SecurityGroup" or "All". To help you decide which:

"SecurityGroup" - groups claim will contain the identifiers of all security groups of which the user is a member.

"All" - groups claim will contain the identifiers of all security groups and all distribution lists of which the user is a member

Now your application will include group claims in your manifest and you can use this fact in your code. References:

<https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/>

NEW QUESTION 48

- (Exam Topic 3)

You are developing a new page for a website that uses Azure Cosmos DB for data storage. The feature uses documents that have the following format:

You must display data for the new page in a specific order. You create the following query for the page:

You need to configure a Cosmos DB policy to support the query.

How should you configure the policy? To answer, drag the appropriate JSON segments to the correct locations. Each JSON segment 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.

| JSON segments | Answer Area |
|------------------|---|
| orderBy | <pre>{ "automatic": true, "ngMode": "Consistent", "includedPaths": [{ "path": "/*" }], "excludedPaths": [], "": [{ "path": "/name", "order": "descending" }, { "path": "/city", "order": " " }] }</pre> |
| sortOrder | |
| ascending | |
| descending | |
| compositeIndexes | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: compositeIndexes

You can order by multiple properties. A query that orders by multiple properties requires a composite index. Box 2: descending

Example: Composite index defined for (name ASC, age ASC):

It is optional to specify the order. If not specified, the order is ascending.

```
{
  "automatic":true, "indexingMode":"Consistent", "includedPaths":[
  {
  "path": "/*"
  }
  ],
  "excludedPaths":[], "compositeIndexes":[ [
  {
  "path": "/name",
  },
  {
  "path": "/age",
  }
  ]
  ]
}
```

NEW QUESTION 53

- (Exam Topic 3)

You are developing a solution that will use Azure messaging services.

You need to ensure that the solution uses a publish-subscribe model and eliminates the need for constant polling.

What are two possible ways to achieve the goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Service Bus
- B. Event Hub
- C. Event Grid
- D. Queue

Answer: AC

Explanation:

It is strongly recommended to use available messaging products and services that support a publish-subscribe model, rather than building your own. In Azure, consider using Service Bus or Event Grid. Other technologies that can be used for pub/sub messaging include Redis, RabbitMQ, and Apache Kafka.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/patterns/publisher-subscriber>

NEW QUESTION 55

- (Exam Topic 3)

You develop a serverless application using several Azure Functions. These functions connect to data from within the code.

You want to configure tracing for an Azure Function App project. You need to change configuration settings in the hostjson file. Which tool should you use?

- A. Azure portal
- B. Azure PowerShell
- C. Azure Functions Core Tools (Azure CLI)
- D. Visual Studio

Answer: A

Explanation:

The function editor built into the Azure portal lets you update the function.json file and the code file for a function. The host.json file, which contains some runtime-specific configurations, is in the root folder of the function app.

References:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-reference#fileupdate>

NEW QUESTION 58

- (Exam Topic 3)

You are building a traffic monitoring system that monitors traffic along six highways. The system produces time series analysis-based reports for each highway.

Data from traffic sensors are stored in Azure Event Hub.

Traffic data is consumed by four departments. Each department has an Azure Web App that displays the time-series-based reports and contains a WebJob that processes the incoming data from Event Hub. All Web Apps run on App Service Plans with three instances.

Data throughout must be maximized. Latency must be minimized. You need to implement the Azure Event Hub.

Which settings should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Setting

Value

Number of partitions

| | |
|----|---|
| | ▼ |
| 3 | |
| 4 | |
| 6 | |
| 12 | |

Partition Key

| | |
|------------|---|
| | ▼ |
| Highway | |
| Department | |
| Timestamp | |
| VM name | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: 6
The number of partitions is specified at creation and must be between 2 and 32. There are 6 highways.
Box 2: Highway References:
<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features>

NEW QUESTION 60

- (Exam Topic 3)
You are writing code to create and run an Azure Batch job. You have created a pool of compute nodes. You need to choose the right class and its method to submit a batch job to the Batch service. Which method should you use?

- A. JobOperations.CreateJobO
- B. CloudJob.Enable(IEnumerable<BatchClientBehavior>)
- C. CloudJob.CommitAsync(IEnumerable<BatchClientBehavior>, CancellationToken)
- D. JobOperations.EnableJob(String, IEnumerable<BatchClientBehavior>)
- E. JobOperations.EnableJobAsync(Strin
- F. IEnumerable<BatchClientBehavior>. CancellationToken)

Answer: C

Explanation:

A Batch job is a logical grouping of one or more tasks. A job includes settings common to the tasks, such as priority and the pool to run tasks on. The app uses the BatchClient.JobOperations.CreateJob method to create a job on your pool. The Commit method submits the job to the Batch service. Initially the job has no tasks.

```
{
CloudJob job = batchClient.JobOperations.CreateJob(); job.Id = JobId;
job.PoolInformation = new PoolInformation { PoolId = PoolId }; job.Commit();
}
```

References:
<https://docs.microsoft.com/en-us/azure/batch/quick-run-dotnet>

NEW QUESTION 64

- (Exam Topic 3)
You are preparing to deploy an application to an Azure Kubernetes Service (AKS) cluster. The application must only be available from within the VNet that includes the cluster. You need to deploy the application. How should you complete the deployment YAML? To answer, drag the appropriate YAML segments to the correct locations. Each YAML segment 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.

| Code segments | Answer Area |
|------------------------------|---|
| Ingress | apiVersion: v1 |
| Service | kind: Code segment |
| LoadBalancer | metadata: |
| Deployment | name: web-app |
| ingress.class | annotations: |
| azure-load-balancer-internal | service.beta.kubernetes.Code segment : "true" |
| | spec: |
| | type: Code segment |
| | ports: |
| | - port: 80 |
| | selector: |
| | app: web-app |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

To create an internal load balancer, create a service manifest named internal-lb.yaml with the service type LoadBalancer and the azure-load-balancer-internal annotation as shown in the following example:

```
YAML:
apiVersion: v1 kind: Service metadata:
name: internal-app annotations:
service.beta.kubernetes.io/azure-load-balancer-internal: "true" spec:
type: LoadBalancer ports:
- port: 80 selector:
```

app: internal-app

References:

<https://docs.microsoft.com/en-us/azure/aks/internal-lb>

NEW QUESTION 65

- (Exam Topic 3)

You are developing a Docker/Go using Azure App Service Web App for Containers. You plan to run the container in an App Service on Linux. You identify a Docker container image to use.

None of your current resource groups reside in a location that supports Linux. You must minimize the number of resource groups required.

You need to create the application and perform an initial deployment.

Which three Azure CLI commands should you use to develop the solution? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the correct order.

| Azure CLI Commands | Answer Area |
|---------------------------|-------------|
| az group create | |
| az group update | |
| az webapp update | ⬅ |
| az webapp create | ⬆ |
| az appservice plan create | ⬆ |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

You can host native Linux applications in the cloud by using Azure Web Apps. To create a Web App for Containers, you must run Azure CLI commands that create a group, then a service plan, and finally the web app itself.

Step 1: az group create

In the Cloud Shell, create a resource group with the az group create command. Step 2: az appservice plan create

In the Cloud Shell, create an App Service plan in the resource group with the az appservice plan create command.

Step 3: az webapp create

In the Cloud Shell, create a web app in the myAppServicePlan App Service plan with the az webapp create command. Don't forget to replace with a unique app name, and <docker-ID> with your Docker ID.

References:

<https://docs.microsoft.com/mt-mt/azure/app-service/containers/quickstart-docker-go?view=sql-server-ver15>

NEW QUESTION 66

- (Exam Topic 3)

You are developing an ASP.NET Core website that can be used to manage photographs which are stored in Azure Blob Storage containers.

Users of the website authenticate by using their Azure Active Directory (Azure AD) credentials.

You implement role-based access control (RBAC) role permission on the containers that store photographs. You assign users to RBAC role.

You need to configure the website's Azure AD Application so that user's permissions can be used with the Azure Blob containers.

How should you configure the application? To answer, drag the appropriate setting to the correct location. Each setting 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.

| Settings | Answer Area | | |
|--------------------|-----------------|------------|---------|
| client_id | | | |
| delegated | | | |
| profile | | | |
| application | | | |
| user_impersonation | | | |
| | API | Permission | Type |
| | Azure Storage | Setting | Setting |
| | Microsoft Graph | User.Read | Setting |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: user_impersonation

Box 2: delegated Example:

- * 1. Select the API permissions section
- * 2. Click the Add a permission button and then: Ensure that the My APIs tab is selected
- * 3. In the list of APIs, select the API TodoListService-aspnetcore.
- * 4. In the Delegated permissions section, ensure that the right permissions are checked: user_impersonation. 5. Select the Add permissions button.

Box 3: delegated Example

- * 1. Select the API permissions section
- * 2. Click the Add a permission button and then, Ensure that the Microsoft APIs tab is selected
- * 3. In the Commonly used Microsoft APIs section, click on Microsoft Graph
- * 4. In the Delegated permissions section, ensure that the right permissions are checked: User.Read. Use the search box if necessary.
- * 5. Select the Add permissions button

References:
<https://docs.microsoft.com/en-us/samples/azure-samples/active-directory-dotnet-webapp-webapi-openidconnect>

NEW QUESTION 70

- (Exam Topic 3)

You are developing an Azure App Service hosted ASP.NET Core web app to deliver video on-demand streaming media. You enable an Azure Content Delivery Network (CDN) Standard for the web endpoint. Customer videos are downloaded from the web app by using the following example URL.:

<http://www.contoso.com/content.mp4?quality=1>

All media content must expire from the cache after one hour. Customer videos with varying quality must be delivered to the closest regional point of presence (POP) node.

You need to configure Azure CDN caching rules.

Which options should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

| Setting | Action |
|-------------------------------|--|
| Caching behavior | <div style="border: 1px solid black; padding: 2px;">▼</div> <div style="border: 1px solid black; padding: 2px;">Bypass cache</div> <div style="border: 1px solid black; padding: 2px;">Override</div> <div style="border: 1px solid black; padding: 2px;">Set if missing</div> |
| Cache expiration duration | <div style="border: 1px solid black; padding: 2px;">▼</div> <div style="border: 1px solid black; padding: 2px;">1 second</div> <div style="border: 1px solid black; padding: 2px;">1 minute</div> <div style="border: 1px solid black; padding: 2px;">1 hour</div> <div style="border: 1px solid black; padding: 2px;">1 day</div> |
| Query string caching behavior | <div style="border: 1px solid black; padding: 2px;">▼</div> <div style="border: 1px solid black; padding: 2px;">Ignore query strings</div> <div style="border: 1px solid black; padding: 2px;">Bypass caching for query strings</div> <div style="border: 1px solid black; padding: 2px;">Cache every unique URL</div> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Override

Override: Ignore origin-provided cache duration; use the provided cache duration instead. This will not override cache-control: no-cache.

Set if missing: Honor origin-provided cache-directive headers, if they exist; otherwise, use the provided cache duration.

Incorrect:

Bypass cache: Do not cache and ignore origin-provided cache-directive headers. Box 2: 1 hour

All media content must expire from the cache after one hour. Box 3: Cache every unique URL

Cache every unique URL: In this mode, each request with a unique URL, including the query string, is treated as a unique asset with its own cache. For example, the response from the origin server for a request for example.ashx?q=test1 is cached at the POP node and returned for subsequent caches with the same query string. A request for example.ashx?q=test2 is cached as a separate asset with its own time-to-live setting.

Reference:

<https://docs.microsoft.com/en-us/azure/cdn/cdn-query-string>

NEW QUESTION 74

- (Exam Topic 3)

You are deploying an Azure Kubernetes Services (AKS) cluster that will use multiple containers.

You need to create the cluster and verify that the services for the containers are configured correctly and available.

Which four commands should you use to develop the solution? To answer, move the appropriate command segments from the list of command segments to the answer area and arrange them in the correct order.

Command segments


az aks get-credentials

az appservice plan create

az aks create

az group create

kubectl apply



Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: az group create

Create a resource group with the az group create command. An Azure resource group is a logical group in which Azure resources are deployed and managed.

Example: The following example creates a resource group named myAKSCluster in the eastus location. az group create --name myAKSCluster --location eastus

Step 2 : az aks create

Use the az aks create command to create an AKS cluster. Step 3: kubectl apply

To deploy your application, use the kubectl apply command. This command parses the manifest file and creates the defined Kubernetes objects.

Step 4: az aks get-credentials

Configure it with the credentials for the new AKS cluster. Example:

az aks get-credentials --name aks-cluster --resource-group aks-resource-group References:

<https://docs.bitnami.com/azure/get-started-aks/>

NEW QUESTION 78

- (Exam Topic 3)

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.

Margie's Travel is an international travel and bookings management service. The company is expanding into restaurant bookings. You are tasked with implementing Azure Search for the restaurants listed in their solution You create the index in Azure Search.

You need to import the restaurant data into the Azure Search service by using the Azure Search NET SDK. Solution:

* 1 Create a SearchIndexClient object to connect to the search index

* 2. Create an IndexBatch that contains the documents which must be added.

* 3. Call the Documents.Index method of the SearchIndexClient and pass the IndexBatch..

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

* 1. The index needs to be populated. To do this, we will need a SearchIndexClient. There are two ways to obtain one: by constructing it, or by calling Indexes.GetClient on the SearchServiceClient. Here we will use the first method.

* 2. Create the indexBatch with the documents Something like:

```
var hotels = new Hotel[];
```

```
{  
    new Hotel()  
    {  
        HotelId = "3",  
        BaseRate = 129.99,  
        Description = "Close to town hall and the river"  
    }  
};  
...
```

```
var batch = IndexBatch.Upload(hotels);
```

* 3. The next step is to populate the newly-created index Example:

```
var batch = IndexBatch.Upload(hotels); try
```

```
{  
    indexClient.Documents.Index(batch);  
}
```

References:

<https://docs.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk>

NEW QUESTION 79

- (Exam Topic 3)

You develop a solution that uses an Azure SQL Database to store user information for a mobile app. The app stores sensitive information about users.

You need to hide sensitive information from developers that query the data for the mobile app.

Which three items must you identify when configuring dynamic data masking? Each correct answer presents a part of the solution.

NOTE: Each correct selection is worth one point.

- A. Column
- B. Table
- C. Trigger
- D. Index
- E. Schema

Answer: ABE

Explanation:

In the Dynamic Data Masking configuration page, you may see some database columns that the recommendations engine has flagged for masking. In order to accept the recommendations, just click Add Mask for one or more columns and a mask is created based on the default type for this column. You can change the masking function by clicking on the masking rule and editing the masking field format to a different format of your choice.

Dynamic Data Masking
demo_database

Save Discard Add Mask

Downlevel clients require the use of Security Enabled Connection Strings.

Masking Rules

| MASK NAME | MASK FUNCTION |
|--|---------------|
| You haven't created any masking rules. | |

SQL users excluded from masking (administrators are always excluded)

SQL users excluded from masking (administrators are always excluded) ✓

Recommended fields to mask

| SCHEMA | TABLE | COLUMN | |
|---------|-----------------|--------------|----------|
| SalesLT | Customer | FirstName | ADD MASK |
| SalesLT | Customer | LastName | ADD MASK |
| SalesLT | Customer | EmailAddress | ADD MASK |
| SalesLT | Customer | Phone | ADD MASK |
| SalesLT | CustomerAddress | AddressID | ADD MASK |

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dynamic-data-masking-get-started-portal>

NEW QUESTION 81

- (Exam Topic 3)

You develop Azure solutions.

You must connect to a No-SQL globally-distributed database by using the .NET API. You need to create an object to configure and execute requests in the database. Which code segment should you use?

- A. new Container(EndpointUri, PrimaryKey);
- B. new Database(Endpoint, PrimaryKey);
- C. new CosmosClient(EndpointUri, PrimaryKey);

Answer: C

Explanation:

Example:

```
// Create a new instance of the Cosmos Client
this.cosmosClient = new CosmosClient(EndpointUri, PrimaryKey)
```

//ADD THIS PART TO YOUR CODE

```
await this.CreateDatabaseAsync();
```

Reference:
<https://docs.microsoft.com/en-us/azure/cosmos-db/sql-api-get-started>

NEW QUESTION 84

- (Exam Topic 3)

You are developing a solution for a hospital to support the following use cases:

•The most recent patient status details must be retrieved even if multiple users in different locations have updated the patient record.
 •Patient health monitoring data retrieved must be the current version or the prior version.
 •After a patient is discharged and all charges have been assessed, the patient billing record contains the final charges.
 You provision a Cosmos DB NoSQL database and set the default consistency level for the database account to Strong. You set the value for Indexing Mode to Consistent.
 You need to minimize latency and any impact to the availability of the solution. You must override the default consistency level at the query level to meet the required consistency guarantees for the scenarios.
 Which consistency levels should you implement? To answer, drag the appropriate consistency levels to the correct requirements. Each consistency level 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.

| Consistency levels | | Answer Area |
|--|--|--|
| <input type="checkbox"/> Strong | <input type="checkbox"/> Bounded Staleness | Return the most recent patient status. |
| <input type="checkbox"/> Consistent Prefix | <input type="checkbox"/> Eventual | Return health monitoring data that is no less than one version behind. |
| | | After patient is discharged and all changes are assessed, retrieve the correct billing data with the final charges |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Strong

Strong: Strong consistency offers a linearizability guarantee. The reads are guaranteed to return the most recent committed version of an item. A client never sees an uncommitted or partial write. Users are always guaranteed to read the latest committed write.

Box 2: Bounded staleness

Bounded staleness: The reads are guaranteed to honor the consistent-prefix guarantee. The reads might lag behind writes by at most "K" versions (that is "updates") of an item or by "t" time interval. When you choose bounded staleness, the "staleness" can be configured in two ways:

The number of versions (K) of the item

The time interval (t) by which the reads might lag behind the writes

Box 3: Eventual

Eventual: There's no ordering guarantee for reads. In the absence of any further writes, the replicas eventually converge.

NEW QUESTION 87

- (Exam Topic 3)

You are creating a hazard notification system that has a single signaling server which triggers audio and visual alarms to start and stop. You implement Azure Service Bus to publish alarms. Each alarm controller uses Azure Service Bus to receive alarm signals as part of a transaction. Alarm events must be recorded for audit purposes. Each transaction record must include information about the alarm type that was activated. You need to implement a reply trail auditing solution. Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Assign the value of the hazard message SessionID property to the SequenceNumber property.
- B. Assign the value of the hazard message SequenceNumber property to the DeliveryCount property.
- C. Assign the value of the hazard message MessageId property to the DeliveryCount property.
- D. Assign the value of the hazard message SessionID property to the ReplyToSessionId property.
- E. Assign the value of the hazard message MessageId property to the SequenceNumber property.
- F. Assign the value of the hazard message MessageId property to the CorrelationId property.

Answer: AB

NEW QUESTION 89

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AZ-204 Exam with Our Prep Materials Via below:

<https://www.certleader.com/AZ-204-dumps.html>