Microsoft

Exam Questions AZ-204

Developing Solutions for Microsoft Azure (beta)

NEW QUESTION 1
- (Exam Topic 1)
You need to resolve the Shipping web site error.
How should you configre the Azure Table Storage service? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

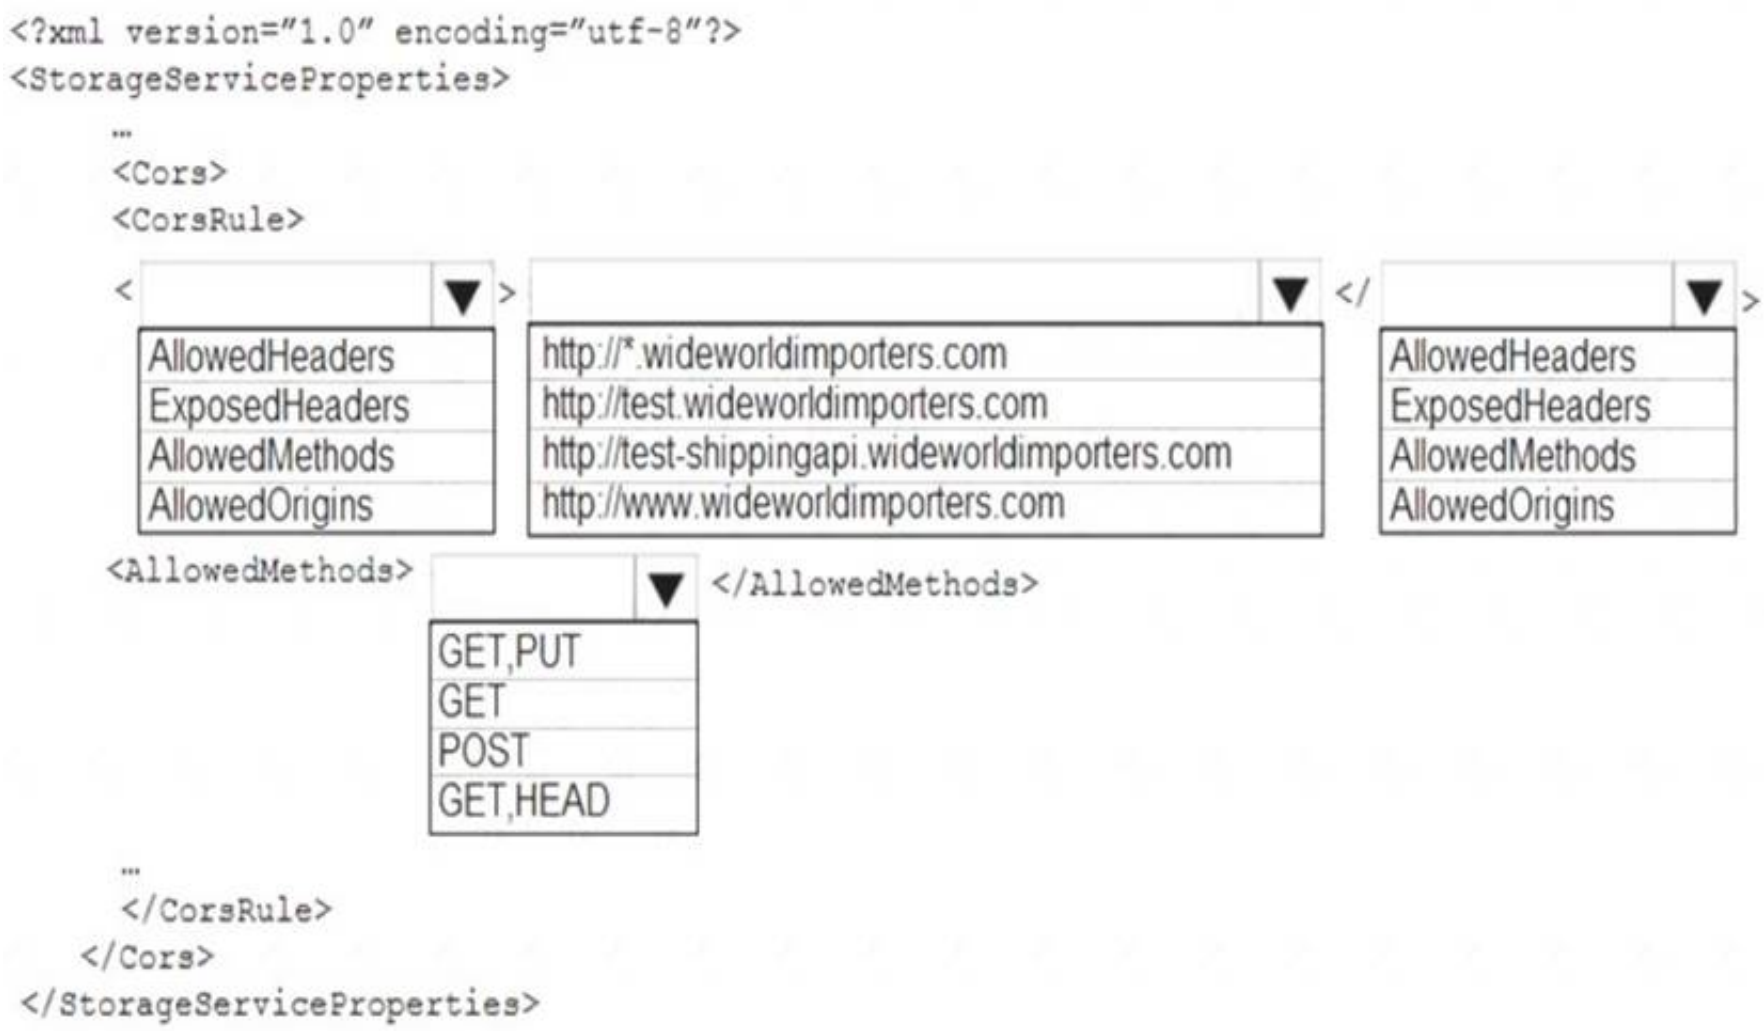

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: AllowedOrigins
A CORS request will fail if Access-Control-Allow-Origin is missing. Scenario:
The following error message displays while you are testing the website:

```
Failed to load http://test-shippingapi.wideworldimporters.com/: No 'Access-Control-Allow-Origin'
header is present on the requested resource. Origin 'http://testwideworldimporters.com/' is
therefore not allowed access.
```

Box 2: http://test-shippingapi.wideworldimporters.com Syntax: Access-Control-Allow-Origin: *
Access-Control-Allow-Origin: <origin> Access-Control-Allow-Origin: null
<origin> Specifies an origin. Only a single origin can be specified. Box 3: AllowedOrigins
Box 4: POST
The only allowed methods are GET, HEAD, and POST. In this case POST is used. "<Corsrule>" "allowedmethods" Failed to load no "Access-control-Origin" header is present References:
https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Access-Control-Allow-Origin

NEW QUESTION 2
- (Exam Topic 1)
You need to support the message processing for the ocean transport workflow.
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.

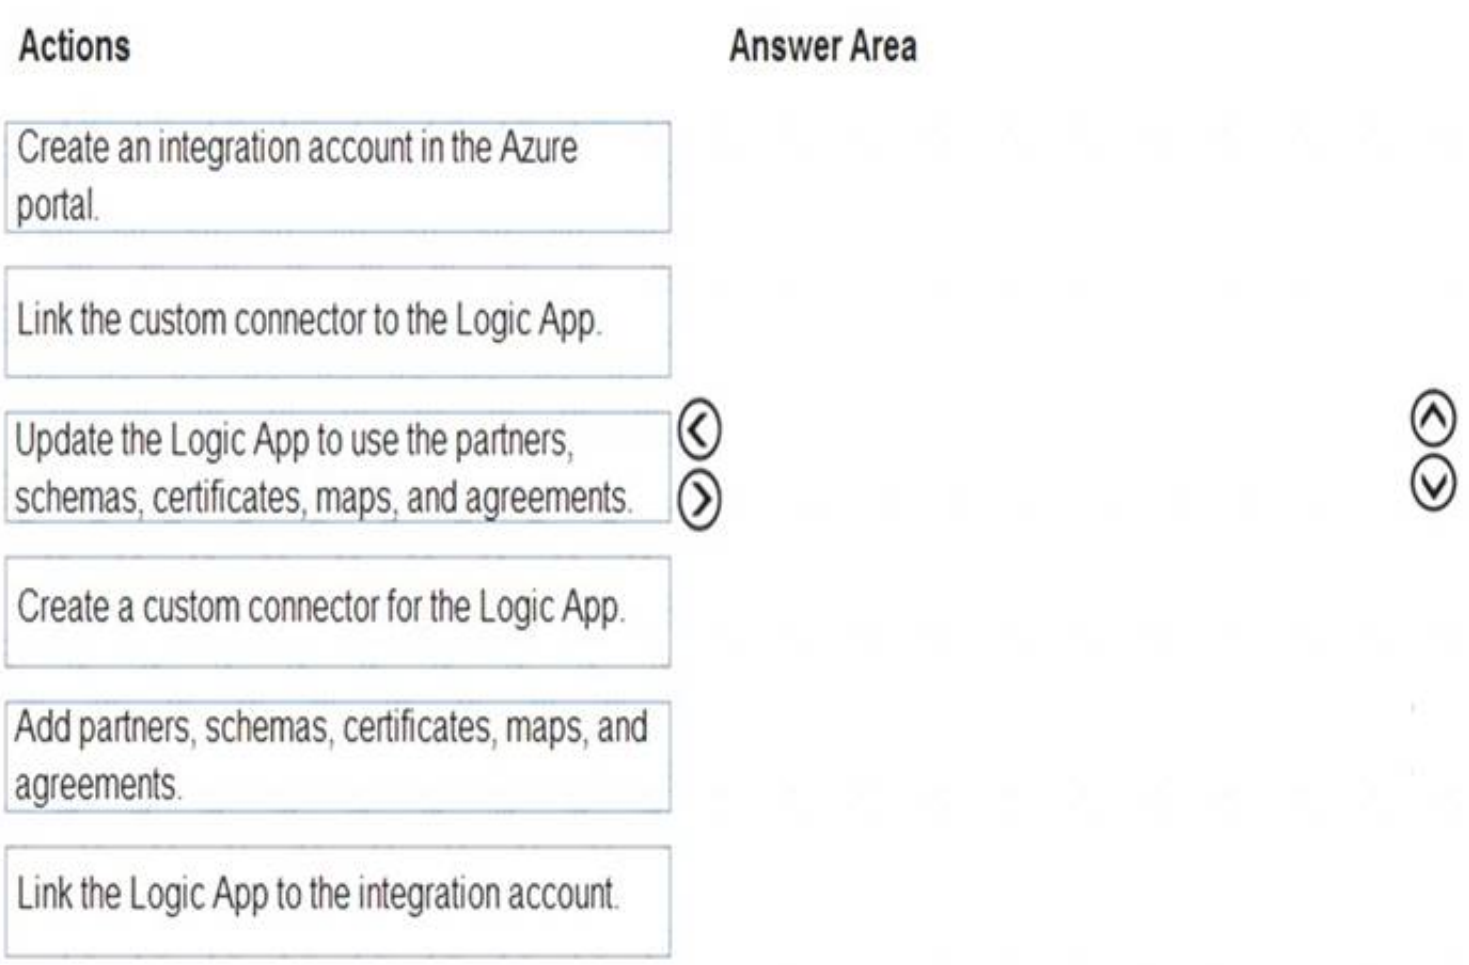

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Create an integration account in the Azure portal
You can define custom metadata for artifacts in integration accounts and get that metadata during runtime for your logic app to use. For example, you can provide metadata for artifacts, such as partners, agreements, schemas, and maps - all store metadata using key-value pairs.
Step 2: Link the Logic App to the integration account
A logic app that's linked to the integration account and artifact metadata you want to use. Step 3: Add partners, schemas, certificates, maps, and agreements
Step 4: Create a custom connector for the Logic App. References:
https://docs.microsoft.com/bs-latn-ba/azure/logic-apps/logic-apps-enterprise-integration-metadata

NEW QUESTION 3
- (Exam Topic 1)
You need to configure Azure CDN for the Shipping web site.
Which configuration options should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

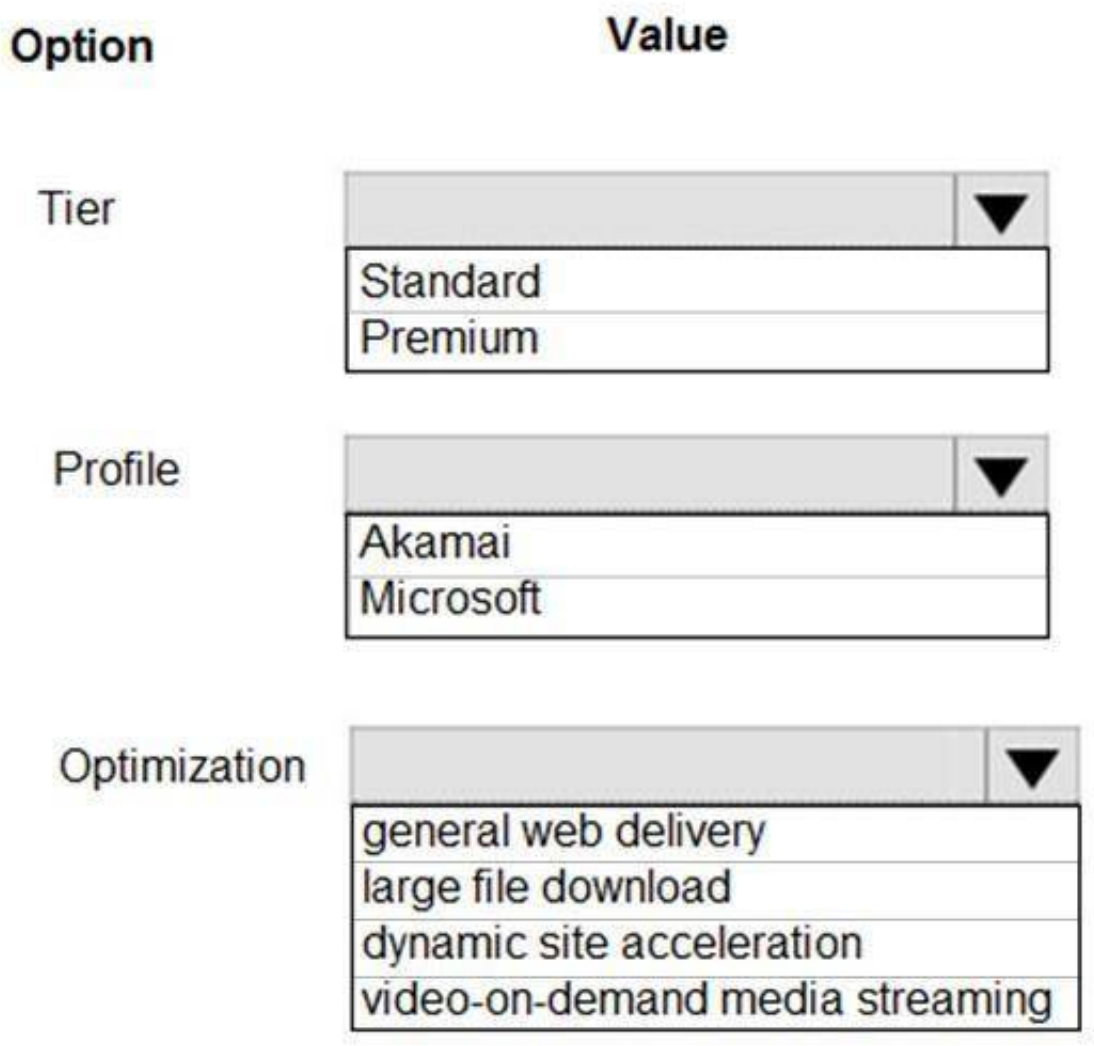

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Scenario: Shipping website
Use Azure Content Delivery Network (CDN) and ensure maximum performance for dynamic content while minimizing latency and costs.
Tier: Standard Profile: Akamai
Optimization: Dynamic site acceleration
Dynamic site acceleration (DSA) is available for Azure CDN Standard from Akamai, Azure CDN Standard from Verizon, and Azure CDN Premium from Verizon profiles.
DSA includes various techniques that benefit the latency and performance of dynamic content. Techniques include route and network optimization, TCP optimization, and more.
You can use this optimization to accelerate a web app that includes numerous responses that aren't cacheable. Examples are search results, checkout transactions, or real-time data. You can continue to use core Azure CDN caching capabilities for static data.
Reference:
https://docs.microsoft.com/en-us/azure/cdn/cdn-optimization-overview

NEW QUESTION 4
- (Exam Topic 1)
You need to secure the Shipping Function app.
How should you configure the app? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Setting	Value
Authorization level	Function / Anonymous / Admin
User claims	JSON Web Token (JWT) / Shared Access Signature (SAS) token / API Key
Trigger type	blob / HTTP / queue / timer

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Scenario: Shipping Function app: Implement secure function endpoints by using app-level security and include Azure Active Directory (Azure AD).
Box 1: Function
Box 2: JSON based Token (JWT)
Azure AD uses JSON based tokens (JWTs) that contain claims Box 3: HTTP
How a web app delegates sign-in to Azure AD and obtains a token
User authentication happens via the browser. The OpenID protocol uses standard HTTP protocol messages. References:
https://docs.microsoft.com/en-us/azure/active-directory/develop/authentication-scenarios

NEW QUESTION 5
- (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 6
- (Exam Topic 3)
You are developing an Azure Function App by using Visual Studio. The app will process orders input by an Azure Web App. The web app places the order information into Azure Queue Storage.
You need to review the Azure Function App code shown below. NOTE: Each correct selection is worth one point.

```
public static class
{
  [FunctionName("ProcessOrders")]
  public static void ProcessOrders([QueueTrigger("incoming-orders")]CloudQueueMessage myQueueItem, [Table
("Orders")]ICollector<Order> tableBindings,
TraceWriter log)
  {
    log.Info($"Processing Order: {myQueueItem.Id}");
    log.Info($"Queue Insertion Time: {myQueueItem.InsertionTime}");
    log.Info($"Queue Expiration Time: {myQueueItem.ExpirationTime}");
    tableBindings.Add(JsonConvert.DeserializeObject<Order>(myQueueItem.AsString));
  }
  [FunctionName("ProcessOrders-Poison")]
  public static void ProcessFailedOrders([QueueTrigger("incoming-orders-poison")]CloudQueueMessage
myQueueItem, TraceWriter log)
  {
    logError($"Failed to process order: {myQueueItem.AsString}");
  }
}
```

	Yes	No
The code will log the time that the order was processed from the queue.	○	○
When the ProcessOrders function fails, the function will retry up to five times for a given order, including the first try.	○	○
When there are multiple orders in the queue, a batch of orders will be retrieved from the queue and the ProcessOrders function will run multiple instances concurrently to process the orders.	○	○
The ProcessOrders function will output the order to an Orders table in Azure Table Storage.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: No
ExpirationTime - The time that the message expires. InsertionTime - The time that the message was added to the queue.
Box 2: Yes
maxDequeueCount - The number of times to try processing a message before moving it to the poison queue.
Default value is 5.
Box 3: Yes
When there are multiple queue messages waiting, the queue trigger retrieves a batch of messages and invokes function instances concurrently to process them.
By default, the batch size is 16. When the number being processed gets down to 8, the runtime gets another batch and starts processing those messages. So the maximum number of concurrent messages being processed per function on one virtual machine (VM) is 24.
Box 4: Yes References:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-queue

NEW QUESTION 7
- (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
Yes	Enable AlwaysOn encryption.	
No	Set the column encryption setting to disabled.	
	Assign users to the Public fixed database role.	
	Store column encryption keys in the system catalog view in the database.	

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 8
- (Exam Topic 3)
You are creating a script that will run a large workload on an Azure Batch pool. Resources will be reused and do not need to be cleaned up after use.
You have the following parameters:
You need to write an Azure CLI script that will create the jobs, tasks, and the pool.
In which order should you arrange the commands to develop the solution? To answer, move the appropriate commands from the list of command segments to the answer area and arrange them in the correct order.

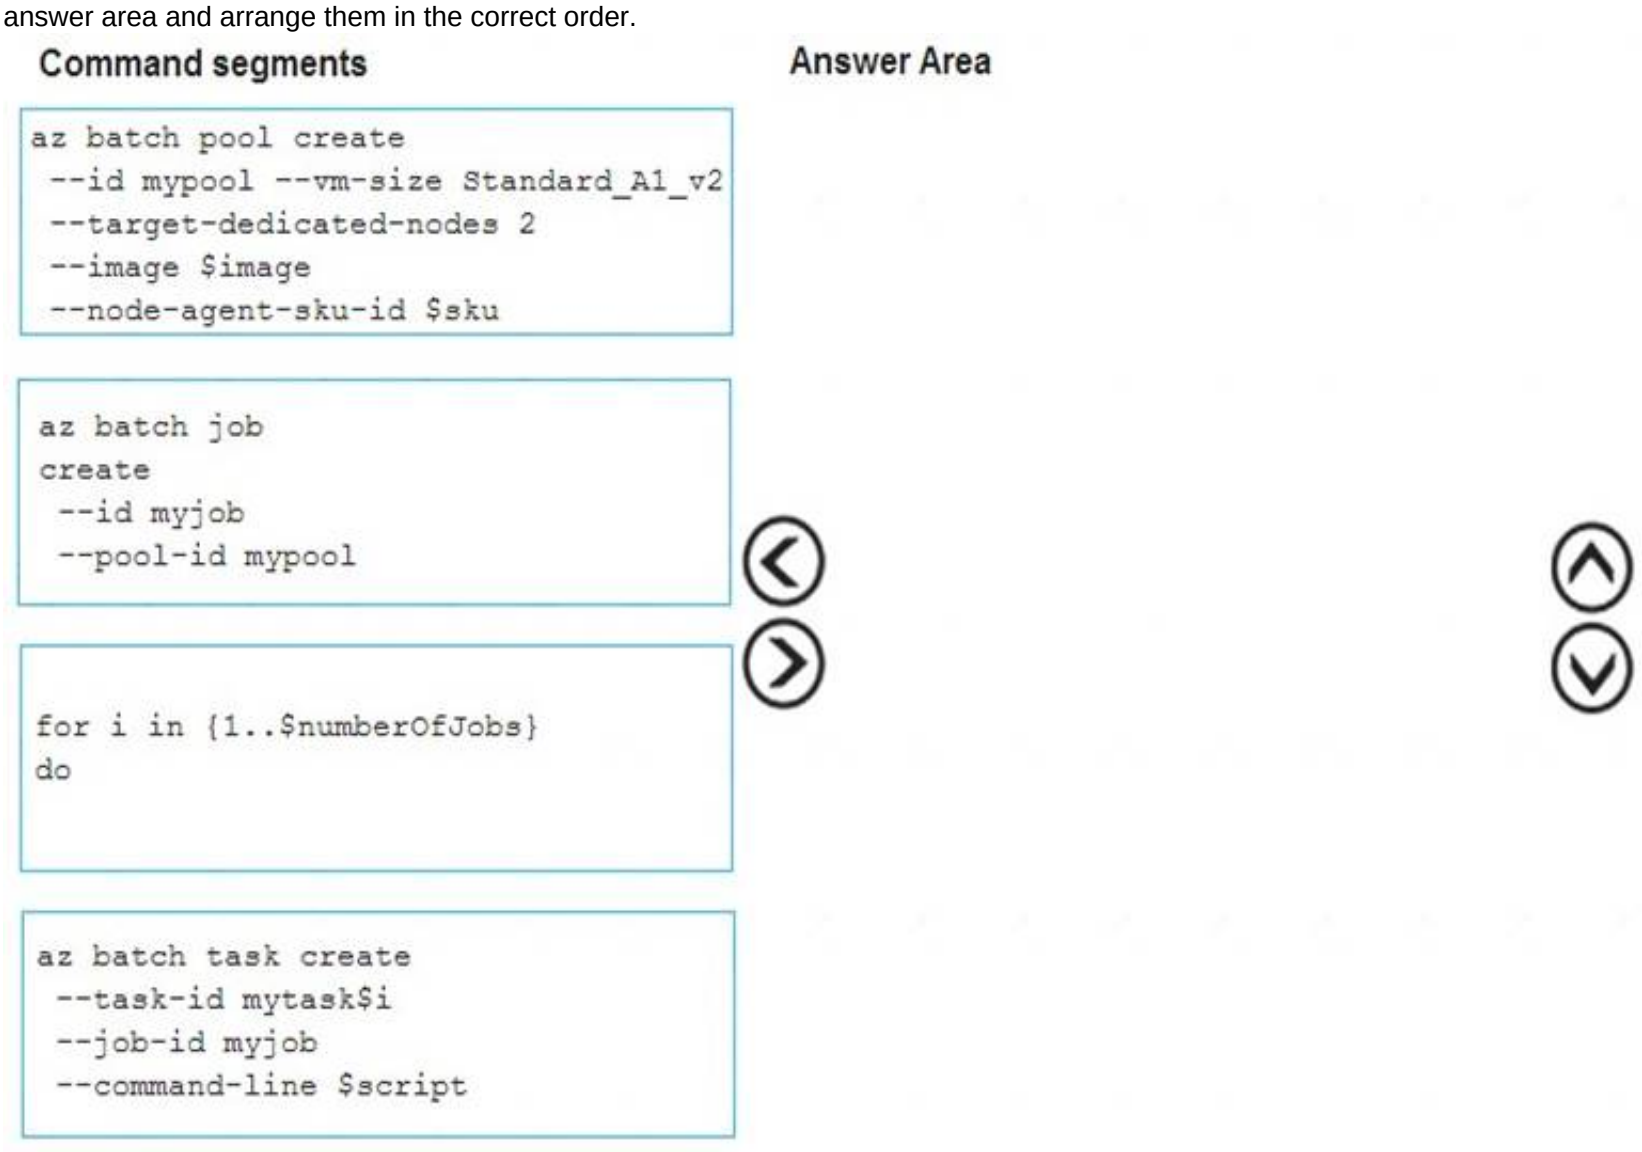

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: az batch pool create
Create a new Linux pool with a virtual machine configuration. az batch pool create \
--id mypool \
--vm-size Standard_A1 \
--target-dedicated 2 \
--image canonical:ubuntuserver:16.04-LTS \
--node-agent-sku-id "batch.node.ubuntu 16.04" Step 2: az batch job create
Create a new job to encapsulate the tasks that are added. az batch job create \
--id myjob \
--pool-id mypool
Step 3: az batch task create
Add tasks to the job. Here the task is a basic shell command. az batch task create \
--job-id myjob \
--task-id task1 \
--command-line "/bin/bash -c 'printenv AZ_BATCH_TASK_WORKING_DIR"' Step 4: for i in {1..$numberOfJobs} do
References:
https://docs.microsoft.com/bs-latn-ba/azure/batch/scripts/batch-cli-sample-run-job

NEW QUESTION 9
- (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 10
- (Exam Topic 3)
You develop an app that allows users to upload photos and videos to Azure storage. The app uses a storage REST API call to upload the media to a blob storage account named Account1. You have blob storage
containers named Container1 and Container2. Uploading of videos occurs on an irregular basis.
You need to copy specific blobs from Container1 to Container2 in real time when specific requirements are met, excluding backup blob copies.
What should you do?

A. Download the blob to a virtual machine and then upload the blob to Container2.
B. Run the Azure PowerShell command Start-AzureStorageBlobCopy.
C. Copy blobs to Container2 by using the Put Blob operation of the Blob Service REST API.
D. Use AzCopy with the Snapshot switch blobs to Container2.

Answer: B

Explanation:
The Start-AzureStorageBlobCopy cmdlet starts to copy a blob. Example 1: Copy a named blob
C:\PS>Start-AzureStorageBlobCopy -SrcBlob "ContosoPlanning2015" -DestContainer "ContosoArchives"
-SrcContainer "ContosoUploads"
This command starts the copy operation of the blob named ContosoPlanning2015 from the container named ContosoUploads to the container named ContosoArchives.
References:
https://docs.microsoft.com/en-us/powershell/module/azure.storage/start-azurestorageblobcopy?view=azurermps

NEW QUESTION 10
- (Exam Topic 3)
You are a developer for a SaaS company that offers many web services. All web services for the company must meet the following requirements:
- Use API Management to access the services
- Use OpenID Connect for authentication
- Prevent anonymous usage

A recent security audit found that several web services can be called without any authentication. Which API Management policy should you implement?

A. jsonp
B. authentication-certificate
C. check-header
D. validate-jwt

Answer: D

Explanation:
Add the validate-jwt policy to validate the OAuth token for every incoming request. Reference:
https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-protect-backend-with-aad

NEW QUESTION 13
- (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 are developing an Azure solution to collect point-of-sale (POS) device data from 2,000 stores located throughout the world. A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data.
You must store the device data in Azure Blob storage. Device data must be correlated based on a device identifier. Additional stores are expected to open in the future.
You need to implement a solution to receive the device data.
Solution: Provision an Azure Service Bus. Configure a topic to receive the device data by using a correlation filter.
Does the solution meet the goal?

A. Yes
B. No

Answer: A

Explanation:
A message is raw data produced by a service to be consumed or stored elsewhere. The Service Bus is for high-value enterprise messaging, and is used for order processing and financial transactions.
Reference:
https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services

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 are developing an Azure solution to collect point-of-sale (POS) device data from 2,000 stores located throughout the world. A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data.
You must store the device data in Azure Blob storage. Device data must be correlated based on a device identifier. Additional stores are expected to open in the future.
You need to implement a solution to receive the device data.
Solution: Provision an Azure Notification Hub. Register all devices with the hub. Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Instead use an Azure Service Bus, which is used order processing and financial transactions.
Reference:
https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services

NEW QUESTION 15
- (Exam Topic 3)
You are developing an Azure Cosmos DB solution by using the Azure Cosmos DB SQL API. The data includes millions of documents. Each document may contain hundreds of properties.
The properties of the documents do not contain distinct values for partitioning. Azure Cosmos DB must scale individual containers in the database to meet the performance needs of the application by spreading the workload evenly across all partitions over time.
You need to select a partition key.
Which two partition keys can you use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. a concatenation of multiple property values with a random suffix appended
B. a single property value that does not appear frequently in the documents
C. a hash suffix appended to a property value
D. a value containing the collection name
E. a single property value that appears frequently in the documents

Answer: AC

Explanation:
You can form a partition key by concatenating multiple property values into a single artificial partitionKey property. These keys are referred to as synthetic keys.
Another possible strategy to distribute the workload more evenly is to append a random number at the end of the partition key value. When you distribute items in this way, you can perform parallel write operations across partitions.
Note: It's the best practice to have a partition key with many distinct values, such as hundreds or thousands. The goal is to distribute your data and workload evenly across the items associated with these partition key values. If such a property doesn't exist in your data, you can construct a synthetic partition key.
References:
https://docs.microsoft.com/en-us/azure/cosmos-db/synthetic-partition-keys

NEW QUESTION 16
- (Exam Topic 3)
You are implementing a software as a service (SaaS) ASP.NET Core web service that will run as an Azure Web App. The web service will use an on-premises SQL Server database for storage. The web service also includes a WebJob that processes data updates. Four customers will use the web service.

- Each instance of the WebJob processes data for a single customer and must run as a singleton instance.
- Each deployment must be tested by using deployment slots prior to serving production data.
- Azure costs must be minimized.
- Azure resources must be located in an isolated network. You need to configure the App Service plan for the Web App.

How should you configure the App Service plan? To answer, select the appropriate settings in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

App service plan setting	Value
Number of VM instances	2 / 4 / 8 / 16
Pricing tier	Isolated / Standard / Premium / Consumption

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Number of VM instances: 4
You are not charged extra for deployment slots. Pricing tier: Isolated
The App Service Environment (ASE) is a powerful feature offering of the Azure App Service that gives network isolation and improved scale capabilities. It is essentially a deployment of the Azure App Service into a subnet of a customer's Azure Virtual Network (VNet).
References:
https://azure.microsoft.com/sv-se/blog/announcing-app-service-isolated-more-power-scale-and-ease-of-use/

NEW QUESTION 19
- (Exam Topic 3)
You are developing an ASP.NET Core Web API web service. The web service uses Azure Application Insights for all telemetry and dependency tracking. The web service reads and writes data to a database other than Microsoft SQL Server.
You need to ensure that dependency tracking works for calls to the third-party database.
Which two Dependency Telemetry properties should you store in the database? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Telemetry.Context.Operation.Id
B. Tetemetry.Context.Cloud.RoleInstance
C. Telemetry.Id
D. Telemetry.ContextSession.Id
E. Telemetry.Name

Answer: AC

Explanation:
References:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/custom-operations-tracking

NEW QUESTION 22
- (Exam Topic 3)
You are building a website to access project data related to terms within your organization. The website does not allow anonymous access. Authentication performed using an Azure Active Directory (Azure AD) app named internal.
The website has the following authentication requirements:
•Azure AD users must be able to login to the website.
•Personalization of the website must be based on membership in Active Directory groups. You need to configure the application's manifest to meet the authentication requirements.
How should you configure the manifest? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

```
{
    ...
    "appId": "d61126e3-089b-4adb-b721-
d5023213df7d",
    [ ▼ ] : "All",
    [ ▼ ] : true
    ...
}
```

- "optionalClaims"
- "groupMembershipClaims"

- "allowPublicClient"
- "oauth2Permissions"
- "requiredResourceAccess"
- "oauth2AllowImplicitFlow"

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: groupMembershipClaims
Scenario: Personalization of the website must be based on membership in Active Directory groups. Group claims can also be configured in the Optional Claims section of the Application Manifest. Enable group membership claims by changing the groupMembershipClaim
The valid values are: "All" "SecurityGroup" "DistributionList" "DirectoryRole"
Box 2: oauth2Permissions
Scenario: Azure AD users must be able to login to the website.
oauth2Permissions specifies the collection of OAuth 2.0 permission scopes that the web API (resource) app exposes to client apps. These permission scopes may be granted to client apps during consent.

NEW QUESTION 24
- (Exam Topic 3)
You have an Azure Batch project that processes and converts files and stores the files in Azure storage. You are developing a function to start the batch job.
You add the following parameters to the function.

Parameter name	Description
fileTasks	a list of tasks to be run
jobId	the identifier that must be assigned to the job
outputContainerSasUrl	a storage SAS URL to store successfully converted files
failedContainerSasUrl	a storage SAS URL to store copies of files that failed to convert.

You must ensure that converted files are placed in the container referenced by the outputContainerSasUrl parameter. Files which fail to convert are places in the container referenced by the failedContainerSasUrl parameter.
You need to ensure the files are correctly processed.
How should you complete the code segment? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

```
public list<CloudTasks> StartTasks(List<FileTask> fileTasks, string jobId,
 string outputContainerSasUrl, string failedContainerSasUrl)
{
  BatchSharedKeyCredentials sharedKeyCredentials =
   new BatchSharedKeyCredentials(batchAccountUrl, batchAccountName,
batchAccountKey);
 List<CloudTask> tasks = new List<CloudTask>();
 using (BatchClient batchClient = BatchClient.Open(sharedKeyCredentials))
{
   CloudJob = batchClient.JobOperations.                    ();
```

- GetJob
- GetTask
- EnableJob
- CreateJob

```
        job.Id = jobId,
        job.PoolInfromation = new PoolInformation { PoolId = poolId };
        job.Commit();
        fileTasks.ForEach((fileTask) =>
        {
         string taskId = $"Task{DateTime.Now.ToFileTimeUtc().ToString()}";
         CloudTask task = new CloudTask (taskId, fileTask.Command);
         List<OutputFile> outputFileList = new List<OutputFile>();
         OutputFileBlobContainerDestination outputContainer =
          new OutputFileBlobContainerDestination(outputContainerSasUrl);
       OutputFileBlobContainerDestination failedContainer =
   new OutputFileBlobContainerDestination (failedContainerSasUrl);
outputFileList.Add(new OutputFile(fileTask.Output,
   new OutputFileDestination(outputContainer),
   new OutputFileUploadOptions(OutputFileUploadCondition.                ));
```

- TaskSuccess
- TaskFailure
- TaskCompletion

```
  outputFileList.Add(new OutputFile(fileTask.Output,
   new OutputFileDestination(failedContainer),
   new OutputFileUploadOptions(OutputFileUploadCondition,                ));
```

- TaskSuccess
- TaskFailure
- TaskCompletion

```
    task                =outputFileList;
```

- OutputFiles
- FilesToStage
- ResourceFiles
- StageFiles

```
      task.Add(task);
      });
      ]
      return tasks,
}
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: CreateJob
Box 2: TaskSuccess
TaskSuccess: Upload the file(s) only after the task process exits with an exit code of 0.
Incorrect: TaskCompletion: Upload the file(s) after the task process exits, no matter what the exit code was. Box 3: TaskFailure
TaskFailure:Upload the file(s) only after the task process exits with a nonzero exit code. Box 4: OutputFiles
To specify output files for a task, create a collection of OutputFile objects and assign it to the CloudTask.OutputFiles property when you create the task.

References:
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.batch.protocol.models.outputfileuploadcondition https://docs.microsoft.com/en-us/azure/batch/batch-task-output-files

NEW QUESTION 27
- (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 28
- (Exam Topic 3)
Contoso, Ltd. provides an API to customers by using Azure API Management (APIM). The API authorizes users with a JWT token.
You must implement response caching for the APIM gateway. The caching mechanism must detect the user ID of the client that accesses data for a given location and cache the response for that user ID.
You need to add the following policies to the policies file:
• a set-variable policy to store the detected user identity
• a cache-lookup-value policy
• a cache-store-value policy
• a find-and-replace policy to update the response body with the user profile information
To which policy section should you add the policies? To answer, drag the appropriate sections to the correct policies. Each section 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

Policy section	Policy	Policy section
	Set-variable	
Inbound	Cache-lookup-value	
Outbound	Cache-store-value	
	Find-and-replace	

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Inbound.
A set-variable policy to store the detected user identity. Example:
<policies>
<inbound>
<!-- How you determine user identity is application dependent -->
<set-variable name="enduserid"
value="@(context.Request.Headers.GetValueOrDefault("Authorization","").Split(' ')[1].AsJwt()?.Subject)" /> Box 2: Inbound
A cache-lookup-value policy Example:
<inbound>
<base />
<cache-lookup vary-by-developer="true | false" vary-by-developer-groups="true | false" downstream-caching-type="none | private | public" must-revalidate="true | false">
<vary-by-query-parameter>parameter name</vary-by-query-parameter> <!-- optional, can repeated several times -->
</cache-lookup>
</inbound>

Box 3: Outbound
A cache-store-value policy. Example:
<outbound>
<base />
<cache-store duration="3600" />
</outbound> Box 4: Outbound
A find-and-replace policy to update the response body with the user profile information. Example:
<outbound>
<!-- Update response body with user profile-->
<find-and-replace from="'$userprofile$"'
to="@((string)context.Variables["userprofile"])" />
<base />
</outbound> Reference:
https://docs.microsoft.com/en-us/azure/api-management/api-management-caching-policies
https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key

NEW QUESTION 30
- (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();
[ ▼ ]
    i["tip"] = 0;
  }
[ ▼ ]
```

- _.value();
- _.readDocument('item');
- getContext().getRequest();
- getContext().getResponse();

- if (!("tip" in i)) {
- if (request.getValue("tip") === null){
- if (isNaN(i)["tip"] || i["tip"]=== null) {
- if (typeof_.pluck("tip") == 'number') {

- r.setBody(i);
- r.setValue(i);
- _.upsertDocument(i);
- _.replaceDocument(i)

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 34
- (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    Table Query<DynamicTableEntity> query - new TableQuery<DynamicTableEntity>().Select(new.string[.] { "Score" })
       .Where(TableQuery.GenerateFilterConditionForInt("Score", QueryComparisons.GretaerThanOrEqual, 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    {
      Console.Writeline($"{scoreItem.Key} {scoreItem.Value}");
7    }
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	○	○
The code will display a maximum of twenty records.	○	○
All records will be sent to the client. The client will display records for scores greater than or equal to 15,000.	○	○
The scoreItem.Key property of the KeyValuePairs that ExecuteQuery returns will contain a value for PlayerID.	○	○

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 36
- (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 solution that will be deployed to an Azure Kubernetes Service (AKS) cluster. The solution will include a custom VNet, Azure Container Registry images, and an Azure Storage account.
The solution must allow dynamic creation and management of all Azure resources within the AKS cluster. You need to configure an AKS cluster for use with the Azure APIs.

Solution: Enable the Azure Policy Add-on for Kubernetes to connect the Azure Policy service to the GateKeeper admission controller for the AKS cluster. Apply a built-in policy to the cluster.
Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Instead create an AKS cluster that supports network policy. Create and apply a network to allow traffic only from within a defined namespace
References:
https://docs.microsoft.com/en-us/azure/aks/use-network-policies

NEW QUESTION 37
- (Exam Topic 3)
You are preparing to deploy an ASP.NET Core website to an Azure Web App from a GitHub repository. The website includes static content generated by a script.
You plan to use the Azure Web App continuous deployment feature.
You need to run the static generation script before the website starts serving traffic.
What are two possible ways to achieve this goal? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. Create a file named .deployment in the root of the repository that calls a script which generates the static content and deploys the website.
B. Add a PreBuild target in the websites csproj project file that runs the static content generation script.
C. Create a file named run.cmd in the folder /run that calls a script which generates the static content and deploys the website.
D. Add the path to the static content generation tool to WEBSITE_RUN_FROM_PACKAGE setting in the host.json file.

Answer: AD

Explanation:
A: To customize your deployment, include a .deployment file in the repository root.
You just need to add a file to the root of your repository with the name .deployment and the content: [config]
command = YOUR COMMAND TO RUN FOR DEPLOYMENT
this command can be just running a script (batch file) that has all that is required for your deployment, like copying files from the repository to the web root directory for example.
D: In Azure, you can run your functions directly from a deployment package file in your function app. The other option is to deploy your files in the d:\home\site\wwwroot directory of your function app (see A above).
To enable your function app to run from a package, you just add a WEBSITE_RUN_FROM_PACKAGE setting to your function app settings.
Note: The host.json metadata file contains global configuration options that affect all functions for a function app.
References:
https://github.com/projectkudu/kudu/wiki/Custom-Deployment-Script
https://docs.microsoft.com/bs-latn-ba/azure/azure-functions/run-functions-from-deployment-package

NEW QUESTION 40
- (Exam Topic 3)
You develop software solutions for a mobile delivery service. You are developing a mobile app that users can use to order from a restaurant in their area. The app uses the following workflow:
* 1. A driver selects the restaurants for which they will deliver orders.
* 2. Orders are sent to all available drivers in an area.
* 3. Only orders for the selected restaurants will appear for the driver.
* 4. The first driver to accept an order removes it from the list of available orders. You need to implement an Azure Service Bus solution.
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.

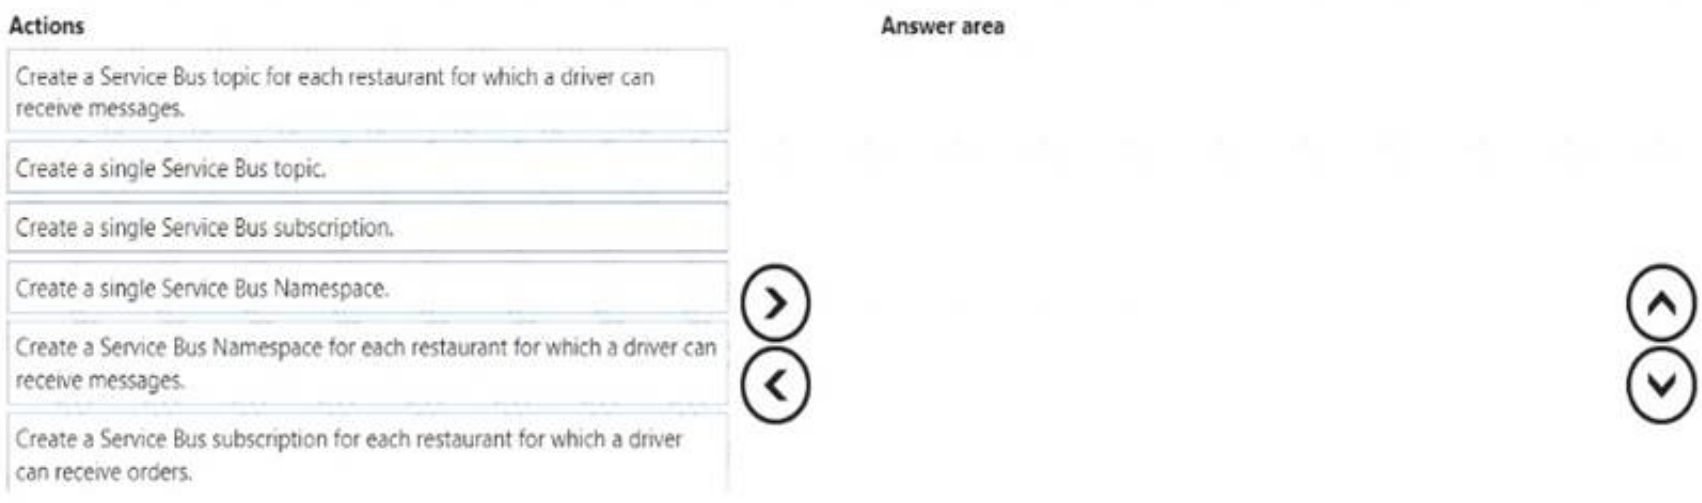

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Create a single Service Bus Namespace
To begin using Service Bus messaging entities in Azure, you must first create a namespace with a name that is unique across Azure. A namespace provides a scoping container for addressing Service Bus resources within your application.
Box 2: Create a Service Bus Topic for each restaurant for which a driver can receive messages. Create topics.
Box 3: Create a Service Bus subscription for each restaurant for which a driver can receive orders. Topics can have multiple, independent subscriptions.
References:

https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-messaging-overview

NEW QUESTION 44
- (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.	○	○
Edge nodes must be purged of all cache assets.	○	○
The compression type is supported.	○	○

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 49
- (Exam Topic 3)
You provide an Azure API Management managed web service lo 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 uses? 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 51
- (Exam Topic 3)
You are creating an app that uses Event Grid to connect with other services. Your app's event data will be sent to a serverless function that checks compliance. This function is maintained by your company.
You write a new event subscription at the scope of your resource. The event must be invalidated after 3 specific period of time. You need to configure Event Grid to ensure security.
What should you implement? To answer, select the appropriate options in [he answer area. NOTE: Each correct selection is worth one point

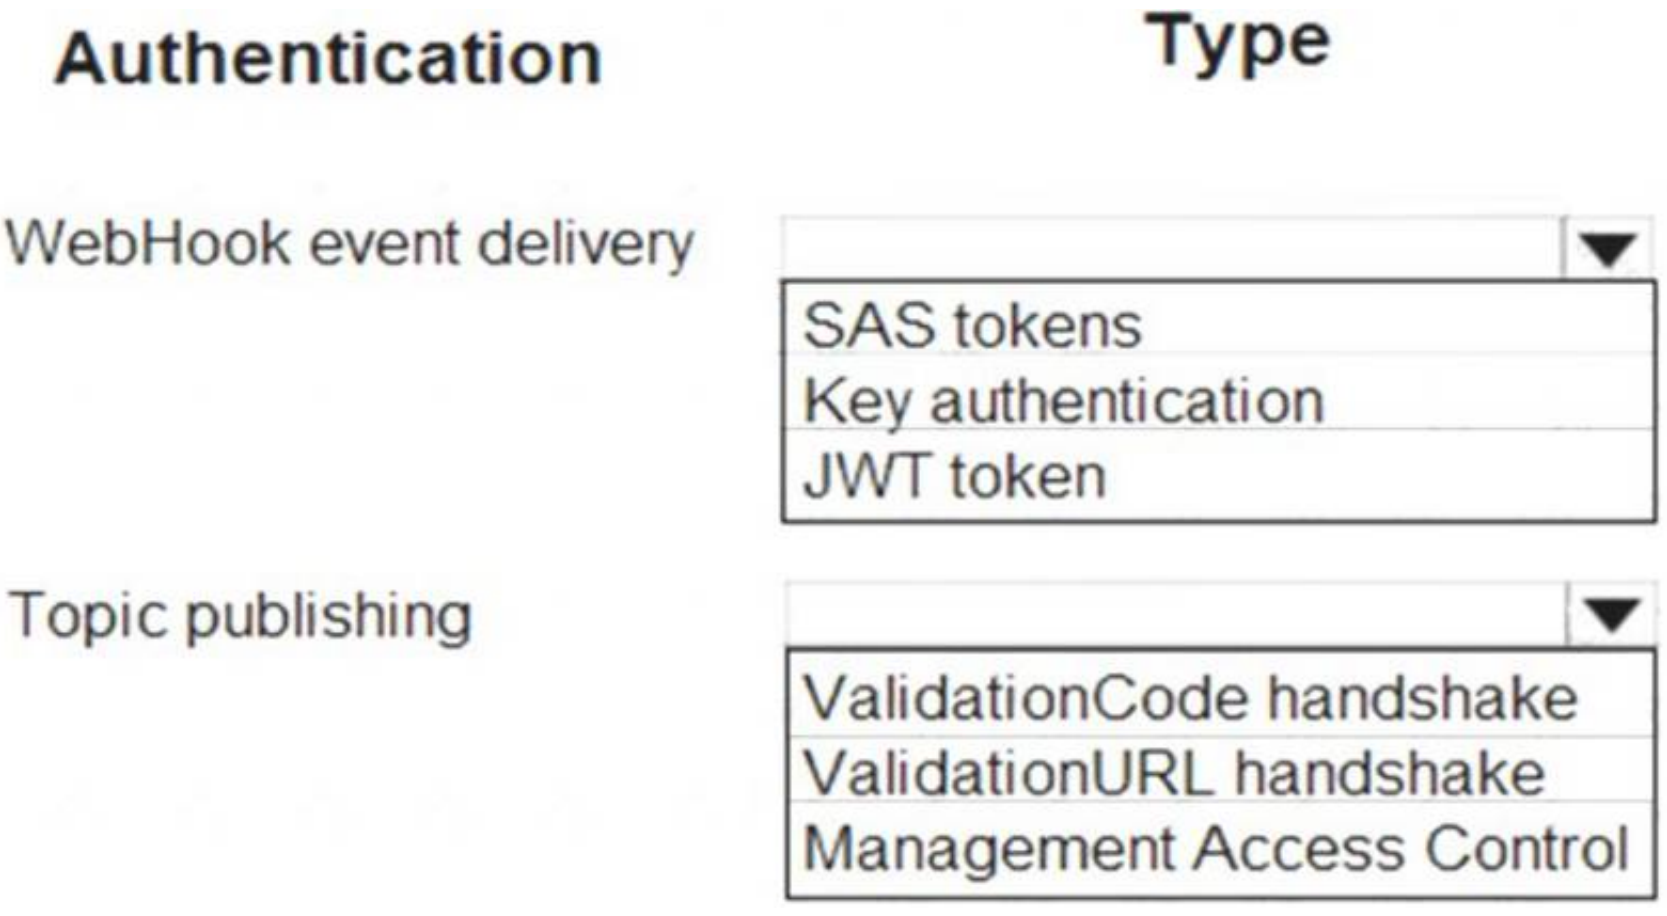

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: SAS tokens
Custom topics use either Shared Access Signature (SAS) or key authentication. Microsoft recommends SAS, but key authentication provides simple programming, and is compatible with many existing webhook publishers.
In this case we need the expiration time provided by SAS tokens. Box 2: ValidationCode handshake
Event Grid supports two ways of validating the subscription: ValidationCode handshake (programmatic) and ValidationURL handshake (manual).
If you control the source code for your endpoint, this method is recommended.

NEW QUESTION 56
- (Exam Topic 3)
A company is implementing a publish-subscribe (Pub/Sub) messaging component by using Azure Service Bus. You are developing the first subscription application.
In the Azure portal you see that messages are being sent to the subscription for each topic. You create and initialize a subscription client object by supplying the correct details, but the subscription application is still not consuming the messages.
You need to complete the source code of the subscription client What should you do?

A. await subscriptionClient.CloseAsync();
B. await subscriptionClient.AddRuleAsync(new RuleDescription(RuleDescription.DefaultRuleName, new TrueFilter()));
C. subscriptionClient.RegisterMessageHandler(ProcessMessagesAsync, messageHandlerOptions);
D. subscriptionClient = new SubscriptionClient(ServiceBusConnectionString, TopicName, SubscriptionName);

Answer: C

Explanation:
Using topic client, call RegisterMessageHandler which is used to receive messages continuously from the entity. It registers a message handler and begins a new thread to receive messages. This handler is waited on every time a new message is received by the receiver.
subscriptionClient.RegisterMessageHandler(ReceiveMessagesAsync, messageHandlerOptions); References:
https://www.c-sharpcorner.com/article/azure-service-bus-topic-and-subscription-pub-sub/

NEW QUESTION 60
- (Exam Topic 3)
You develop a gateway solution for a public facing news API. The news API back end is implemented as a RESTful service and uses an OpenAPI specification.
You need to ensure that you can access the news API by using an Azure API Management service instance. Which Azure PowerShell command should you run?

A. Import-AzureRmApiManagementApi –Context $ApiMgmtContext –SpecificationFormat "Swagger" -SpecificationPath $SwaggerPath –Path $Path
B. New-AzureRmApiManagementBackend -Context $ApiMgmtContext -Url $Url -Protocol http
C. New-AzureRmApiManagement –ResourceGroupName $ResourceGroup –Name $Name – Location $Location –Organization $Org –AdminEmail $AdminEmail
D. New-AzureRmApiManagementBackendProxy –Url $ApiUrl

Answer: D

Explanation:
New-AzureRmApiManagementBackendProxy creates a new Backend Proxy Object which can be piped when creating a new Backend entity.
Example: Create a Backend Proxy In-Memory Object
PS C:\>$secpassword = ConvertTo-SecureString "PlainTextPassword" -AsPlainText -Force
PS C:\>$proxyCreds = New-Object System.Management.Automation.PSCredential ("foo", $secpassword) PS C:\>$credential = New-AzureRmApiManagementBackendProxy -Url "http://12.168.1.1:8080"
-ProxyCredential $proxyCreds

PS C:\>$apimContext = New-AzureRmApiManagementContext -ResourceGroupName "Api-Default-WestUS" -ServiceName "contoso"
PS C:\>$backend = New-AzureRmApiManagementBackend -Context $apimContext -BackendId 123 -Url 'https://contoso.com/awesomeapi' -Protocol http -Title "first backend" -SkipCertificateChainValidation $true
-Proxy $credential -Description "backend with proxy server"
Creates a Backend Proxy Object and sets up Backend

NEW QUESTION 61
- (Exam Topic 3)
You are developing a mobile instant messaging app for a company. The mobile app must meet the following requirements:
• Support offline data sync.
• Update the latest messages during normal sync cycles. You need to implement Offline Data Sync.
Which two actions should you perform? Each conn I answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Retrieve records from Offline Data Sync on every call to the PullAsync method.
B. Retrieve records from Offline Data Sync using an Incremental Sync.
C. Push records to Offline Data Sync using an Incremental Sync.
D. Return the updatedAt column from the Mobile Service Backend and implement sorting by using the column.
E. Return the updatedAt column from the Mobile Service Backend and implement sorting by the message id.

Answer: BE

Explanation:
B: Incremental Sync: the first parameter to the pull operation is a query name that is used only on the client. If you use a non-null query name, the Azure Mobile SDK performs an incremental sync. Each time a pull operation returns a set of results, the latest updatedAt timestamp from that result set is stored in the SDK local system tables. Subsequent pull operations retrieve only records after that timestamp.
E (not D): To use incremental sync, your server must return meaningful updatedAt values and must also support sorting by this field. However, since the SDK adds its own sort on the updatedAt field, you cannot use a pull query that has its own orderBy clause.
References:
https://docs.microsoft.com/en-us/azure/app-service-mobile/app-service-mobile-offline-data-sync

NEW QUESTION 62
- (Exam Topic 3)
Fourth Coffee has an ASP.NET Core web app that runs in Docker. The app is mapped to the www.fourthcoffee.com domain.
Fourth Coffee is migrating this application to Azure.
You need to provision an App Service Web App to host this docker image and map the custom domain to the App Service web app.
A resource group named FourthCoffeePublicWebResourceGroup has been created in the WestUS region that contains an App Service Plan named AppServiceLinuxDockerPlan.
Which order should the CLI commands be used to develop the solution? To answer, move all of the Azure CLI command from the list of commands to the answer area and arrange them in the correct order.

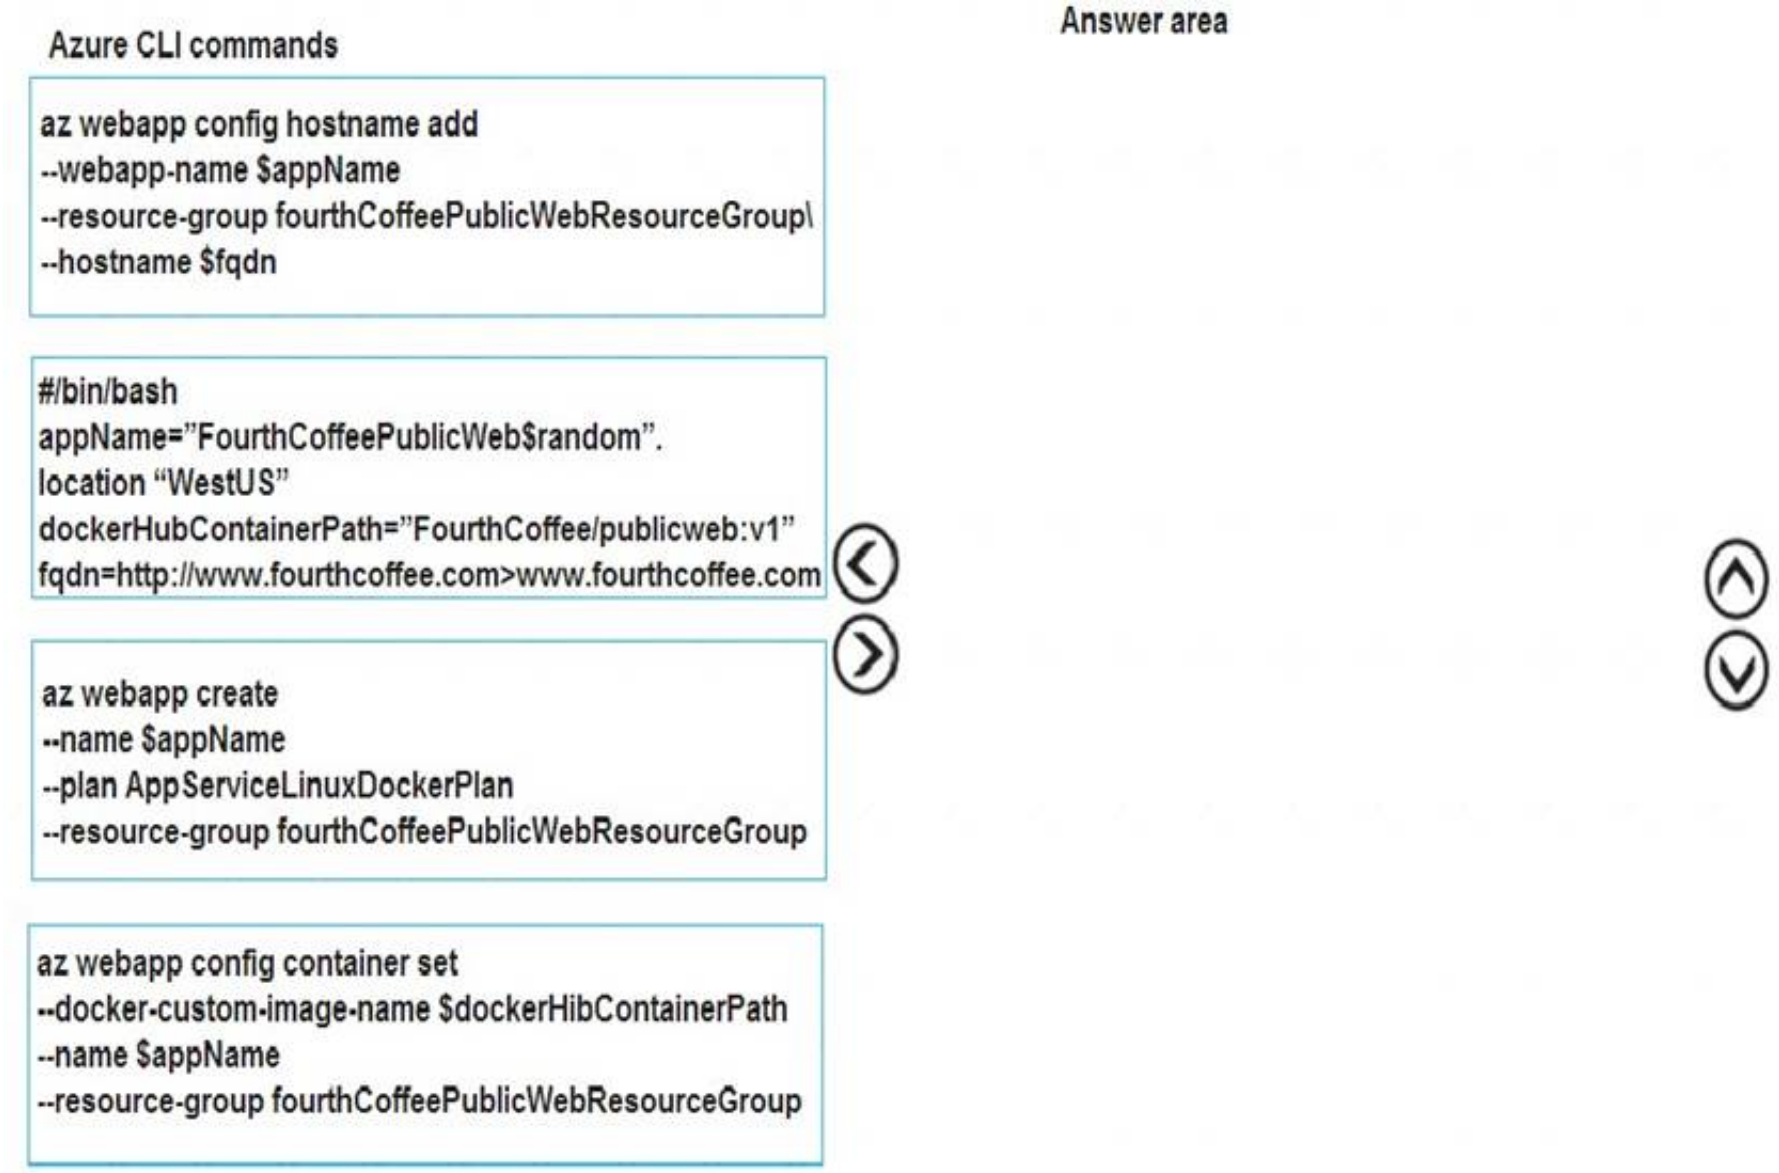

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Step 1: #bin/bash
The appName is used when the webapp-name is created in step 2. Step 2: az webapp config hostname add
The webapp-name is used when the webapp is created in step 3. Step 3: az webapp create
Create a web app. In the Cloud Shell, create a web app in the myAppServicePlan App Service plan with the az webapp create command.
Step : az webapp confing container set
In Create a web app, you specified an image on Docker Hub in the az webapp create command. This is good enough for a public image. To use a private image, you need to configure your Docker account ID and password in your Azure web app.
In the Cloud Shell, follow the az webapp create command with az webapp config container set.
References:
https://docs.microsoft.com/en-us/azure/app-service/containers/tutorial-custom-docker-image

NEW QUESTION 63
- (Exam Topic 3)
You are creating a CLI script that creates an Azure web app related services in Azure App Service. The web app uses the following variables:

Variable name	Value
$gitrepo	https://github.com/Contos/webapp
$webappname	Webapp1103

You need to automatically deploy code from GitHub to the newly created web app.
How should you complete the script? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

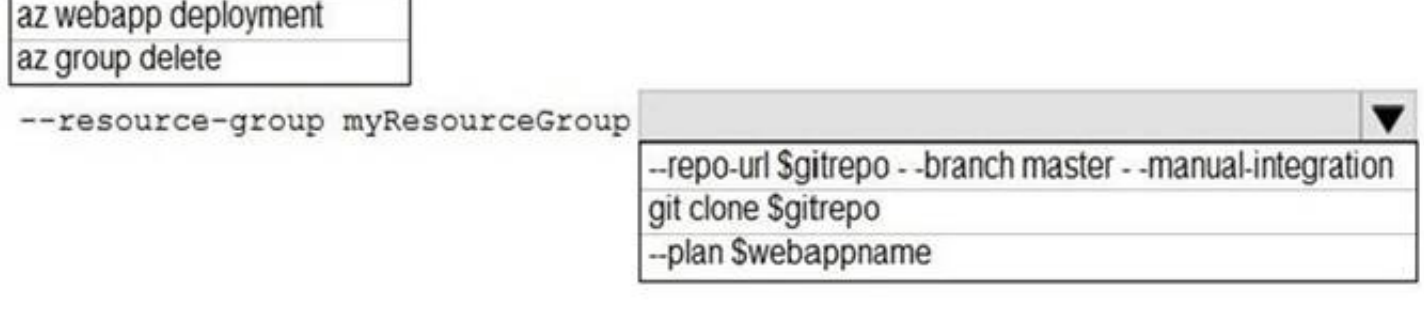

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: az appservice plan create
The azure group creates command successfully returns JSON result. Now we can use resource group to create a azure app service plan
Box 2: az webapp create Create a new web app..
Box 3: --plan $webappname
with the serviceplan we created in step 1. Box 4: az webapp deployment
Continuous Delivery with GitHub. Example:
az webapp deployment source config --name firstsamplewebsite1 --resource-group websites--repo-url $gitrepo
--branch master --git-token $token
Box 5: --repo-url $gitrepo --branch master --manual-integration Reference:
https://medium.com/@satish1v/devops-your-way-to-azure-web-apps-with-azure-cli-206ed4b3e9b1

NEW QUESTION 67
- (Exam Topic 3)
You develop a website. You plan to host the website in Azure. You expect the website to experience high traffic volumes after it is published. You must ensure that

the website remains available and responsive while minimizing cost. You need to deploy the website. What should you do?

A. Deploy the website to an App Service that uses the Shared service tie
B. Configure the App Service plan to automatically scale when the CPU load is high.
C. Deploy the website to a virtual machin
D. Configure the virtual machine to automatically scale when the CPU load is high.
E. Deploy the website to an App Service that uses the Standard service tie
F. Configure the App Service plan to automatically scale when the CPU load is high.
G. Deploy the website to a virtual machin
H. Configure a Scale Set to increase the virtual machine instance count when the CPU load

Answer: C

Explanation:
Windows Azure Web Sites (WAWS) offers 3 modes: Standard, Free, and Shared.
Standard mode carries an enterprise-grade SLA (Service Level Agreement) of 99.9% monthly, even for sites with just one instance.
Standard mode runs on dedicated instances, making it different from the other ways to buy Windows Azure Web Sites.

NEW QUESTION 70
- (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 are developing an Azure solution to collect point-of-sale fPOS) device data from 2,000 stores located throughout the world. A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data.
You must store the device data in Azure Blob storage. Device data must be correlated based on a device identifier. Additional stores are expected to open in the future.
You need to implement a solution to receive the device data.
Solution: Provision an Azure Event Hub. Configure the machine identifier as the partition key and enable capture.

A. Yes
B. No

Answer: A

Explanation:
References:
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-programming-guide

NEW QUESTION 72
- (Exam Topic 3)
Your company has several websites that use a company logo image. You use Azure Content Delivery Network (CDN) to store the static image.
You need to determine the correct process of how the CDN and the Point of Presence (POP) server will distribute the image and list the items in the correct order.
In which order do the actions occur? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

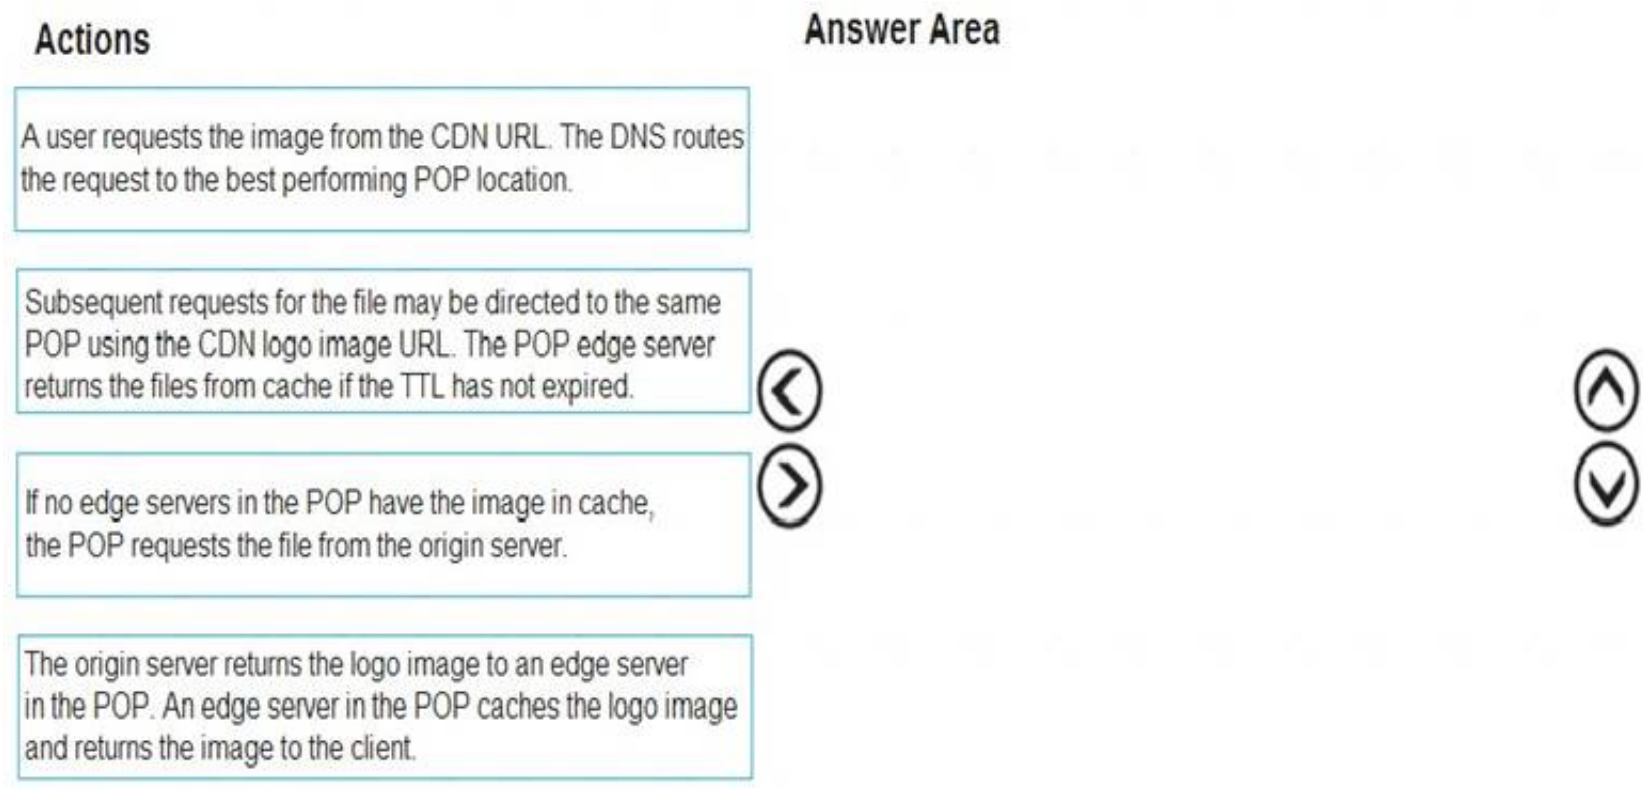

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: A user requests the image..
A user requests a file (also called an asset) by using a URL with a special domain name, such as <endpoint name>.azureedge.net. This name can be an endpoint hostname or a custom domain. The DNS routes the request to the best performing POP location, which is usually the POP that is geographically closest to the user.
Step 2: If no edge servers in the POP have the..
If no edge servers in the POP have the file in their cache, the POP requests the file from the origin server. The origin server can be an Azure Web App, Azure

Cloud Service, Azure Storage account, or any publicly accessible web server.
Step 3: The origin server returns the..
The origin server returns the file to an edge server in the POP.
An edge server in the POP caches the file and returns the file to the original requestor (Alice). The file remains cached on the edge server in the POP until the time-to-live (TTL) specified by its HTTP headers expires. If the origin server didn't specify a TTL, the default TTL is seven days.
Step 4: Subsequent requests for..
Additional users can then request the same file by using the same URL that the original user used, and can also be directed to the same POP.
If the TTL for the file hasn't expired, the POP edge server returns the file directly from the cache. This process results in a faster, more responsive user experience.
References:
https://docs.microsoft.com/en-us/azure/cdn/cdn-overview

NEW QUESTION 76
- (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 77
- (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 79
- (Exam Topic 3)
You have an application that provides weather forecasting data to external partners. You use Azure API Management to publish APIs.
You must change the behavior of the API to meet the following requirements:
• Support alternative input parameters.
• Remove formatting text from responses.
• Provide additional context to back-end services.
Which types of policies should you implement? To answer, drag the policy types to the correct scenarios. Each policy 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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 81
- (Exam Topic 3)
You are developing a software solution for an autonomous transportation system. The solution uses large data sets and Azure Batch processing to simulate navigation sets for entire fleets of vehicles.
You need to create compute nodes for the solution on Azure Batch. What should you do?

A. In Python, implement the class: TaskAddParameter
B. In Python, implement the class: JobAddParameter
C. In the Azure portal, create a Batch account
D. In a .NET method, call the method: BatchClient.PoolOperations.CreateJob

Answer: D

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.
Note:
Step 1: Create a pool of compute nodes. When you create a pool, you specify the number of compute nodes for the pool, their size, and the operating system. When each task in your job runs, it's assigned to execute on one of the nodes in your pool.
Step 2 : Create a job. A job manages a collection of tasks. You associate each job to a specific pool where that job's tasks will run.
Step 3: Add tasks to the job. Each task runs the application or script that you uploaded to process the data files it downloads from your Storage account. As each task completes, it can upload its output to Azure Storage.

NEW QUESTION 86
- (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.

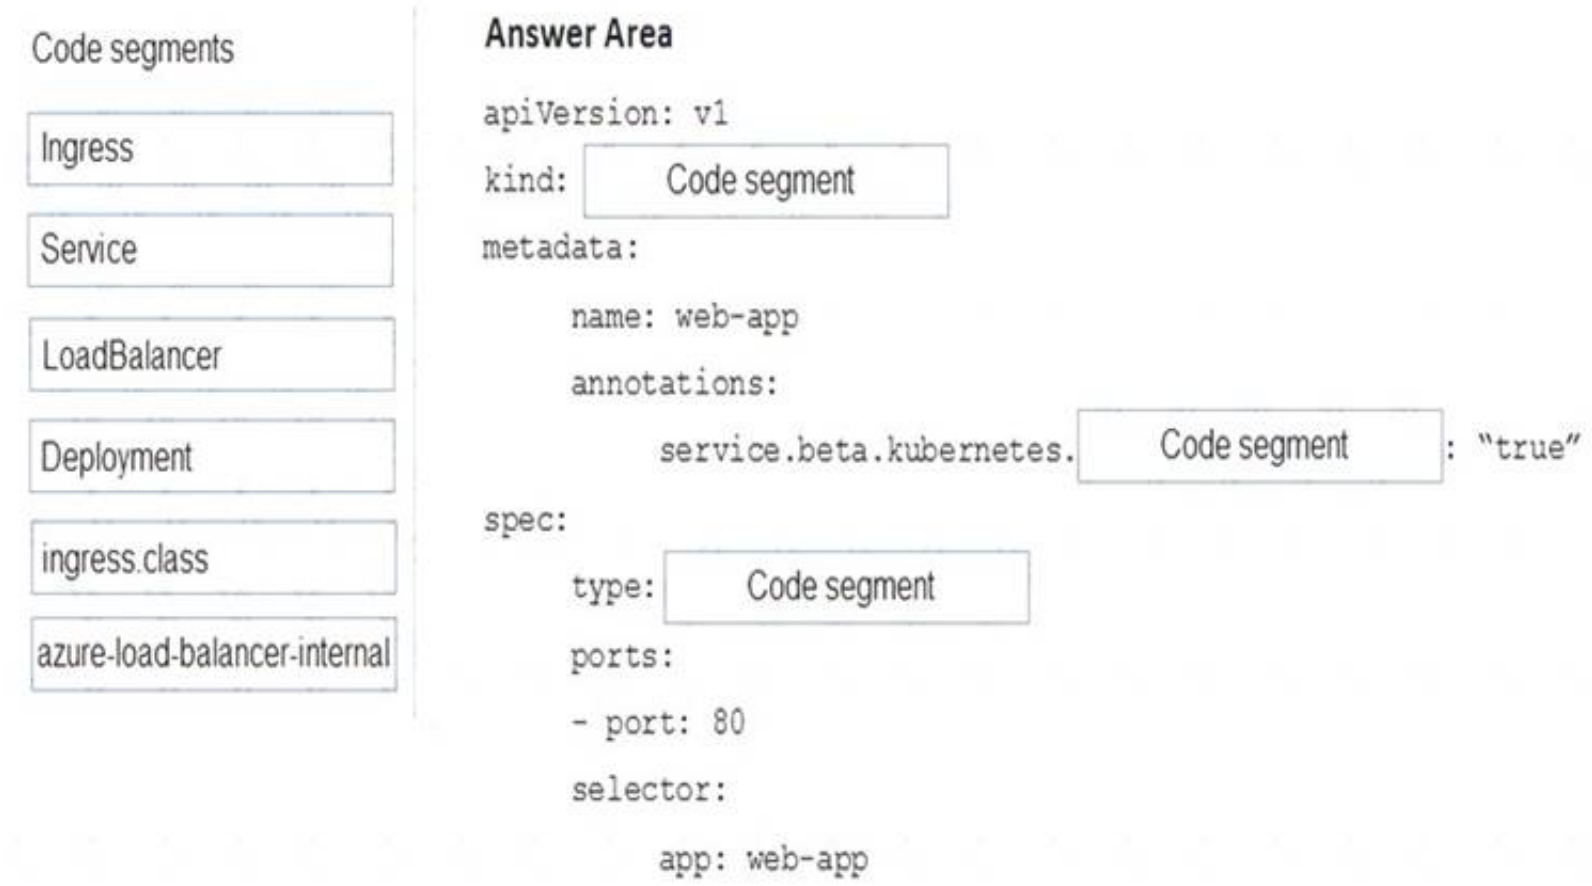

A. Mastered
B. Not Mastered

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 91
- (Exam Topic 3)
You have an Azure App Services Web App. Azure SQL Database instance. Azure Storage Account and an Azure Redis Cache instance in a resource group.
A developer must be able to publish code to the web app. You must grant the developer the Contribute role to the web app
You need to grant the role.
What two commands can you use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. New-AzureRmRoleAssignment
B. az role assignment create
C. az role definition create
D. New-AzureRmRoleDefinition

Answer: AB

Explanation:
References:
https://docs.microsoft.com/en-us/cli/azure/role/assignment?view=azure-cli-latest#az-role-assignment-create https://docs.microsoft.com/en-us/powershell/module/azurerm.resources/new-azurermroleassignment?view=azur

NEW QUESTION 94
- (Exam Topic 3)
A company uses Azure SQL Database to store data for an app. The data includes sensitive information.
You need to implement measures that allow only members of the managers group to see sensitive information. Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Include the managers group.
B. Exclude the managers group.
C. Exclude the administrators group.
D. Navigate to the following URL:

```
PUT https://management.azure.com/subscriptions/00000000-1111-2222-3333-444444444444
/resourceGroups/rg01/providers/Microsoft.Sql/servers/server01/databases/customers
/transparentDataEncryption/current?api-version=2014-04-01
```

E. Run the following Azure PowerShell command:

```
New-AzureRmSqlDatabaseDataMaskingRule -SchemaName "dbo" -TableName "customers" '
-ColumnName "ssn" -MaskingFunction "Default"
```

A. Option A
B. Option B
C. Option C
D. Option D
E. Option E

Answer: BE

Explanation:
Dynamic data masking helps prevent unauthorized access to sensitive data by enabling customers to designate how much of the sensitive data to reveal with minimal impact on the application layer.
SQL users excluded from masking - A set of SQL users or AAD identities that get unmasked data in the SQL query results.
Note: The New-AzureRmSqlDatabaseDataMaskingRule cmdlet creates a data masking rule for an Azure SQL database.
References:
https://docs.microsoft.com/en-us/powershell/module/azurerm.sql/new-azurermsqldatabasedatamaskingrule?view

NEW QUESTION 95
- (Exam Topic 3)

A company backs up all manufacturing data to Azure Blob Storage. Admins move blobs from hot storage to archive tier storage every month.
You must automatically move blocks to Archive tier after they have not been accessed for 180 days. The path for any item that is not archived must be placed in an existing queue. This operation must be performed automatically once a month. You set the value of TierAgeInDays to 180.
How should you configure the Logic App? To answer, drag the appropriate triggers or action blocks to the correct trigger or action slots. Each trigger or action block 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.

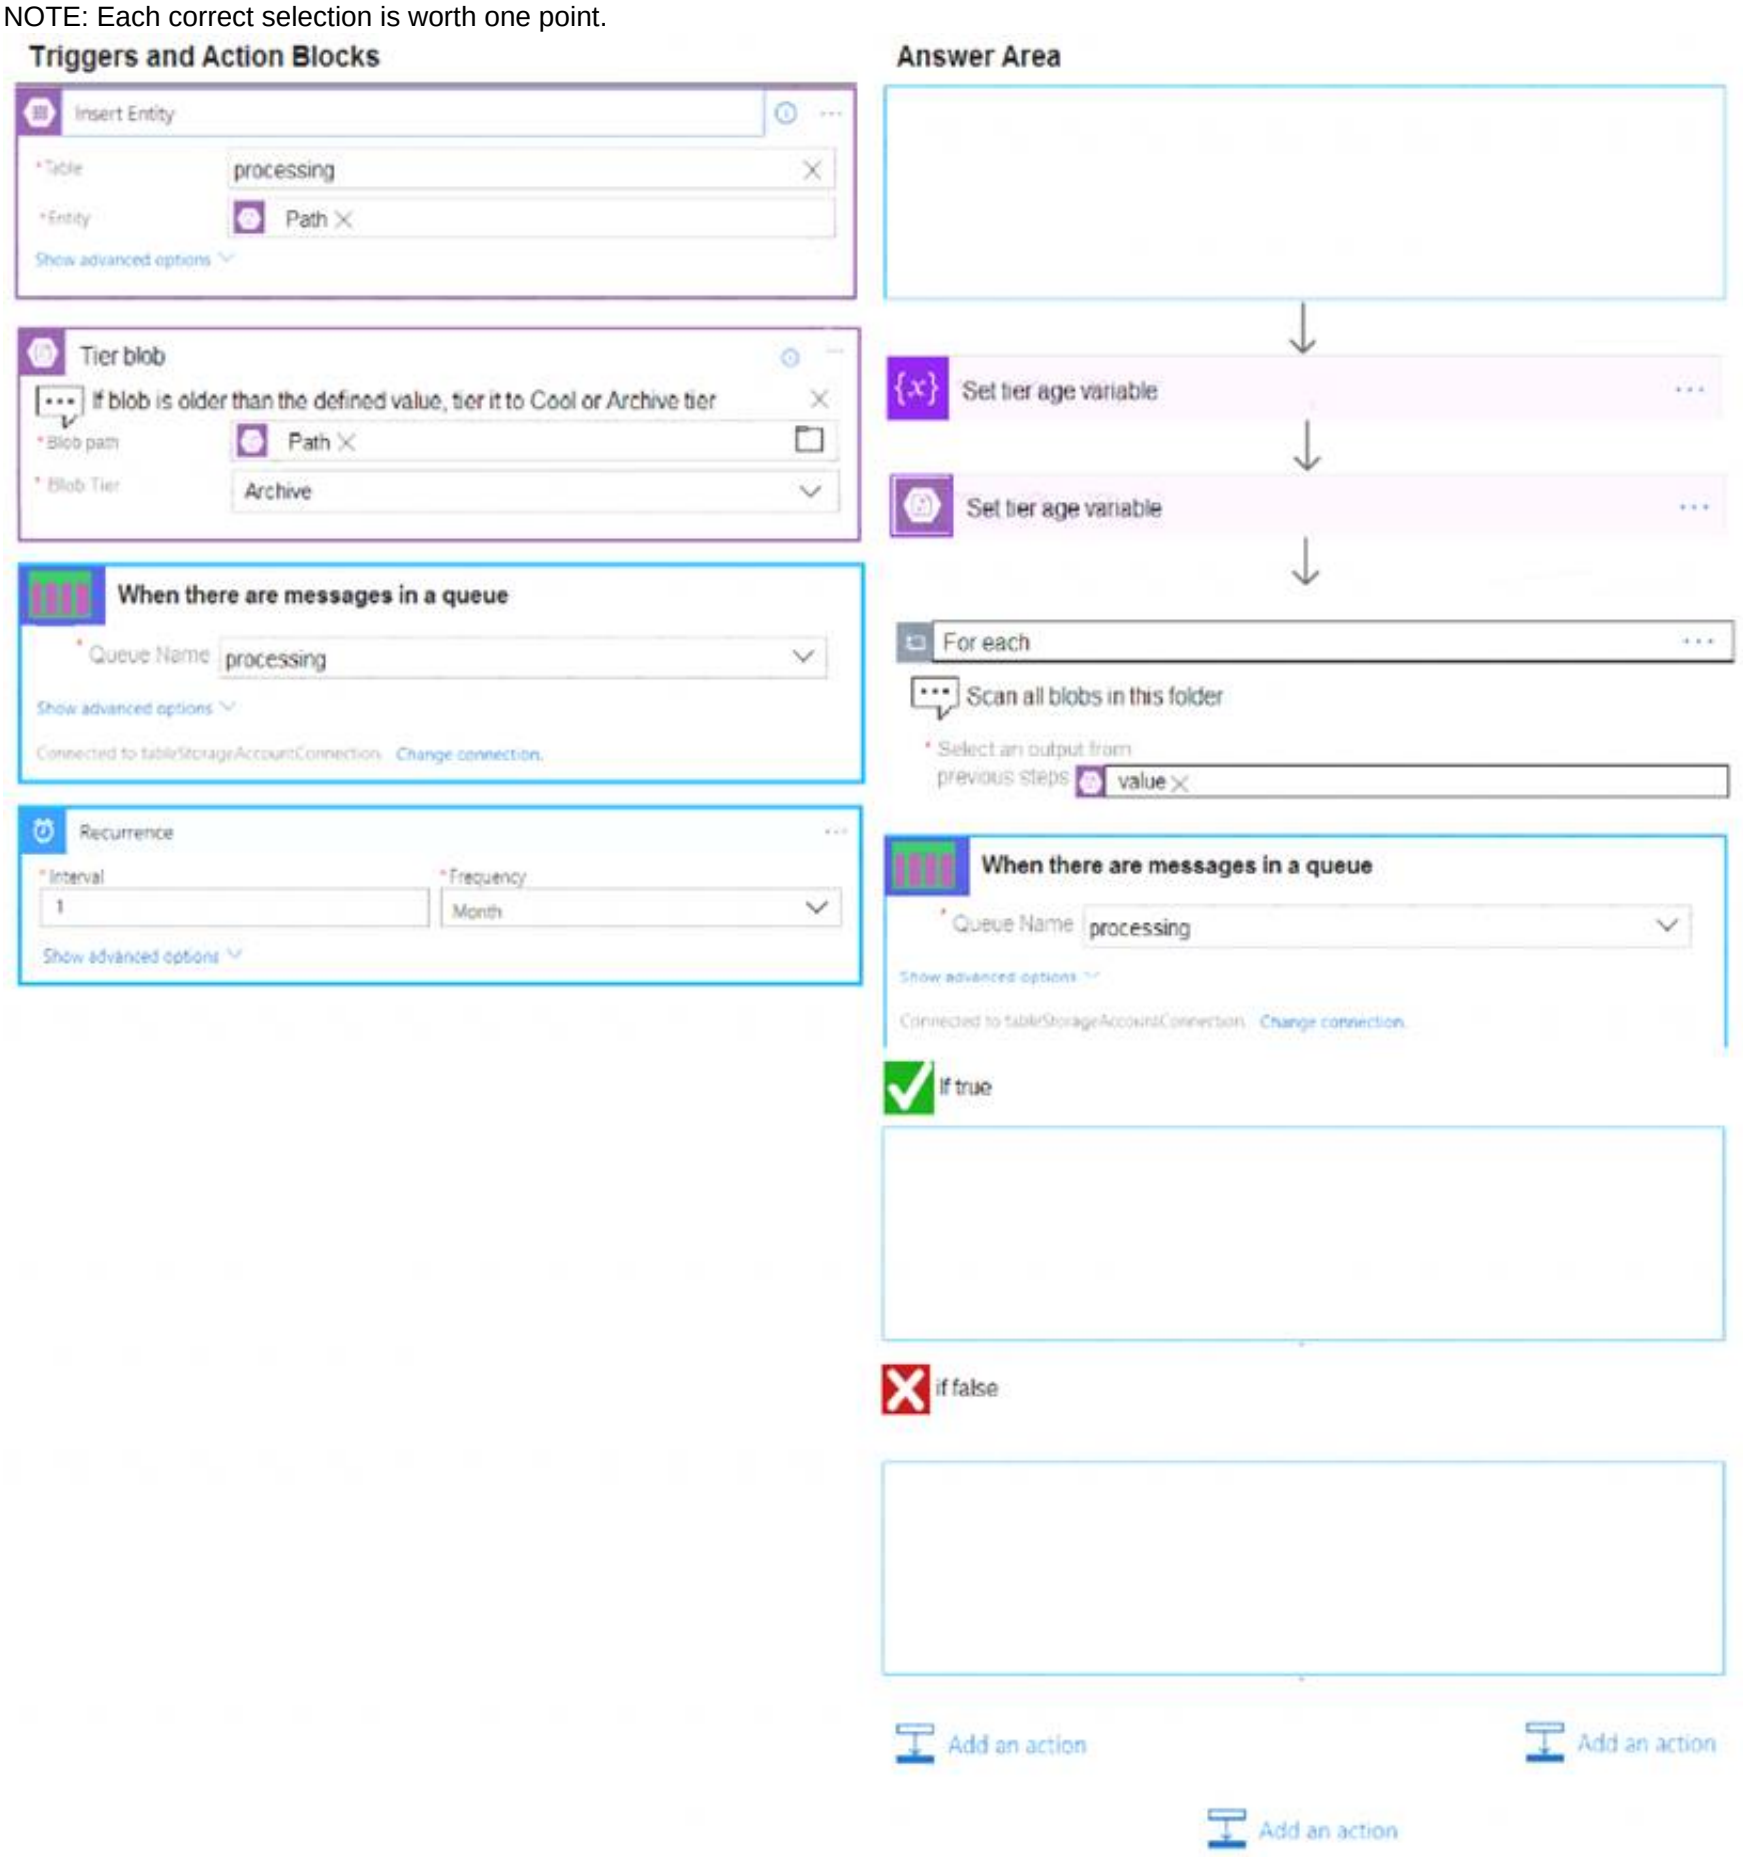

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Recurrence Box 2: Insert Entity
Box 3 (if true): Tier Blob Box 4: (if false):
Leave blank. References:
https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-perform-data-operations

NEW QUESTION 97
- (Exam Topic 3)
You are working for a company that designs mobile applications. They maintain a server where player records are assigned to their different games. The tracking system is new and in development.
The application uses Entity Framework to connect to an Azure Database. The database holds a Player table and Game table.
When adding a player, the code should insert a new player record, and add a relationship between an existing game record and the new player record.
The application will call CreatePlayerWithGame with the correct gameIdand the playerId to start the process. (Line numbers are included for reference only.)

```
01. namespace ContosoCradt
02. {
03.   public class PlayerDbContext : DbContext
04.   {
05.  public PlayerDbContext() : base ("name-dBConnString") { }
06.  public DbSet<Player> Players { get ; set ; }
07.  public DbSet<Game> Games { get ; set ]
08.  protected override void DmModelCreating{DBModelBuilder modelBuilder}
09.  {
10.   modelBuilder.Entity<Player>().MesMany(x => x.Games). WithMany (x => x Players);
11.  }
12.  }
13.   internal series class dbConfiguration : DbMigrationConfiguration<PlayerDbContext>
14.  {
15.  public dbConfiguration() . {AutomaticMigrationsEnabled = true ; }
16.  {
17.    public class mp
18.    {
19.  public void CreatePlayerWithGame(int playerId, int gameId) => AddPlayer(playerId, GetGame[gameId));
20.  public game GetGame(int gameId)
21.  {
22.    using (var db = new PlayerDbContext())
23.  {
24.     return db.Games.FirstOrDefault(x => x.GameId == gameId);
25.  }
26.   }
27.   public Player AddPlayer (int playerId, Game game)
28.   {
29.  using (var db = new PlayerDbContext())
30.  {
31.    var player = new Player
32.     {
33.      PlayerId = playerId,
34.      Games = new List <Game> {game },
35.  };
36.  db.Players.Add(player);
37.  db.SaveChanges();
38.  return player;
39.    }
40. }
41. public class Player
42. {
43.  public int PlayerId { get ; set; }
44.  public string PlayerName { get ; set; }
45.  public virtual List<Game> Games { get ; set; }
46.  }
47.  public class Game
48.  {
49.  public int GameIs { get ; set }
50.  public string Title { get ; set; }
51.   public string Platform { get ; set; ]
52.   public virtual List<Player> Players { get ; set; }
53.      }
54.   }
```

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 successfully insert a player record.	○	○
The code has a bug and will insert an additional copy of the Game record with a new Id.	○	○
The code has a bug and will insert the wrong gameId value.	○	○
There is a valid many-to-many relationship between Players and Games.	○	○

A. Mastered

B. Not Mastered

Answer: A

Explanation:
Many-to-many relationships without an entity class to represent the join table are not yet supported. However, you can represent a many-to-many relationship by including an entity class for the join table and mapping two separate one-to-many relationships.
protected override void OnModelCreating(ModelBuilder modelBuilder)
{
modelBuilder.Entity<PostTag>() HasKey(t => new { t.PostId, t.TagId }); modelBuilder.Entity<PostTag>() HasOne(pt => pt.Post)
WithMany(p => p.PostTags) HasForeignKey(pt => pt.PostId); modelBuilder.Entity<PostTag>() HasOne(pt => pt.Tag) WithMany(t => t.PostTags) HasForeignKey(pt => pt.TagId);
}
}

NEW QUESTION 102
- (Exam Topic 3)
You manage several existing Logic Apps.
You need to change definitions, add new logic, and optimize these apps on a regular basis.
What should you use? To answer, drag the appropriate tools to the correct functionalities. Each tool 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

Tools	Functionality	Tool
Logic Apps Designer	Edit B2B workflows	
Code View Editor	Edit definitions in JSON	
Enterprise Integration Pack	Visually and functionality	

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Enterprise Integration Pack
After you create an integration account that has partners and agreements, you are ready to create a business to business (B2B) workflow for your logic app with the Enterprise Integration Pack.
Box 2: Code View Editor
To work with logic app definitions in JSON, open the Code View editor when working in the Azure portal or in Visual Studio, or copy the definition into any editor that you want.
Box 3: Logical Apps Designer
You can build your logic apps visually with the Logic Apps Designer, which is available in the Azure portal through your browser and in Visual Studio.
References:
https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-enterprise-integration-b2b https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-overview

NEW QUESTION 105
- (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.

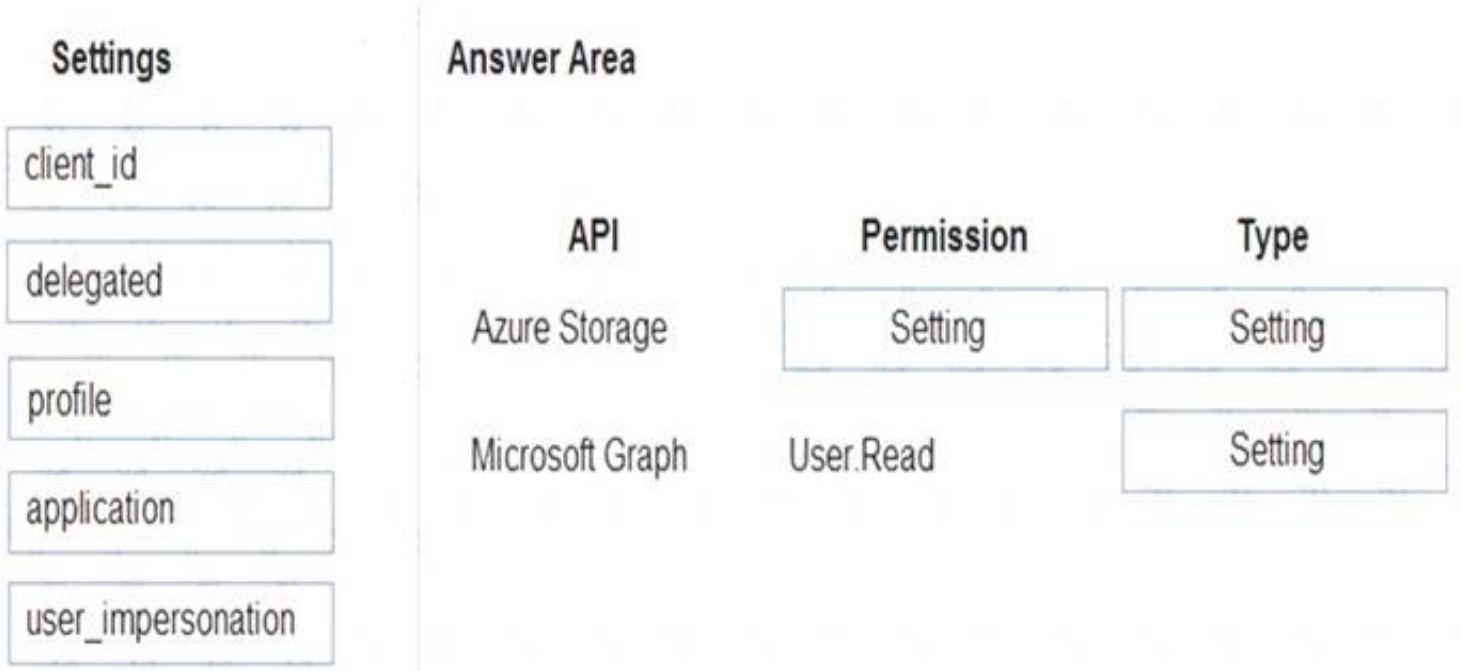

A. Mastered
B. Not Mastered

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 106
- (Exam Topic 3)
You develop an Azure web app. You monitor performance of the web app by using Application Insights. You need to ensure the cost for Application Insights does not exceed a preset budget. What should you do?

A. Implement ingestion sampling using the Azure portal.
B. Set a daily cap for the Application Insights instance.
C. Implement adaptive sampling using the Azure portal.
D. Implement adaptive sampling using the Application Insights SDK.
E. Implement ingestion sampling using the Application Insights SDK.

Answer: D

Explanation:
Sampling is an effective way to reduce charges and stay within your monthly quota.
You can set sampling manually, either in the portal on the Usage and estimated costs page; or in the ASP.NET SDK in the .config file; or in the Java SDK in the ApplicationInsights.xml file, to also reduce the network traffic.
Adaptive sampling is the default for the ASP.NET SDK. Adaptive sampling automatically adjusts to the volume of telemetry that your app sends. It operates automatically in the SDK in your web app so that telemetry traffic on the network is reduced.
References:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/sampling

NEW QUESTION 110
- (Exam Topic 3)
You are working for Contoso, Ltd.
You define an API Policy object by using the following XML markup:

```
<set-variable name= "bodySize" value="@(context.Request.Headers["Content-Length"] [0])"/>
<choose>
 <when condition= "@(int.Parse(context.Variables.GetValueOrDefault<string> ("bodySize"))<512000)">
</when>
<otherwise>
     <rewrite-uri template= "/put"/>
     <set-backend-service base-url= "http://contoso.com/api/9.1/"/>
</otherwise>
</choose>
```

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 Area

Statement	Yes	No
The XML segment belongs in the <inbound> section of the policy.	○	○
If the body size is >256k, an error will occur.	○	○
If the request is http://contoso.com/api/9.2/, the policy will retain the higher version.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Yes
Use the set-backend-service policy to redirect an incoming request to a different backend than the one specified in the API settings for that operation. Syntax: <set-backend-service base-url="base URL of the backend service" />
Box 2: No
The condition is on 512k, not on 256k. Box 3: No
The set-backend-service policy changes the backend service base URL of the incoming request to the one specified in the policy.
Reference:
https://docs.microsoft.com/en-us/azure/api-management/api-management-transformation-policies

NEW QUESTION 114
- (Exam Topic 3)
You are developing a ticket reservation system for an airline.
The storage solution for the application must meet the following requirements:

- Ensure at least 99.99% availability and provide low latency.
- Accept reservations event when localized network outages or other unforeseen failures occur.
- Process reservations in the exact sequence as reservations are submitted to minimize overbooking or selling the same seat to multiple travelers.
- Allow simultaneous and out-of-order reservations with a maximum five-second tolerance window. You provision a resource group named airlineResourceGroup in the Azure South-Central US region. You need to provision a SQL SPI Cosmos DB account to support the app.

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

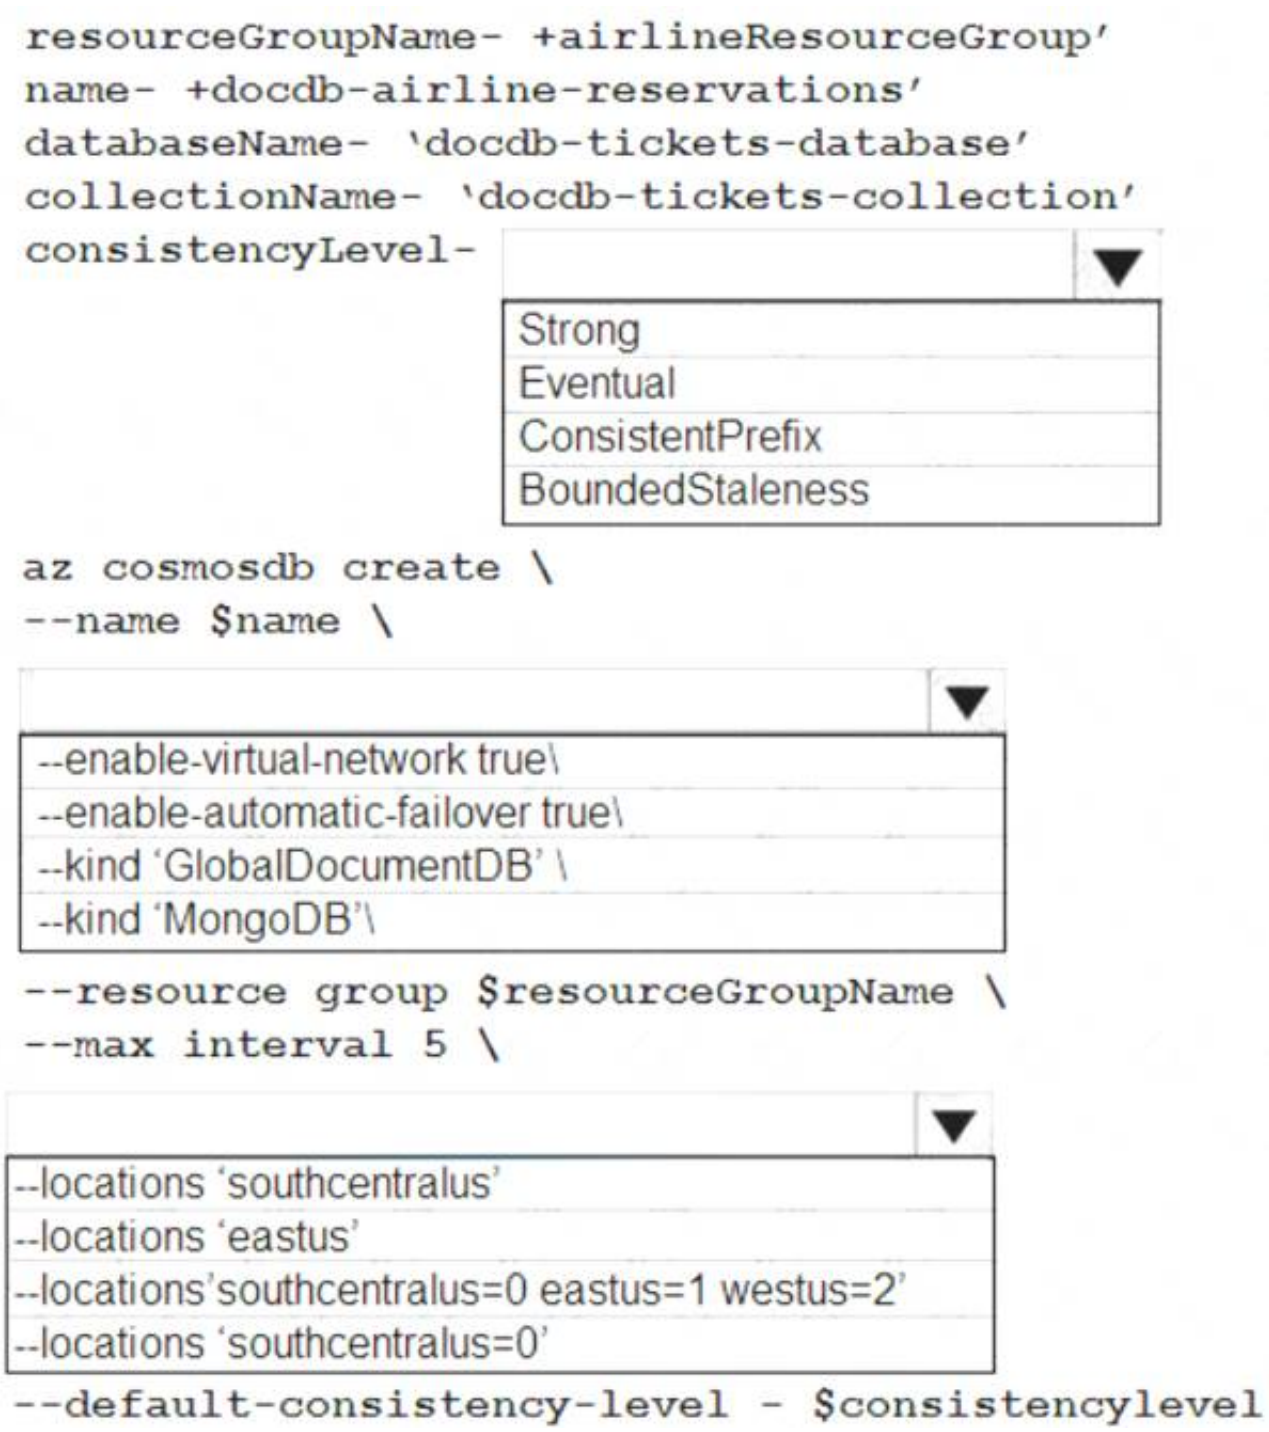

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: BoundedStaleness
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. In other words, 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 Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/cosmos-db/manage-with-cli.md

NEW QUESTION 116
- (Exam Topic 3)
You are developing an application. You have an Azure user account that has access to two subscriptions. You need to retrieve a storage account key secret from Azure Key Vault.
In which order should you arrange the PowerShell commands to develop the solution? To answer, move all commands from the list of commands to the answer area and arrange them in the correct order.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Get-AzSubscription
If you have multiple subscriptions, you might have to specify the one that was used to create your key vault. Enter the following to see the subscriptions for your account:
Get-AzSubscription
Step 2: Set-AzContext -SubscriptionId
To specify the subscription that's associated with the key vault you'll be logging, enter: Set-AzContext -SubscriptionId <subscriptionID>
Step 3: Get-AzStorageAccountKey You must get that storage account key.
Step 4: $secretvalue = ConvertTo-SecureString <storageAccountKey> -AsPlainText -Force
Set-AzKeyVaultSecret -VaultName <vaultName> -Name <secretName> -SecretValue $secretvalue After retrieving your secret (in this case, your storage account key), you must convert that key to a secure
string, and then create a secret with that value in your key vault.
Step 5: Get-AzKeyVaultSecret
Next, get the URI for the secret you created. You'll need this URI in a later step to call the key vault and retrieve your secret. Run the following PowerShell command and make note of the ID value, which is the secret's URI:
Get-AzKeyVaultSecret –VaultName <vaultName> Reference:
https://docs.microsoft.com/bs-latn-ba/Azure/key-vault/key-vault-key-rotation-log-monitoring

NEW QUESTION 118
- (Exam Topic 3)
You have a web app named MainApp. You are developing a triggered App Service background task by using the WebJobs SDK. This task automatically invokes a function code whenever any new data is received in a queue.
You need to configure the services.
Which service should you use for each scenario? To answer, drag the appropriate services to the correct scenarios. Each service 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.

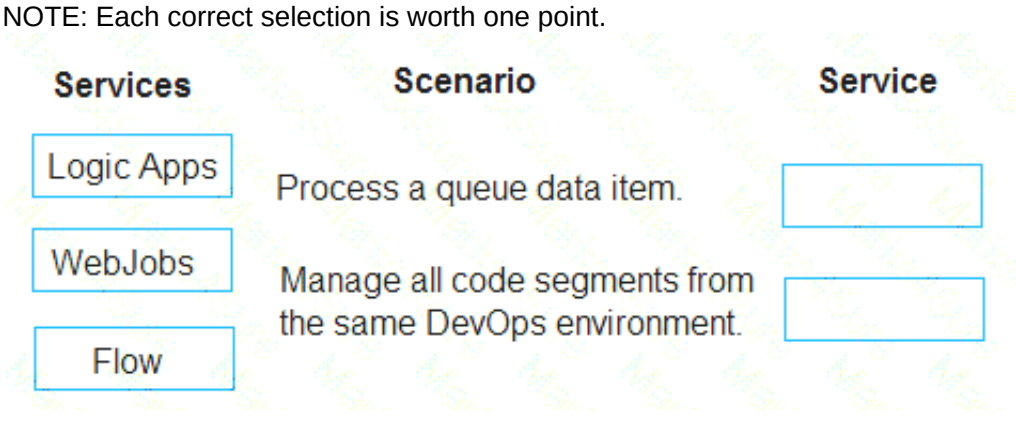

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: WebJobs
A WebJob is a simple way to set up a background job, which can process continuously or on a schedule. WebJobs differ from a cloud service as it gives you get less fine-grained control over your processing environment, making it a more true PaaS service.
Box 2: Flow

NEW QUESTION 120
- (Exam Topic 3)
You are developing Azure WebJobs.
You need to recommend a WebJob type for each scenario.
Which WebJob type should you recommend? To answer, drag the appropriate WebJob types to the correct scenarios. Each WebJob 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.

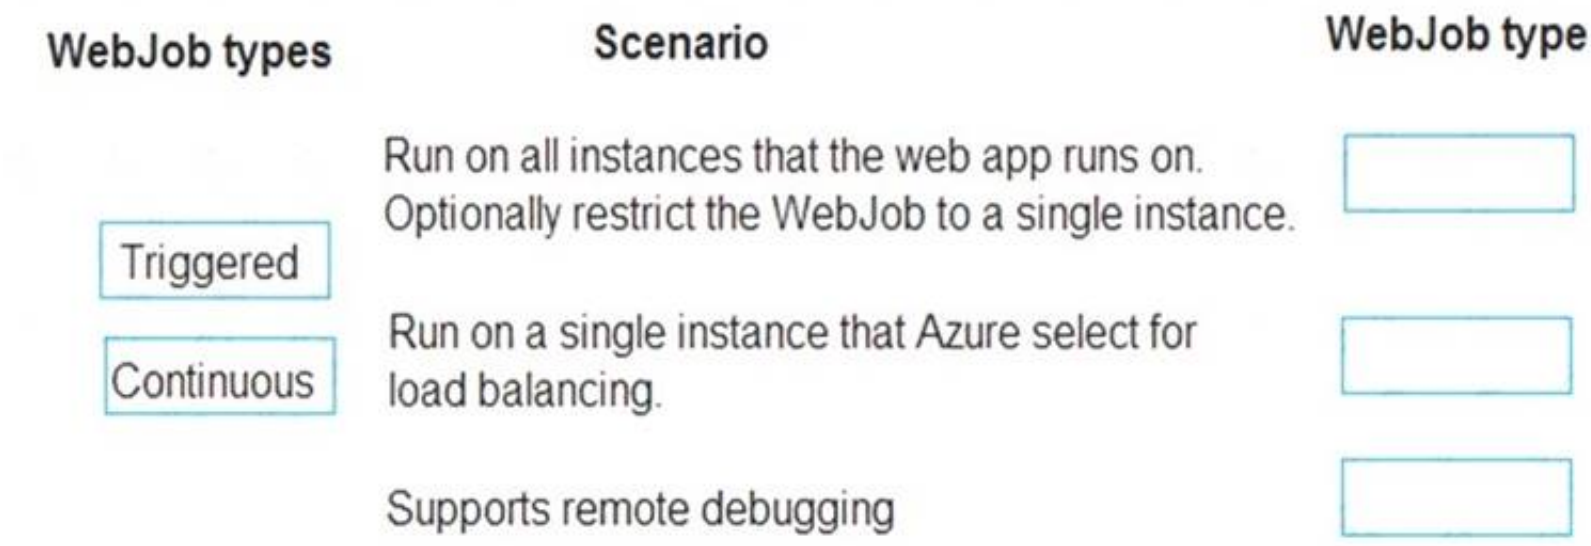

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Continuous
Continuous runs on all instances that the web app runs on. You can optionally restrict the WebJob to a single instance.
Box 2: Triggered
Triggered runs on a single instance that Azure selects for load balancing. Box 3: Continuous
Continuous supports remote debugging. Note:
The following table describes the differences between continuous and triggered WebJobs.

Continuous	Triggered
Starts immediately when the WebJob is created. To keep the job from ending, the program or script typically does its work inside an endless loop. If the job does end, you can restart it.	Starts only when triggered manually or on a schedule.
Runs on all instances that the web app runs on. You can optionally restrict the WebJob to a single instance.	Runs on a single instance that Azure selects for load balancing.
Supports remote debugging.	Doesn't support remote debugging.

References:
https://docs.microsoft.com/en-us/azure/app-service/web-sites-create-web-jobs

NEW QUESTION 125
- (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.

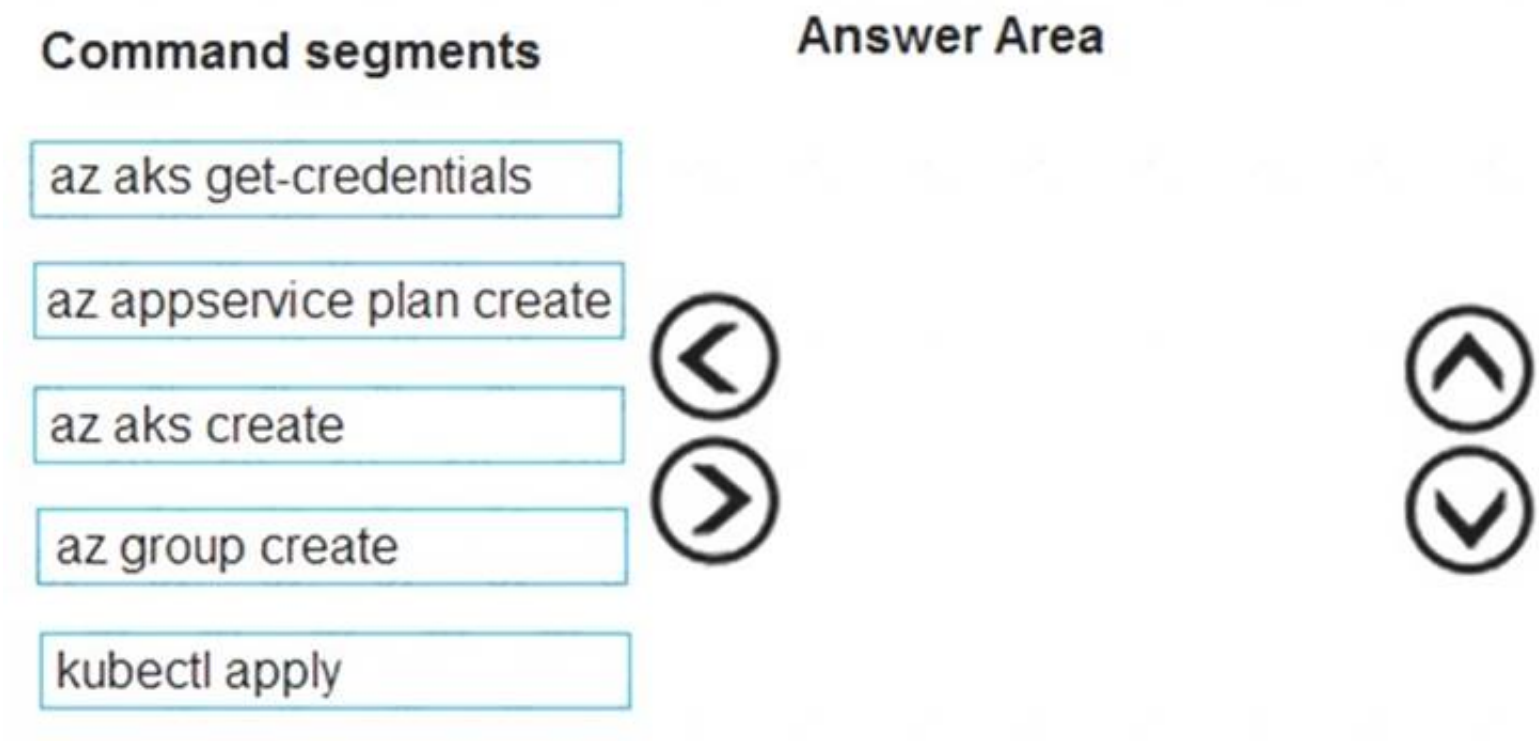

A. Mastered
B. Not Mastered

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 129
- (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 131
- (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.

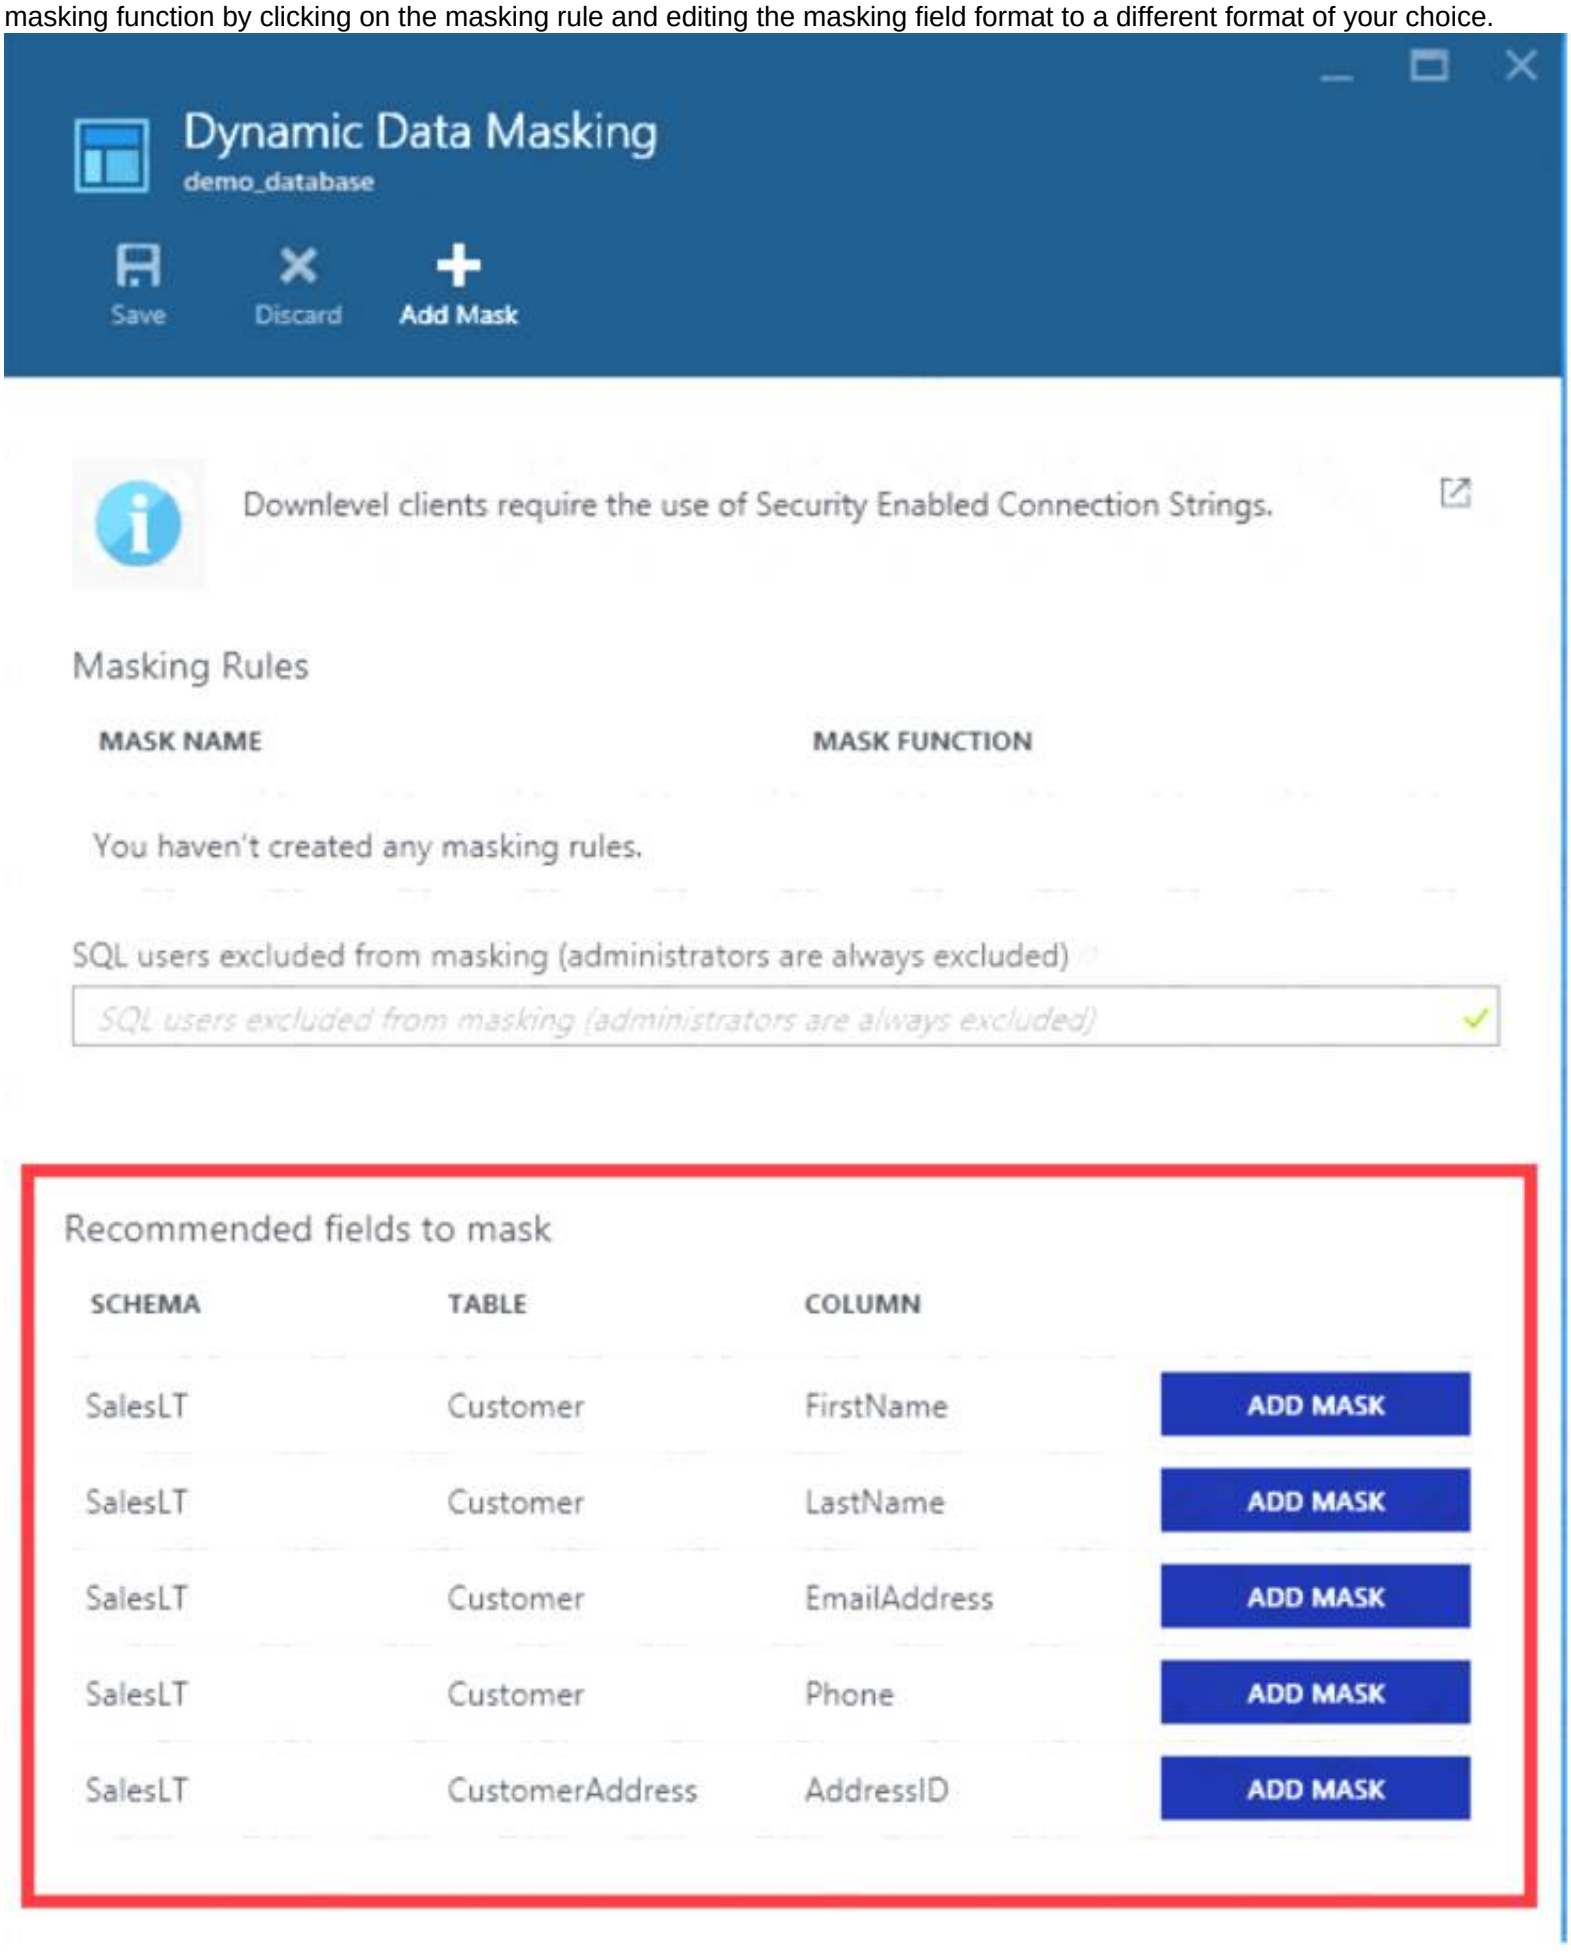

References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dynamic-data-masking-get-started-portal

NEW QUESTION 135
- (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 propert
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 139
- (Exam Topic 3)
You are developing an internal website for employees to view sensitive data. The website uses Azure Active Directory (AAD) for authentication. You need to implement multifactor authentication for the website.
What should you do? Each correct answer presents part of the solution. NOTE; Each correct selection is worth one point.

A. In Azure AD, create a new conditional access policy.
B. In Azure AD, enable application proxy.
C. Configure the website to use Azure AD B2C.
D. In Azure AD conditional access, enable the baseline policy.
E. Upgrade to Azure AD Premium.

Answer: AE

Explanation:
References:
https://docs.microsoft.com/en-us/azure/active-directory/authentication/howto-mfa-getstarted

NEW QUESTION 142
.....

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1
- (Exam Topic 1)
You need to resolve the Shipping web site error.
How should you configre the Azure Table Storage service? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

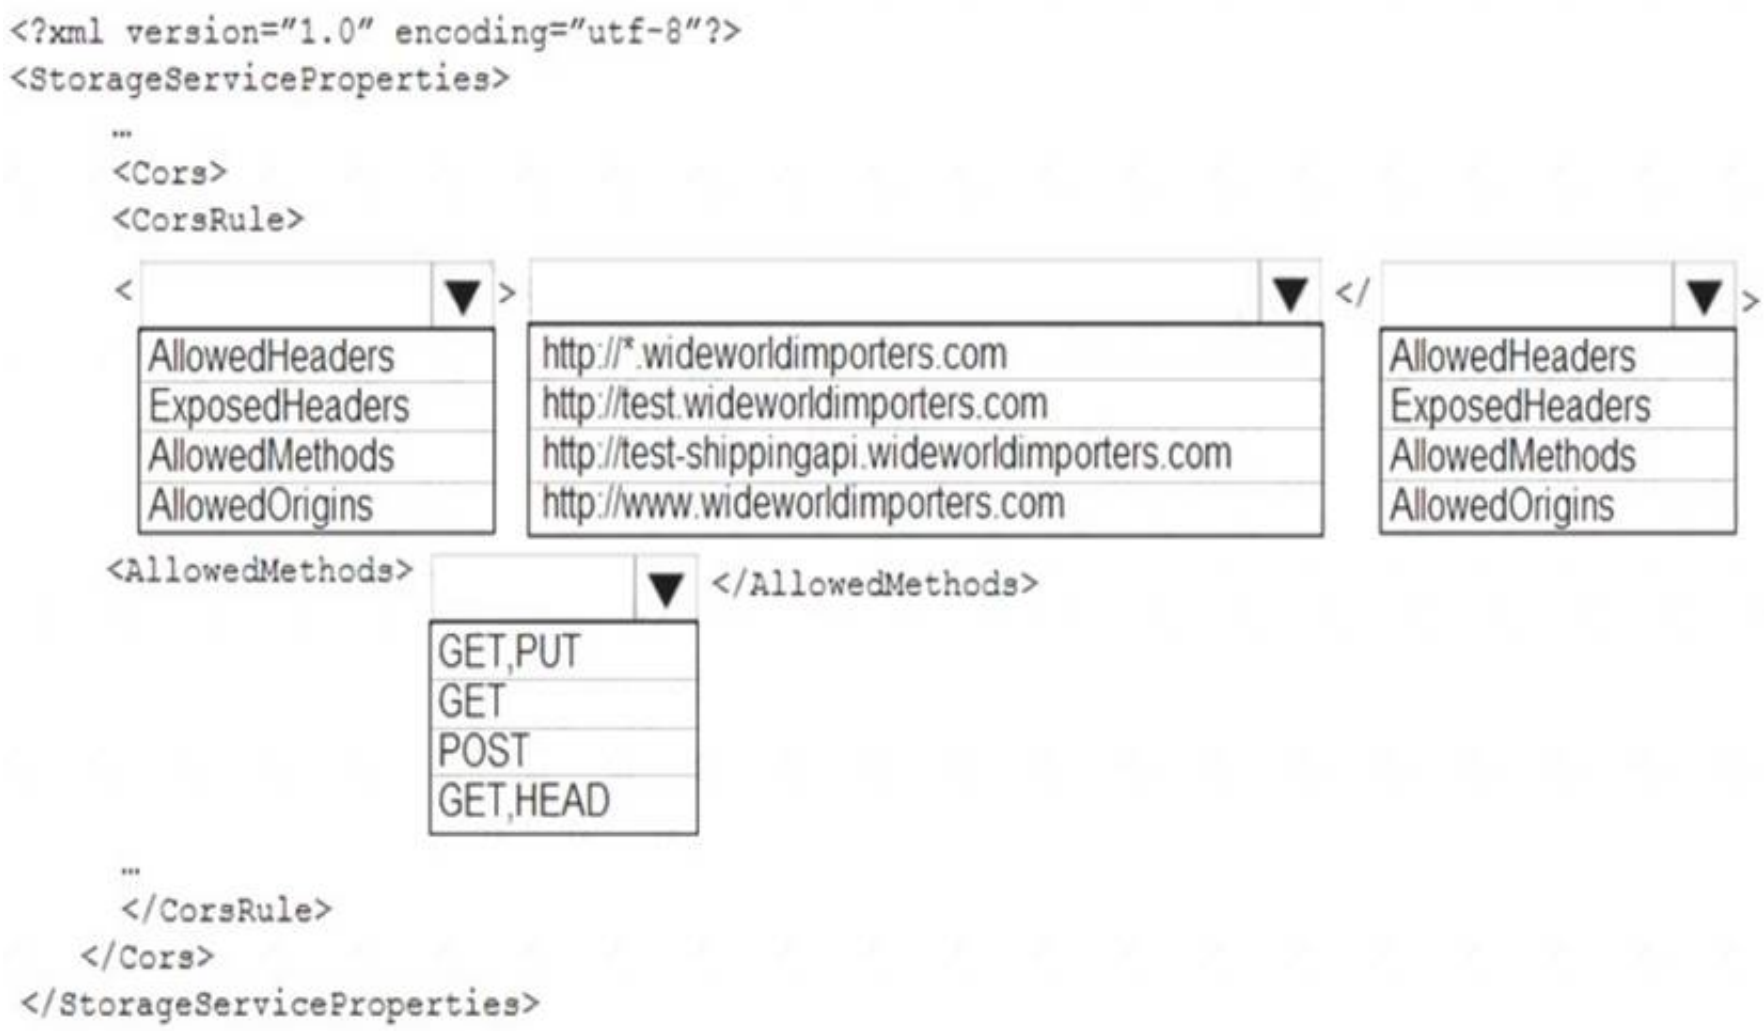

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: AllowedOrigins
A CORS request will fail if Access-Control-Allow-Origin is missing. Scenario:
The following error message displays while you are testing the website:

```
Failed to load http://test-shippingapi.wideworldimporters.com/: No 'Access-Control-Allow-Origin'
header is present on the requested resource. Origin 'http://testwideworldimporters.com/' is
therefore not allowed access.
```

Box 2: http://test-shippingapi.wideworldimporters.com Syntax: Access-Control-Allow-Origin: *
Access-Control-Allow-Origin: <origin> Access-Control-Allow-Origin: null
<origin> Specifies an origin. Only a single origin can be specified. Box 3: AllowedOrigins
Box 4: POST
The only allowed methods are GET, HEAD, and POST. In this case POST is used. "<Corsrule>" "allowedmethods" Failed to load no "Access-control-Origin" header is present References:
https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Access-Control-Allow-Origin

NEW QUESTION 2
- (Exam Topic 1)
You need to support the message processing for the ocean transport workflow.
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.

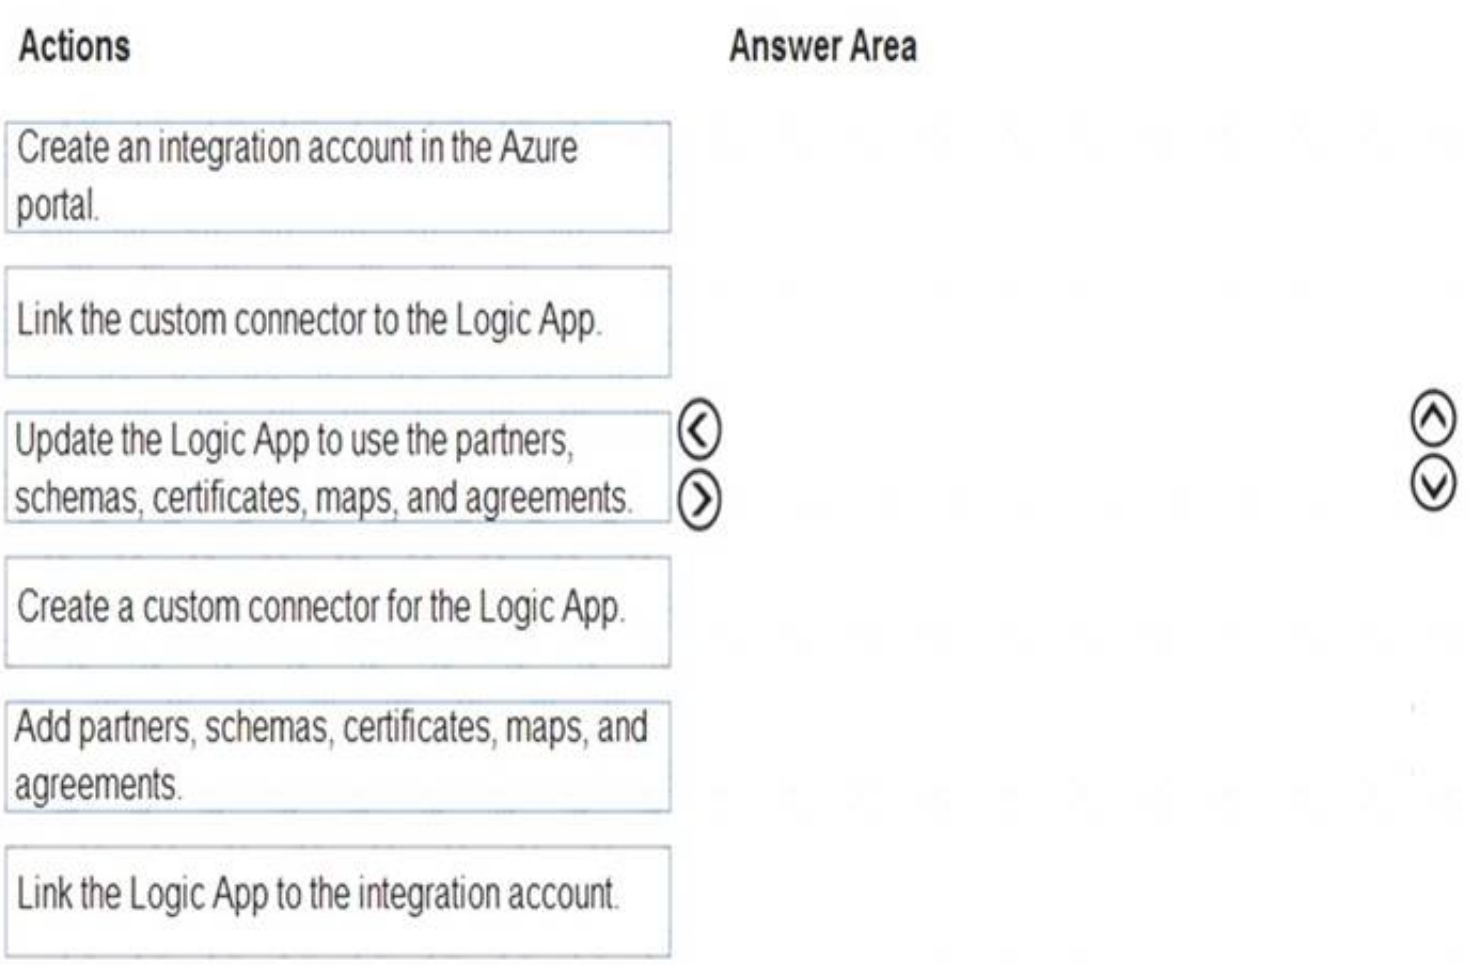

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Create an integration account in the Azure portal
You can define custom metadata for artifacts in integration accounts and get that metadata during runtime for your logic app to use. For example, you can provide metadata for artifacts, such as partners, agreements, schemas, and maps - all store metadata using key-value pairs.
Step 2: Link the Logic App to the integration account
A logic app that's linked to the integration account and artifact metadata you want to use. Step 3: Add partners, schemas, certificates, maps, and agreements
Step 4: Create a custom connector for the Logic App. References:
https://docs.microsoft.com/bs-latn-ba/azure/logic-apps/logic-apps-enterprise-integration-metadata

NEW QUESTION 3
- (Exam Topic 1)
You need to configure Azure CDN for the Shipping web site.
Which configuration options should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

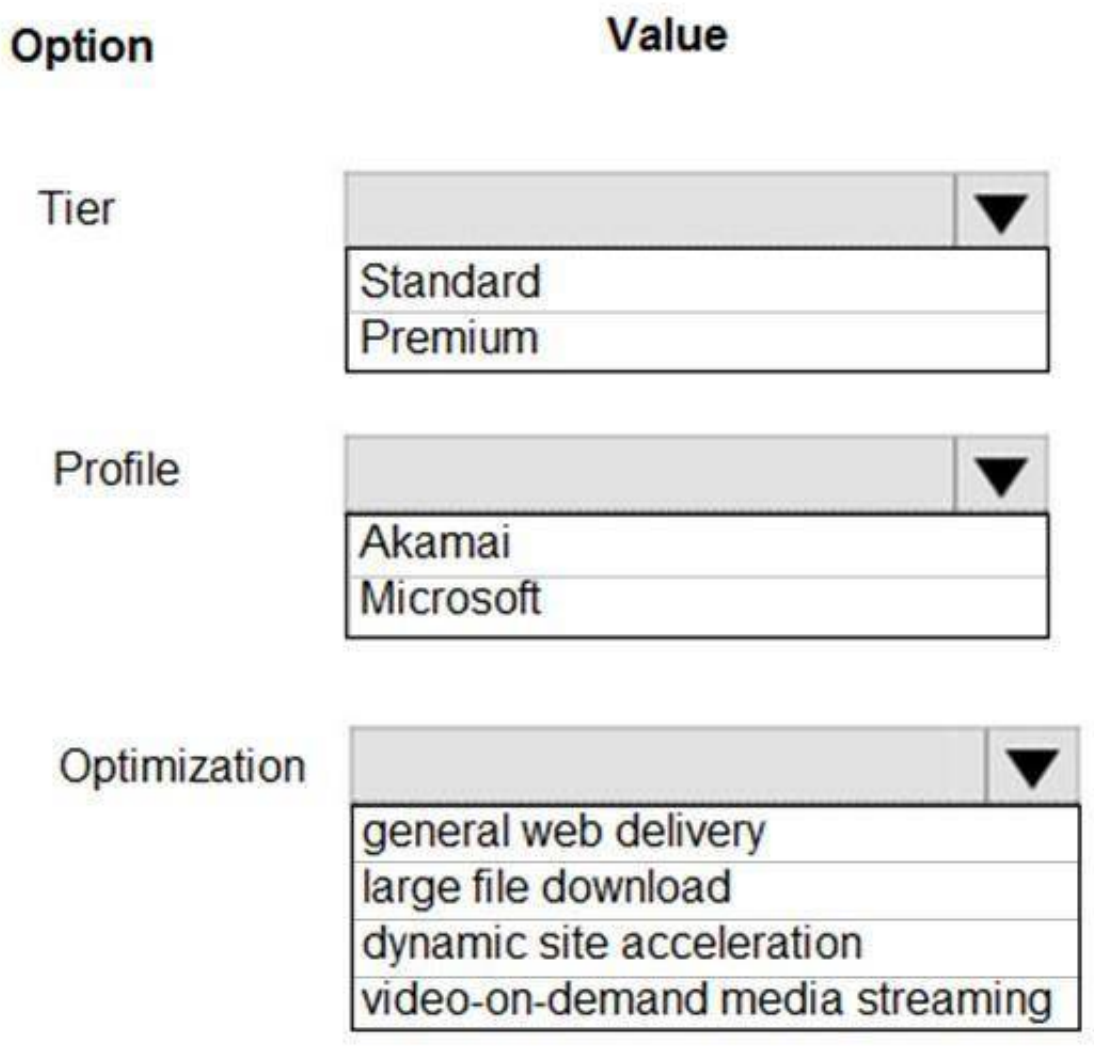

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Scenario: Shipping website
Use Azure Content Delivery Network (CDN) and ensure maximum performance for dynamic content while minimizing latency and costs.
Tier: Standard Profile: Akamai
Optimization: Dynamic site acceleration
Dynamic site acceleration (DSA) is available for Azure CDN Standard from Akamai, Azure CDN Standard from Verizon, and Azure CDN Premium from Verizon profiles.
DSA includes various techniques that benefit the latency and performance of dynamic content. Techniques include route and network optimization, TCP optimization, and more.
You can use this optimization to accelerate a web app that includes numerous responses that aren't cacheable. Examples are search results, checkout transactions, or real-time data. You can continue to use core Azure CDN caching capabilities for static data.
Reference:
https://docs.microsoft.com/en-us/azure/cdn/cdn-optimization-overview

NEW QUESTION 4
- (Exam Topic 1)
You need to secure the Shipping Function app.
How should you configure the app? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Setting	Value
Authorization level	Function / Anonymous / Admin
User claims	JSON Web Token (JWT) / Shared Access Signature (SAS) token / API Key
Trigger type	blob / HTTP / queue / timer

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Scenario: Shipping Function app: Implement secure function endpoints by using app-level security and include Azure Active Directory (Azure AD).
Box 1: Function
Box 2: JSON based Token (JWT)
Azure AD uses JSON based tokens (JWTs) that contain claims Box 3: HTTP
How a web app delegates sign-in to Azure AD and obtains a token
User authentication happens via the browser. The OpenID protocol uses standard HTTP protocol messages. References:
https://docs.microsoft.com/en-us/azure/active-directory/develop/authentication-scenarios

NEW QUESTION 5
- (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 6
- (Exam Topic 3)
You are developing an Azure Function App by using Visual Studio. The app will process orders input by an Azure Web App. The web app places the order information into Azure Queue Storage.
You need to review the Azure Function App code shown below. NOTE: Each correct selection is worth one point.

```
public static class
{
  [FunctionName("ProcessOrders")]
  public static void ProcessOrders([QueueTrigger("incoming-orders")]CloudQueueMessage myQueueItem, [Table
("Orders")]ICollector<Order> tableBindings,
TraceWriter log)
  {
    log.Info($"Processing Order: {myQueueItem.Id}");
    log.Info($"Queue Insertion Time: {myQueueItem.InsertionTime}");
    log.Info($"Queue Expiration Time: {myQueueItem.ExpirationTime}");
    tableBindings.Add(JsonConvert.DeserializeObject<Order>(myQueueItem.AsString));
  }
  [FunctionName("ProcessOrders-Poison")]
  public static void ProcessFailedOrders([QueueTrigger("incoming-orders-poison")]CloudQueueMessage
myQueueItem, TraceWriter log)
  {
    logError($"Failed to process order: {myQueueItem.AsString}");
  }
}
```

	Yes	No
The code will log the time that the order was processed from the queue.	○	○
When the ProcessOrders function fails, the function will retry up to five times for a given order, including the first try.	○	○
When there are multiple orders in the queue, a batch of orders will be retrieved from the queue and the ProcessOrders function will run multiple instances concurrently to process the orders.	○	○
The ProcessOrders function will output the order to an Orders table in Azure Table Storage.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: No
ExpirationTime - The time that the message expires. InsertionTime - The time that the message was added to the queue.
Box 2: Yes
maxDequeueCount - The number of times to try processing a message before moving it to the poison queue.
Default value is 5.
Box 3: Yes
When there are multiple queue messages waiting, the queue trigger retrieves a batch of messages and invokes function instances concurrently to process them.
By default, the batch size is 16. When the number being processed gets down to 8, the runtime gets another batch and starts processing those messages. So the maximum number of concurrent messages being processed per function on one virtual machine (VM) is 24.
Box 4: Yes References:
https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-queue

NEW QUESTION 7
- (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
Yes	Enable AlwaysOn encryption.	
No	Set the column encryption setting to disabled.	
	Assign users to the Public fixed database role.	
	Store column encryption keys in the system catalog view in the database.	

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 8
- (Exam Topic 3)
You are creating a script that will run a large workload on an Azure Batch pool. Resources will be reused and do not need to be cleaned up after use.
You have the following parameters:
You need to write an Azure CLI script that will create the jobs, tasks, and the pool.
In which order should you arrange the commands to develop the solution? To answer, move the appropriate commands from the list of command segments to the answer area and arrange them in the correct order.

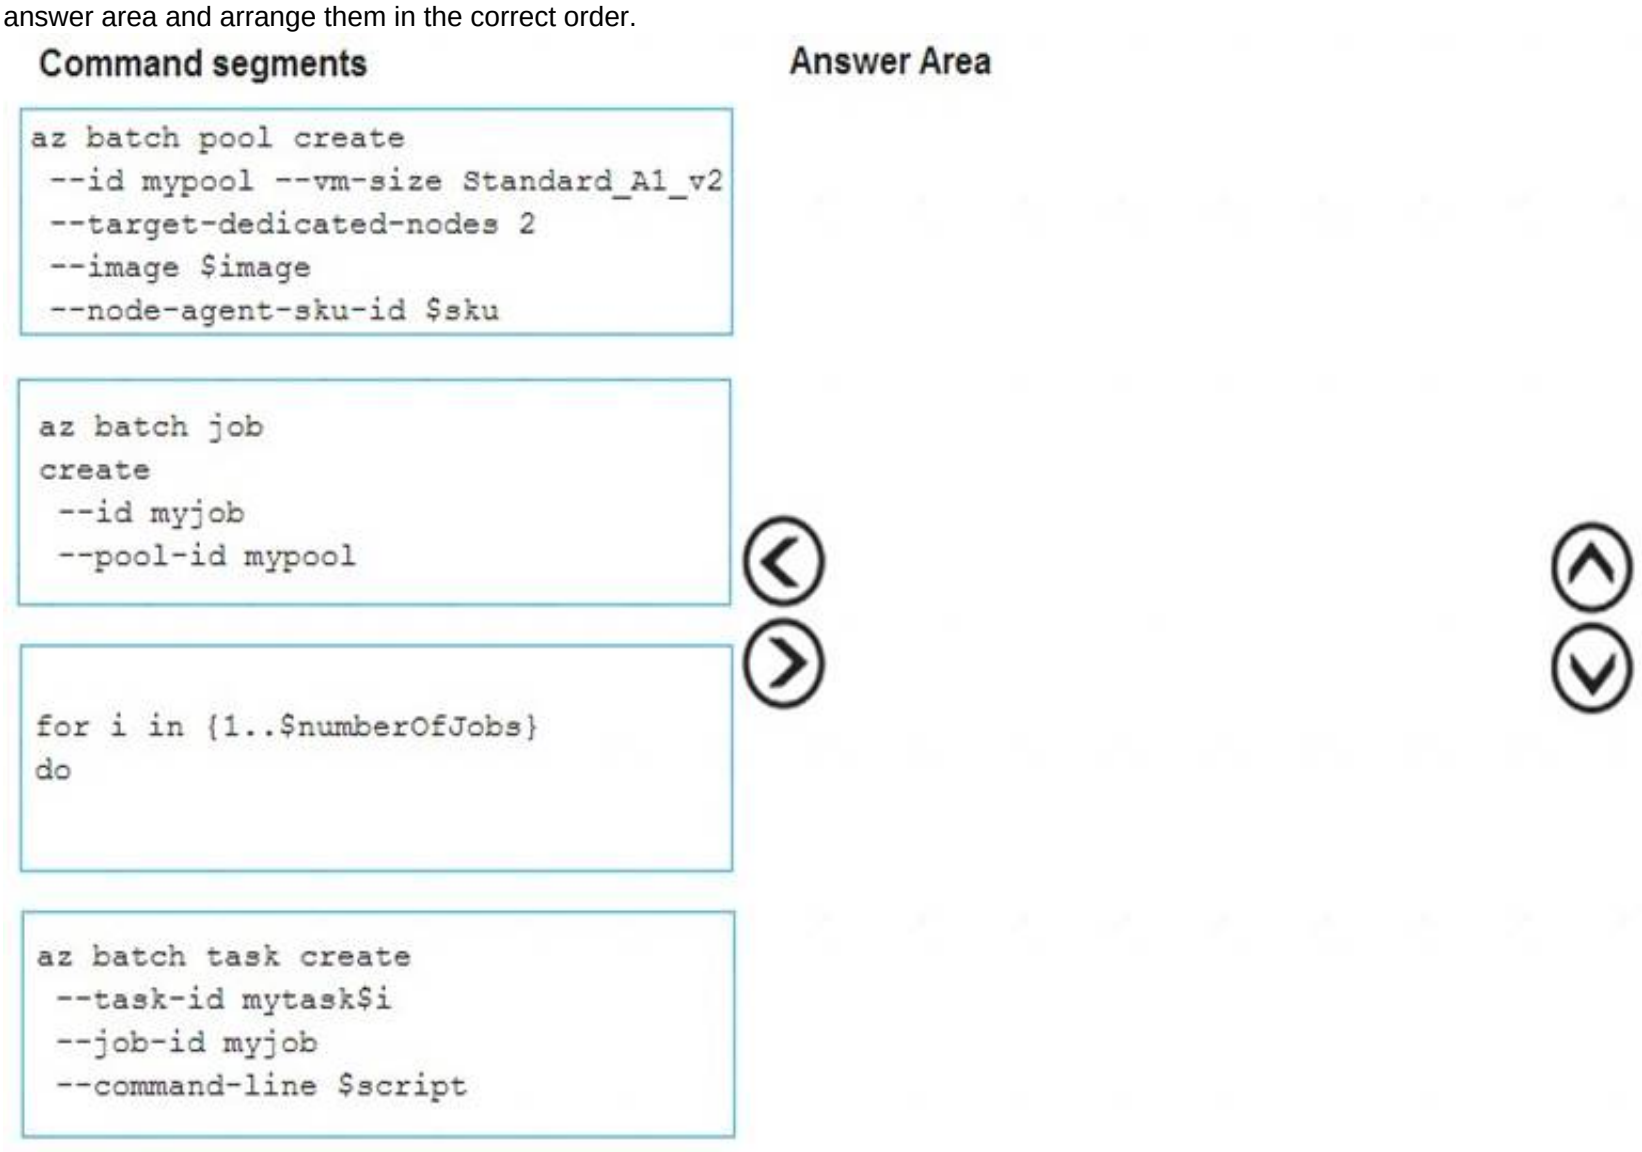

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: az batch pool create
Create a new Linux pool with a virtual machine configuration. az batch pool create \
--id mypool \
--vm-size Standard_A1 \
--target-dedicated 2 \
--image canonical:ubuntuserver:16.04-LTS \
--node-agent-sku-id "batch.node.ubuntu 16.04" Step 2: az batch job create
Create a new job to encapsulate the tasks that are added. az batch job create \
--id myjob \
--pool-id mypool
Step 3: az batch task create
Add tasks to the job. Here the task is a basic shell command. az batch task create \
--job-id myjob \
--task-id task1 \
--command-line "/bin/bash -c 'printenv AZ_BATCH_TASK_WORKING_DIR"' Step 4: for i in {1..$numberOfJobs} do
References:
https://docs.microsoft.com/bs-latn-ba/azure/batch/scripts/batch-cli-sample-run-job

NEW QUESTION 9
- (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 10
- (Exam Topic 3)
You develop an app that allows users to upload photos and videos to Azure storage. The app uses a storage REST API call to upload the media to a blob storage account named Account1. You have blob storage
containers named Container1 and Container2. Uploading of videos occurs on an irregular basis.
You need to copy specific blobs from Container1 to Container2 in real time when specific requirements are met, excluding backup blob copies.
What should you do?

A. Download the blob to a virtual machine and then upload the blob to Container2.
B. Run the Azure PowerShell command Start-AzureStorageBlobCopy.
C. Copy blobs to Container2 by using the Put Blob operation of the Blob Service REST API.
D. Use AzCopy with the Snapshot switch blobs to Container2.

Answer: B

Explanation:
The Start-AzureStorageBlobCopy cmdlet starts to copy a blob. Example 1: Copy a named blob
C:\PS>Start-AzureStorageBlobCopy -SrcBlob "ContosoPlanning2015" -DestContainer "ContosoArchives"
-SrcContainer "ContosoUploads"
This command starts the copy operation of the blob named ContosoPlanning2015 from the container named ContosoUploads to the container named ContosoArchives.
References:
https://docs.microsoft.com/en-us/powershell/module/azure.storage/start-azurestorageblobcopy?view=azurermps

NEW QUESTION 10
- (Exam Topic 3)
You are a developer for a SaaS company that offers many web services. All web services for the company must meet the following requirements:
- Use API Management to access the services
- Use OpenID Connect for authentication
- Prevent anonymous usage

A recent security audit found that several web services can be called without any authentication. Which API Management policy should you implement?

A. jsonp
B. authentication-certificate
C. check-header
D. validate-jwt

Answer: D

Explanation:
Add the validate-jwt policy to validate the OAuth token for every incoming request. Reference:
https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-protect-backend-with-aad

NEW QUESTION 13
- (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 are developing an Azure solution to collect point-of-sale (POS) device data from 2,000 stores located throughout the world. A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data.
You must store the device data in Azure Blob storage. Device data must be correlated based on a device identifier. Additional stores are expected to open in the future.
You need to implement a solution to receive the device data.
Solution: Provision an Azure Service Bus. Configure a topic to receive the device data by using a correlation filter.
Does the solution meet the goal?

A. Yes
B. No

Answer: A

Explanation:
A message is raw data produced by a service to be consumed or stored elsewhere. The Service Bus is for high-value enterprise messaging, and is used for order processing and financial transactions.
Reference:
https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services

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 are developing an Azure solution to collect point-of-sale (POS) device data from 2,000 stores located throughout the world. A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data.
You must store the device data in Azure Blob storage. Device data must be correlated based on a device identifier. Additional stores are expected to open in the future.
You need to implement a solution to receive the device data.
Solution: Provision an Azure Notification Hub. Register all devices with the hub. Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Instead use an Azure Service Bus, which is used order processing and financial transactions.
Reference:
https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services

NEW QUESTION 15
- (Exam Topic 3)
You are developing an Azure Cosmos DB solution by using the Azure Cosmos DB SQL API. The data includes millions of documents. Each document may contain hundreds of properties.
The properties of the documents do not contain distinct values for partitioning. Azure Cosmos DB must scale individual containers in the database to meet the performance needs of the application by spreading the workload evenly across all partitions over time.
You need to select a partition key.
Which two partition keys can you use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. a concatenation of multiple property values with a random suffix appended
B. a single property value that does not appear frequently in the documents
C. a hash suffix appended to a property value
D. a value containing the collection name
E. a single property value that appears frequently in the documents

Answer: AC

Explanation:
You can form a partition key by concatenating multiple property values into a single artificial partitionKey property. These keys are referred to as synthetic keys.
Another possible strategy to distribute the workload more evenly is to append a random number at the end of the partition key value. When you distribute items in this way, you can perform parallel write operations across partitions.
Note: It's the best practice to have a partition key with many distinct values, such as hundreds or thousands. The goal is to distribute your data and workload evenly across the items associated with these partition key values. If such a property doesn't exist in your data, you can construct a synthetic partition key.
References:
https://docs.microsoft.com/en-us/azure/cosmos-db/synthetic-partition-keys

NEW QUESTION 16
- (Exam Topic 3)
You are implementing a software as a service (SaaS) ASP.NET Core web service that will run as an Azure Web App. The web service will use an on-premises SQL Server database for storage. The web service also includes a WebJob that processes data updates. Four customers will use the web service.

- Each instance of the WebJob processes data for a single customer and must run as a singleton instance.
- Each deployment must be tested by using deployment slots prior to serving production data.
- Azure costs must be minimized.
- Azure resources must be located in an isolated network. You need to configure the App Service plan for the Web App.

How should you configure the App Service plan? To answer, select the appropriate settings in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

App service plan setting	Value
Number of VM instances	2 / 4 / 8 / 16
Pricing tier	Isolated / Standard / Premium / Consumption

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Number of VM instances: 4
You are not charged extra for deployment slots. Pricing tier: Isolated
The App Service Environment (ASE) is a powerful feature offering of the Azure App Service that gives network isolation and improved scale capabilities. It is essentially a deployment of the Azure App Service into a subnet of a customer's Azure Virtual Network (VNet).
References:
https://azure.microsoft.com/sv-se/blog/announcing-app-service-isolated-more-power-scale-and-ease-of-use/

NEW QUESTION 19
- (Exam Topic 3)
You are developing an ASP.NET Core Web API web service. The web service uses Azure Application Insights for all telemetry and dependency tracking. The web service reads and writes data to a database other than Microsoft SQL Server.
You need to ensure that dependency tracking works for calls to the third-party database.
Which two Dependency Telemetry properties should you store in the database? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Telemetry.Context.Operation.Id
B. Tetemetry.Context.Cloud.RoleInstance
C. Telemetry.Id
D. Telemetry.ContextSession.Id
E. Telemetry.Name

Answer: AC

Explanation:
References:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/custom-operations-tracking

NEW QUESTION 22
- (Exam Topic 3)
You are building a website to access project data related to terms within your organization. The website does not allow anonymous access. Authentication performed using an Azure Active Directory (Azure AD) app named internal.
The website has the following authentication requirements:
•Azure AD users must be able to login to the website.
•Personalization of the website must be based on membership in Active Directory groups. You need to configure the application's manifest to meet the authentication requirements.
How should you configure the manifest? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

```
{
    ...
    "appId": "d61126e3-089b-4adb-b721-
d5023213df7d",
    [ ▼ ] : "All",
    [ ▼ ] : true
    ...
}
```

- "optionalClaims"
- "groupMembershipClaims"

- "allowPublicClient"
- "oauth2Permissions"
- "requiredResourceAccess"
- "oauth2AllowImplicitFlow"

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: groupMembershipClaims
Scenario: Personalization of the website must be based on membership in Active Directory groups. Group claims can also be configured in the Optional Claims section of the Application Manifest. Enable group membership claims by changing the groupMembershipClaim
The valid values are: "All" "SecurityGroup" "DistributionList" "DirectoryRole"
Box 2: oauth2Permissions
Scenario: Azure AD users must be able to login to the website.
oauth2Permissions specifies the collection of OAuth 2.0 permission scopes that the web API (resource) app exposes to client apps. These permission scopes may be granted to client apps during consent.

NEW QUESTION 24
- (Exam Topic 3)
You have an Azure Batch project that processes and converts files and stores the files in Azure storage. You are developing a function to start the batch job.
You add the following parameters to the function.

Parameter name	Description
fileTasks	a list of tasks to be run
jobId	the identifier that must be assigned to the job
outputContainerSasUrl	a storage SAS URL to store successfully converted files
failedContainerSasUrl	a storage SAS URL to store copies of files that failed to convert.

You must ensure that converted files are placed in the container referenced by the outputContainerSasUrl parameter. Files which fail to convert are places in the container referenced by the failedContainerSasUrl parameter.
You need to ensure the files are correctly processed.
How should you complete the code segment? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

```
public list<CloudTasks> StartTasks(List<FileTask> fileTasks, string jobId,
 string outputContainerSasUrl, string failedContainerSasUrl)
{
  BatchSharedKeyCredentials sharedKeyCredentials =
   new BatchSharedKeyCredentials(batchAccountUrl, batchAccountName,
batchAccountKey);
 List<CloudTask> tasks = new List<CloudTask>();
 using (BatchClient batchClient = BatchClient.Open(sharedKeyCredentials))
{
   CloudJob = batchClient.JobOperations.                    ();
```

- GetJob
- GetTask
- EnableJob
- CreateJob

```
        job.Id = jobId,
        job.PoolInfromation = new PoolInformation { PoolId = poolId };
        job.Commit();
        fileTasks.ForEach((fileTask) =>
        {
         string taskId = $"Task{DateTime.Now.ToFileTimeUtc().ToString()}";
         CloudTask task = new CloudTask (taskId, fileTask.Command);
         List<OutputFile> outputFileList = new List<OutputFile>();
         OutputFileBlobContainerDestination outputContainer =
          new OutputFileBlobContainerDestination(outputContainerSasUrl);
       OutputFileBlobContainerDestination failedContainer =
   new OutputFileBlobContainerDestination (failedContainerSasUrl);
outputFileList.Add(new OutputFile(fileTask.Output,
   new OutputFileDestination(outputContainer),
   new OutputFileUploadOptions(OutputFileUploadCondition.                ));
```

- TaskSuccess
- TaskFailure
- TaskCompletion

```
  outputFileList.Add(new OutputFile(fileTask.Output,
   new OutputFileDestination(failedContainer),
   new OutputFileUploadOptions(OutputFileUploadCondition,                ));
```

- TaskSuccess
- TaskFailure
- TaskCompletion

```
    task                =outputFileList;
```

- OutputFiles
- FilesToStage
- ResourceFiles
- StageFiles

```
      task.Add(task);
      });
      ]
      return tasks,
}
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: CreateJob
Box 2: TaskSuccess
TaskSuccess: Upload the file(s) only after the task process exits with an exit code of 0.
Incorrect: TaskCompletion: Upload the file(s) after the task process exits, no matter what the exit code was. Box 3: TaskFailure
TaskFailure:Upload the file(s) only after the task process exits with a nonzero exit code. Box 4: OutputFiles
To specify output files for a task, create a collection of OutputFile objects and assign it to the CloudTask.OutputFiles property when you create the task.

References:
https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.batch.protocol.models.outputfileuploadcondition https://docs.microsoft.com/en-us/azure/batch/batch-task-output-files

NEW QUESTION 27
- (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 28
- (Exam Topic 3)
Contoso, Ltd. provides an API to customers by using Azure API Management (APIM). The API authorizes users with a JWT token.
You must implement response caching for the APIM gateway. The caching mechanism must detect the user ID of the client that accesses data for a given location and cache the response for that user ID.
You need to add the following policies to the policies file:
• a set-variable policy to store the detected user identity
• a cache-lookup-value policy
• a cache-store-value policy
• a find-and-replace policy to update the response body with the user profile information
To which policy section should you add the policies? To answer, drag the appropriate sections to the correct policies. Each section 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

Policy section	Policy	Policy section
	Set-variable	
Inbound	Cache-lookup-value	
Outbound	Cache-store-value	
	Find-and-replace	

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Inbound.
A set-variable policy to store the detected user identity. Example:
<policies>
<inbound>
<!-- How you determine user identity is application dependent -->
<set-variable name="enduserid"
value="@(context.Request.Headers.GetValueOrDefault("Authorization","").Split(' ')[1].AsJwt()?.Subject)" /> Box 2: Inbound
A cache-lookup-value policy Example:
<inbound>
<base />
<cache-lookup vary-by-developer="true | false" vary-by-developer-groups="true | false" downstream-caching-type="none | private | public" must-revalidate="true | false">
<vary-by-query-parameter>parameter name</vary-by-query-parameter> <!-- optional, can repeated several times -->
</cache-lookup>
</inbound>

Box 3: Outbound
A cache-store-value policy. Example:
<outbound>
<base />
<cache-store duration="3600" />
</outbound> Box 4: Outbound
A find-and-replace policy to update the response body with the user profile information. Example:
<outbound>
<!-- Update response body with user profile-->
<find-and-replace from="'$userprofile$"'
to="@((string)context.Variables["userprofile"])" />
<base />
</outbound> Reference:
https://docs.microsoft.com/en-us/azure/api-management/api-management-caching-policies
https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key

NEW QUESTION 30
- (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();
[ ▼ ]
    i["tip"] = 0;
  }
[ ▼ ]
```

- _.value();
- _.readDocument('item');
- getContext().getRequest();
- getContext().getResponse();

- if (!("tip" in i)) {
- if (request.getValue("tip") === null){
- if (isNaN(i)["tip"] || i["tip"]=== null) {
- if (typeof_.pluck("tip") == 'number') {

- r.setBody(i);
- r.setValue(i);
- _.upsertDocument(i);
- _.replaceDocument(i)

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 34
- (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    Table Query<DynamicTableEntity> query - new TableQuery<DynamicTableEntity>().Select(new.string[.] { "Score" })
       .Where(TableQuery.GenerateFilterConditionForInt("Score", QueryComparisons.GretaerThanOrEqual, 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    {
      Console.Writeline($"{scoreItem.Key} {scoreItem.Value}");
7    }
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	○	○
The code will display a maximum of twenty records.	○	○
All records will be sent to the client. The client will display records for scores greater than or equal to 15,000.	○	○
The scoreItem.Key property of the KeyValuePairs that ExecuteQuery returns will contain a value for PlayerID.	○	○

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 36
- (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 solution that will be deployed to an Azure Kubernetes Service (AKS) cluster. The solution will include a custom VNet, Azure Container Registry images, and an Azure Storage account.
The solution must allow dynamic creation and management of all Azure resources within the AKS cluster. You need to configure an AKS cluster for use with the Azure APIs.

Solution: Enable the Azure Policy Add-on for Kubernetes to connect the Azure Policy service to the GateKeeper admission controller for the AKS cluster. Apply a built-in policy to the cluster.
Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Instead create an AKS cluster that supports network policy. Create and apply a network to allow traffic only from within a defined namespace
References:
https://docs.microsoft.com/en-us/azure/aks/use-network-policies

NEW QUESTION 37
- (Exam Topic 3)
You are preparing to deploy an ASP.NET Core website to an Azure Web App from a GitHub repository. The website includes static content generated by a script.
You plan to use the Azure Web App continuous deployment feature.
You need to run the static generation script before the website starts serving traffic.
What are two possible ways to achieve this goal? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. Create a file named .deployment in the root of the repository that calls a script which generates the static content and deploys the website.
B. Add a PreBuild target in the websites csproj project file that runs the static content generation script.
C. Create a file named run.cmd in the folder /run that calls a script which generates the static content and deploys the website.
D. Add the path to the static content generation tool to WEBSITE_RUN_FROM_PACKAGE setting in the host.json file.

Answer: AD

Explanation:
A: To customize your deployment, include a .deployment file in the repository root.
You just need to add a file to the root of your repository with the name .deployment and the content: [config]
command = YOUR COMMAND TO RUN FOR DEPLOYMENT
this command can be just running a script (batch file) that has all that is required for your deployment, like copying files from the repository to the web root directory for example.
D: In Azure, you can run your functions directly from a deployment package file in your function app. The other option is to deploy your files in the d:\home\site\wwwroot directory of your function app (see A above).
To enable your function app to run from a package, you just add a WEBSITE_RUN_FROM_PACKAGE setting to your function app settings.
Note: The host.json metadata file contains global configuration options that affect all functions for a function app.
References:
https://github.com/projectkudu/kudu/wiki/Custom-Deployment-Script
https://docs.microsoft.com/bs-latn-ba/azure/azure-functions/run-functions-from-deployment-package

NEW QUESTION 40
- (Exam Topic 3)
You develop software solutions for a mobile delivery service. You are developing a mobile app that users can use to order from a restaurant in their area. The app uses the following workflow:
* 1. A driver selects the restaurants for which they will deliver orders.
* 2. Orders are sent to all available drivers in an area.
* 3. Only orders for the selected restaurants will appear for the driver.
* 4. The first driver to accept an order removes it from the list of available orders. You need to implement an Azure Service Bus solution.
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.

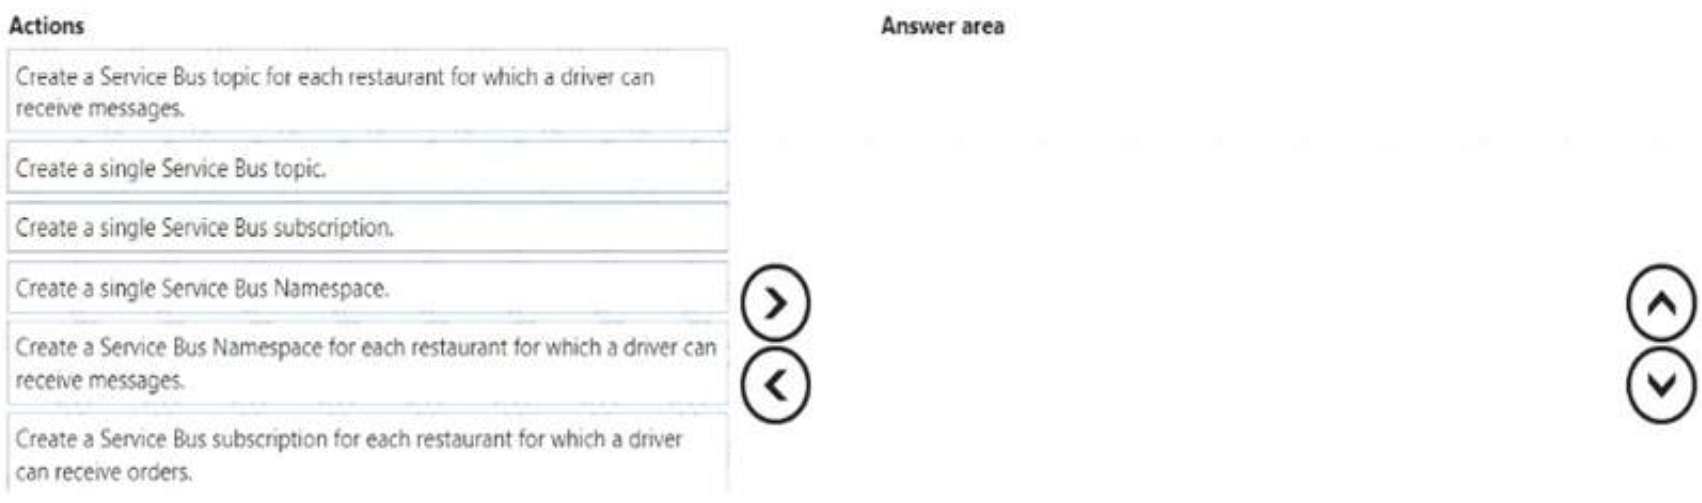

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Create a single Service Bus Namespace
To begin using Service Bus messaging entities in Azure, you must first create a namespace with a name that is unique across Azure. A namespace provides a scoping container for addressing Service Bus resources within your application.
Box 2: Create a Service Bus Topic for each restaurant for which a driver can receive messages. Create topics.
Box 3: Create a Service Bus subscription for each restaurant for which a driver can receive orders. Topics can have multiple, independent subscriptions.
References:

https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-messaging-overview

NEW QUESTION 44
- (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.	○	○
Edge nodes must be purged of all cache assets.	○	○
The compression type is supported.	○	○

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 49
- (Exam Topic 3)
You provide an Azure API Management managed web service lo 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 uses? 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 51
- (Exam Topic 3)
You are creating an app that uses Event Grid to connect with other services. Your app's event data will be sent to a serverless function that checks compliance. This function is maintained by your company.
You write a new event subscription at the scope of your resource. The event must be invalidated after 3 specific period of time. You need to configure Event Grid to ensure security.
What should you implement? To answer, select the appropriate options in [he answer area. NOTE: Each correct selection is worth one point

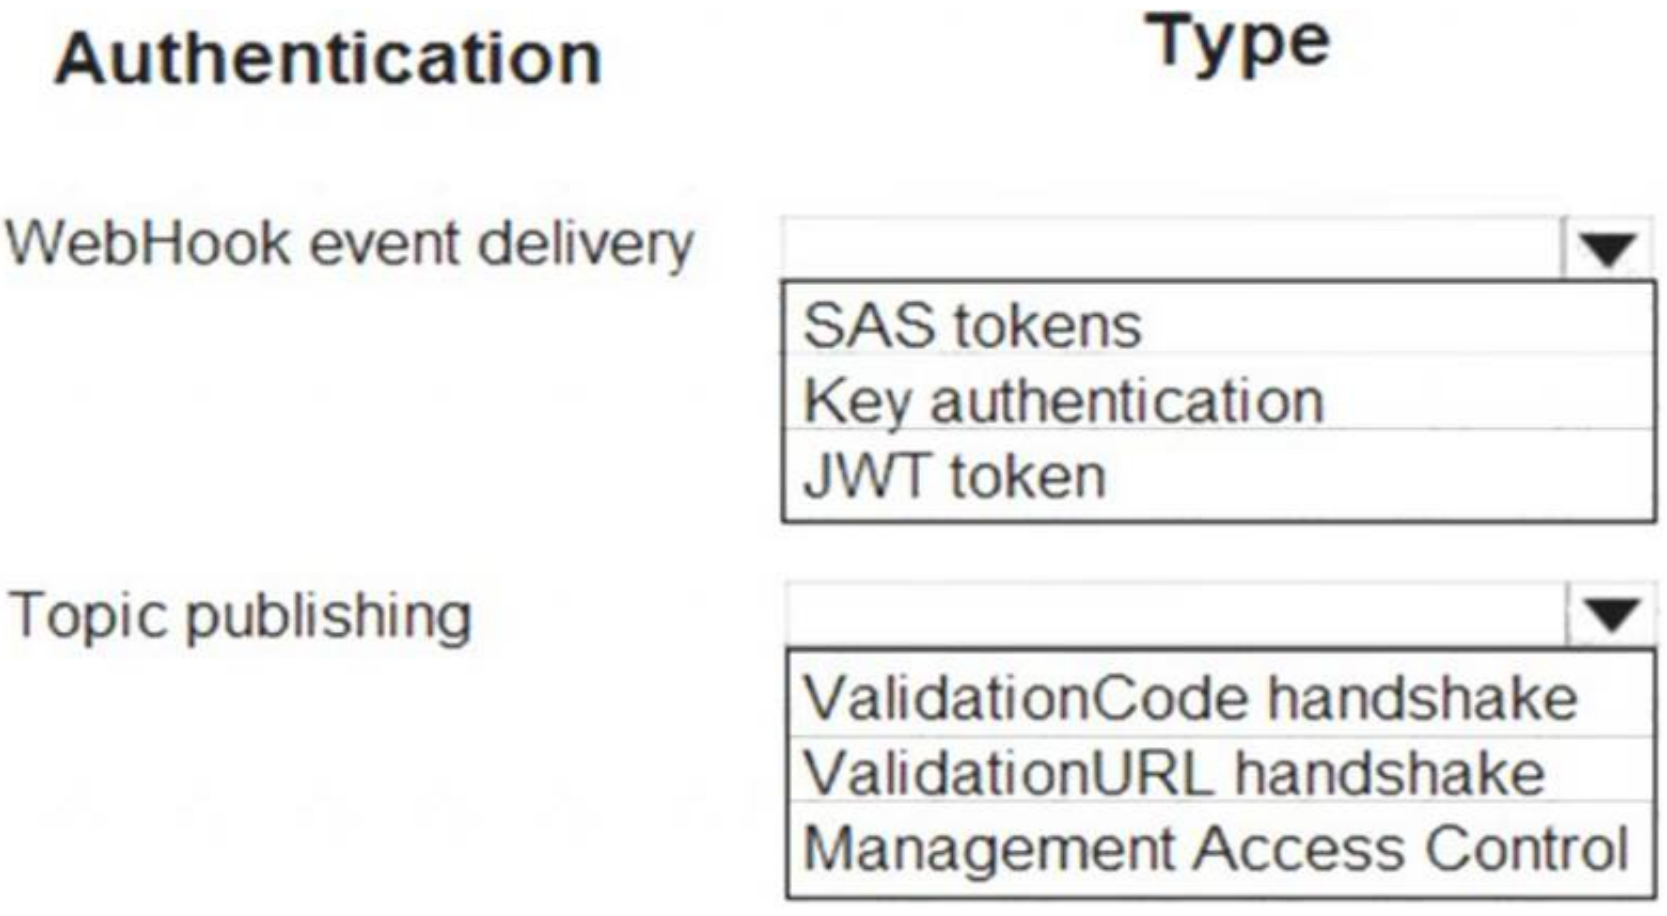

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: SAS tokens
Custom topics use either Shared Access Signature (SAS) or key authentication. Microsoft recommends SAS, but key authentication provides simple programming, and is compatible with many existing webhook publishers.
In this case we need the expiration time provided by SAS tokens. Box 2: ValidationCode handshake
Event Grid supports two ways of validating the subscription: ValidationCode handshake (programmatic) and ValidationURL handshake (manual).
If you control the source code for your endpoint, this method is recommended.

NEW QUESTION 56
- (Exam Topic 3)
A company is implementing a publish-subscribe (Pub/Sub) messaging component by using Azure Service Bus. You are developing the first subscription application.
In the Azure portal you see that messages are being sent to the subscription for each topic. You create and initialize a subscription client object by supplying the correct details, but the subscription application is still not consuming the messages.
You need to complete the source code of the subscription client What should you do?

A. await subscriptionClient.CloseAsync();
B. await subscriptionClient.AddRuleAsync(new RuleDescription(RuleDescription.DefaultRuleName, new TrueFilter()));
C. subscriptionClient.RegisterMessageHandler(ProcessMessagesAsync, messageHandlerOptions);
D. subscriptionClient = new SubscriptionClient(ServiceBusConnectionString, TopicName, SubscriptionName);

Answer: C

Explanation:
Using topic client, call RegisterMessageHandler which is used to receive messages continuously from the entity. It registers a message handler and begins a new thread to receive messages. This handler is waited on every time a new message is received by the receiver.
subscriptionClient.RegisterMessageHandler(ReceiveMessagesAsync, messageHandlerOptions); References:
https://www.c-sharpcorner.com/article/azure-service-bus-topic-and-subscription-pub-sub/

NEW QUESTION 60
- (Exam Topic 3)
You develop a gateway solution for a public facing news API. The news API back end is implemented as a RESTful service and uses an OpenAPI specification.
You need to ensure that you can access the news API by using an Azure API Management service instance. Which Azure PowerShell command should you run?

A. Import-AzureRmApiManagementApi –Context $ApiMgmtContext –SpecificationFormat "Swagger" -SpecificationPath $SwaggerPath –Path $Path
B. New-AzureRmApiManagementBackend -Context $ApiMgmtContext -Url $Url -Protocol http
C. New-AzureRmApiManagement –ResourceGroupName $ResourceGroup –Name $Name – Location $Location –Organization $Org –AdminEmail $AdminEmail
D. New-AzureRmApiManagementBackendProxy –Url $ApiUrl

Answer: D

Explanation:
New-AzureRmApiManagementBackendProxy creates a new Backend Proxy Object which can be piped when creating a new Backend entity.
Example: Create a Backend Proxy In-Memory Object
PS C:\>$secpassword = ConvertTo-SecureString "PlainTextPassword" -AsPlainText -Force
PS C:\>$proxyCreds = New-Object System.Management.Automation.PSCredential ("foo", $secpassword) PS C:\>$credential = New-AzureRmApiManagementBackendProxy -Url "http://12.168.1.1:8080"
-ProxyCredential $proxyCreds

PS C:\>$apimContext = New-AzureRmApiManagementContext -ResourceGroupName "Api-Default-WestUS" -ServiceName "contoso"
PS C:\>$backend = New-AzureRmApiManagementBackend -Context $apimContext -BackendId 123 -Url 'https://contoso.com/awesomeapi' -Protocol http -Title "first backend" -SkipCertificateChainValidation $true
-Proxy $credential -Description "backend with proxy server"
Creates a Backend Proxy Object and sets up Backend

NEW QUESTION 61
- (Exam Topic 3)
You are developing a mobile instant messaging app for a company. The mobile app must meet the following requirements:
• Support offline data sync.
• Update the latest messages during normal sync cycles. You need to implement Offline Data Sync.
Which two actions should you perform? Each conn I answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Retrieve records from Offline Data Sync on every call to the PullAsync method.
B. Retrieve records from Offline Data Sync using an Incremental Sync.
C. Push records to Offline Data Sync using an Incremental Sync.
D. Return the updatedAt column from the Mobile Service Backend and implement sorting by using the column.
E. Return the updatedAt column from the Mobile Service Backend and implement sorting by the message id.

Answer: BE

Explanation:
B: Incremental Sync: the first parameter to the pull operation is a query name that is used only on the client. If you use a non-null query name, the Azure Mobile SDK performs an incremental sync. Each time a pull operation returns a set of results, the latest updatedAt timestamp from that result set is stored in the SDK local system tables. Subsequent pull operations retrieve only records after that timestamp.
E (not D): To use incremental sync, your server must return meaningful updatedAt values and must also support sorting by this field. However, since the SDK adds its own sort on the updatedAt field, you cannot use a pull query that has its own orderBy clause.
References:
https://docs.microsoft.com/en-us/azure/app-service-mobile/app-service-mobile-offline-data-sync

NEW QUESTION 62
- (Exam Topic 3)
Fourth Coffee has an ASP.NET Core web app that runs in Docker. The app is mapped to the www.fourthcoffee.com domain.
Fourth Coffee is migrating this application to Azure.
You need to provision an App Service Web App to host this docker image and map the custom domain to the App Service web app.
A resource group named FourthCoffeePublicWebResourceGroup has been created in the WestUS region that contains an App Service Plan named AppServiceLinuxDockerPlan.
Which order should the CLI commands be used to develop the solution? To answer, move all of the Azure CLI command from the list of commands to the answer area and arrange them in the correct order.

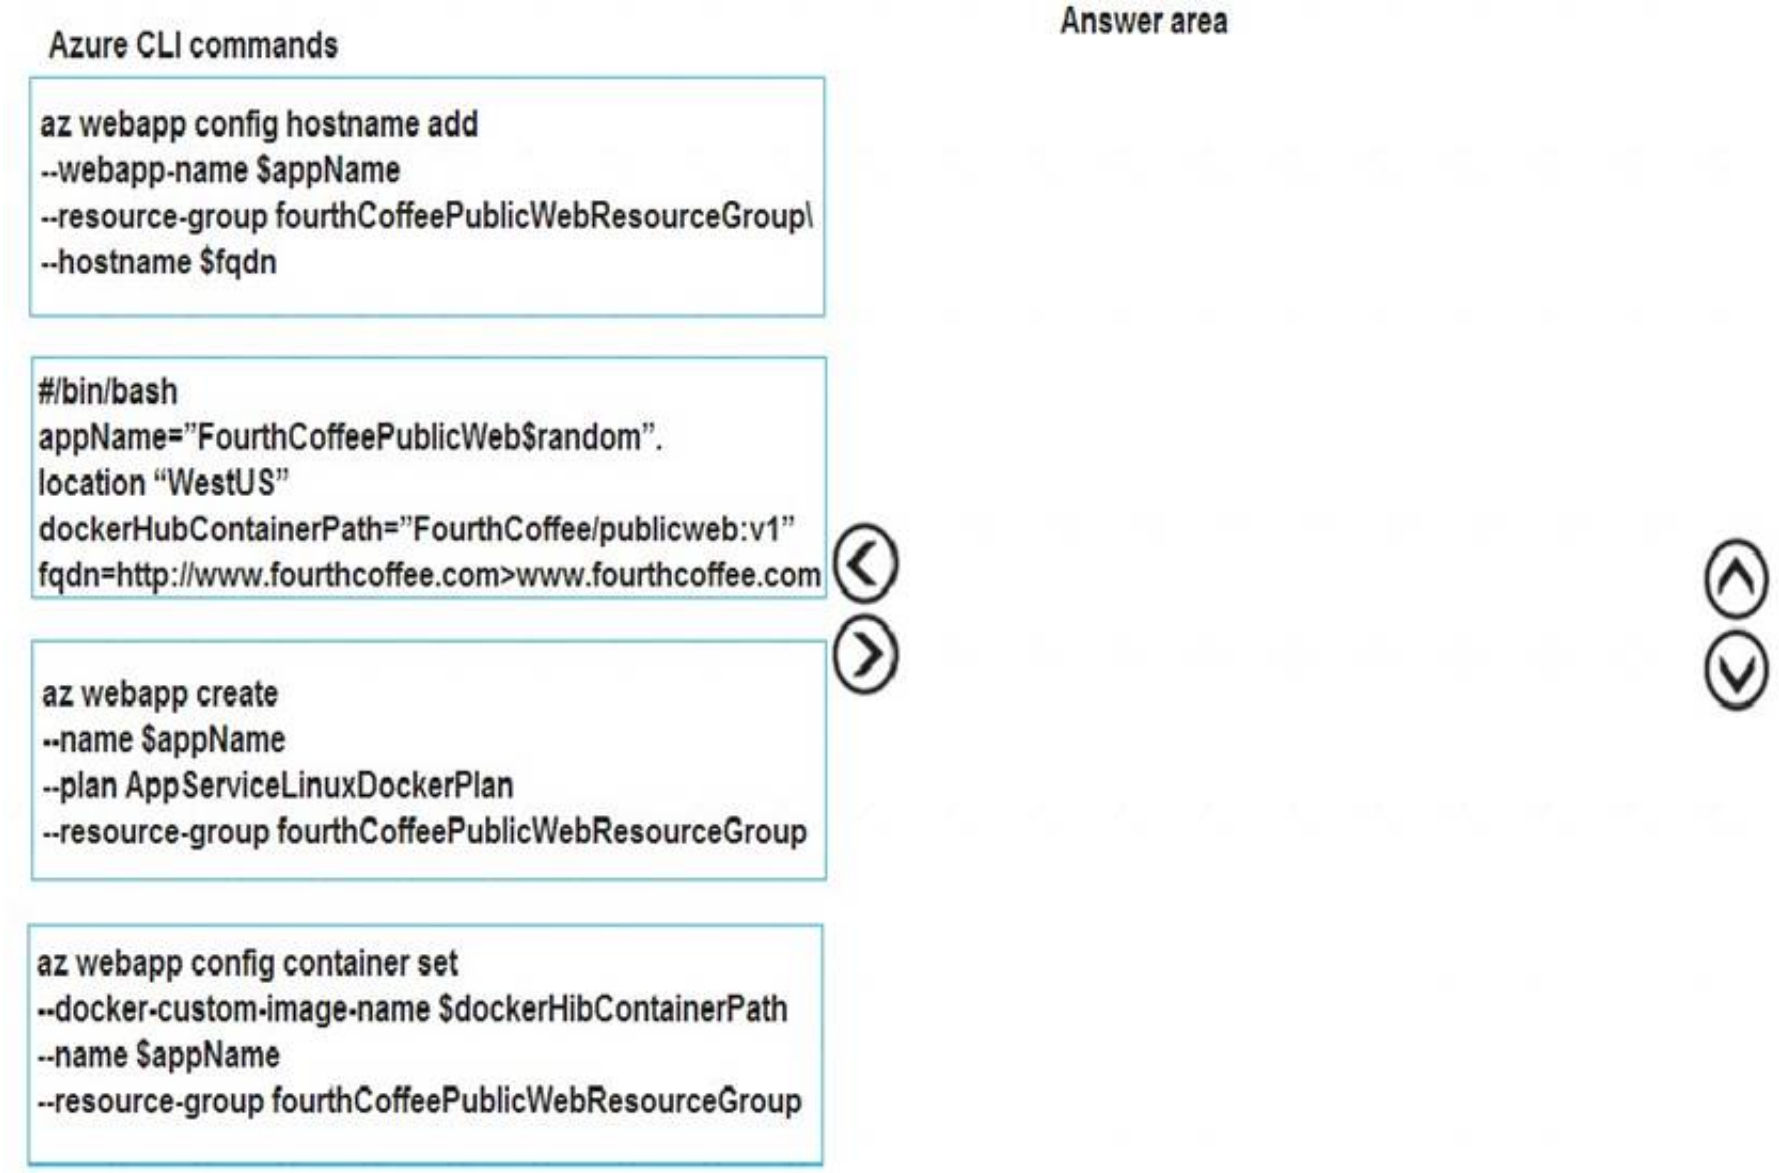

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Step 1: #bin/bash
The appName is used when the webapp-name is created in step 2. Step 2: az webapp config hostname add
The webapp-name is used when the webapp is created in step 3. Step 3: az webapp create
Create a web app. In the Cloud Shell, create a web app in the myAppServicePlan App Service plan with the az webapp create command.
Step : az webapp confing container set
In Create a web app, you specified an image on Docker Hub in the az webapp create command. This is good enough for a public image. To use a private image, you need to configure your Docker account ID and password in your Azure web app.
In the Cloud Shell, follow the az webapp create command with az webapp config container set.
References:
https://docs.microsoft.com/en-us/azure/app-service/containers/tutorial-custom-docker-image

NEW QUESTION 63
- (Exam Topic 3)
You are creating a CLI script that creates an Azure web app related services in Azure App Service. The web app uses the following variables:

Variable name	Value
$gitrepo	https://github.com/Contos/webapp
$webappname	Webapp1103

You need to automatically deploy code from GitHub to the newly created web app.
How should you complete the script? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

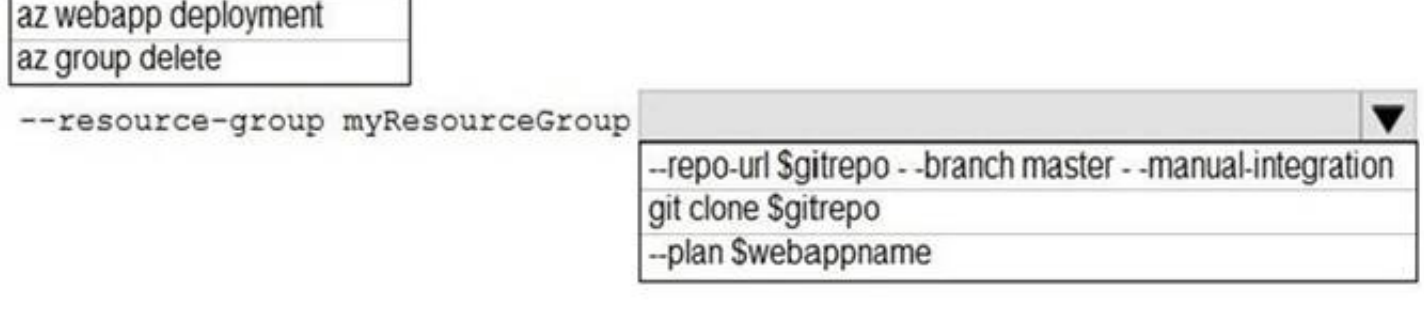

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: az appservice plan create
The azure group creates command successfully returns JSON result. Now we can use resource group to create a azure app service plan
Box 2: az webapp create Create a new web app..
Box 3: --plan $webappname
with the serviceplan we created in step 1. Box 4: az webapp deployment
Continuous Delivery with GitHub. Example:
az webapp deployment source config --name firstsamplewebsite1 --resource-group websites--repo-url $gitrepo
--branch master --git-token $token
Box 5: --repo-url $gitrepo --branch master --manual-integration Reference:
https://medium.com/@satish1v/devops-your-way-to-azure-web-apps-with-azure-cli-206ed4b3e9b1

NEW QUESTION 67
- (Exam Topic 3)
You develop a website. You plan to host the website in Azure. You expect the website to experience high traffic volumes after it is published. You must ensure that

the website remains available and responsive while minimizing cost. You need to deploy the website. What should you do?

A. Deploy the website to an App Service that uses the Shared service tie
B. Configure the App Service plan to automatically scale when the CPU load is high.
C. Deploy the website to a virtual machin
D. Configure the virtual machine to automatically scale when the CPU load is high.
E. Deploy the website to an App Service that uses the Standard service tie
F. Configure the App Service plan to automatically scale when the CPU load is high.
G. Deploy the website to a virtual machin
H. Configure a Scale Set to increase the virtual machine instance count when the CPU load

Answer: C

Explanation:
Windows Azure Web Sites (WAWS) offers 3 modes: Standard, Free, and Shared.
Standard mode carries an enterprise-grade SLA (Service Level Agreement) of 99.9% monthly, even for sites with just one instance.
Standard mode runs on dedicated instances, making it different from the other ways to buy Windows Azure Web Sites.

NEW QUESTION 70
- (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 are developing an Azure solution to collect point-of-sale fPOS) device data from 2,000 stores located throughout the world. A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data.
You must store the device data in Azure Blob storage. Device data must be correlated based on a device identifier. Additional stores are expected to open in the future.
You need to implement a solution to receive the device data.
Solution: Provision an Azure Event Hub. Configure the machine identifier as the partition key and enable capture.

A. Yes
B. No

Answer: A

Explanation:
References:
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-programming-guide

NEW QUESTION 72
- (Exam Topic 3)
Your company has several websites that use a company logo image. You use Azure Content Delivery Network (CDN) to store the static image.
You need to determine the correct process of how the CDN and the Point of Presence (POP) server will distribute the image and list the items in the correct order.
In which order do the actions occur? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

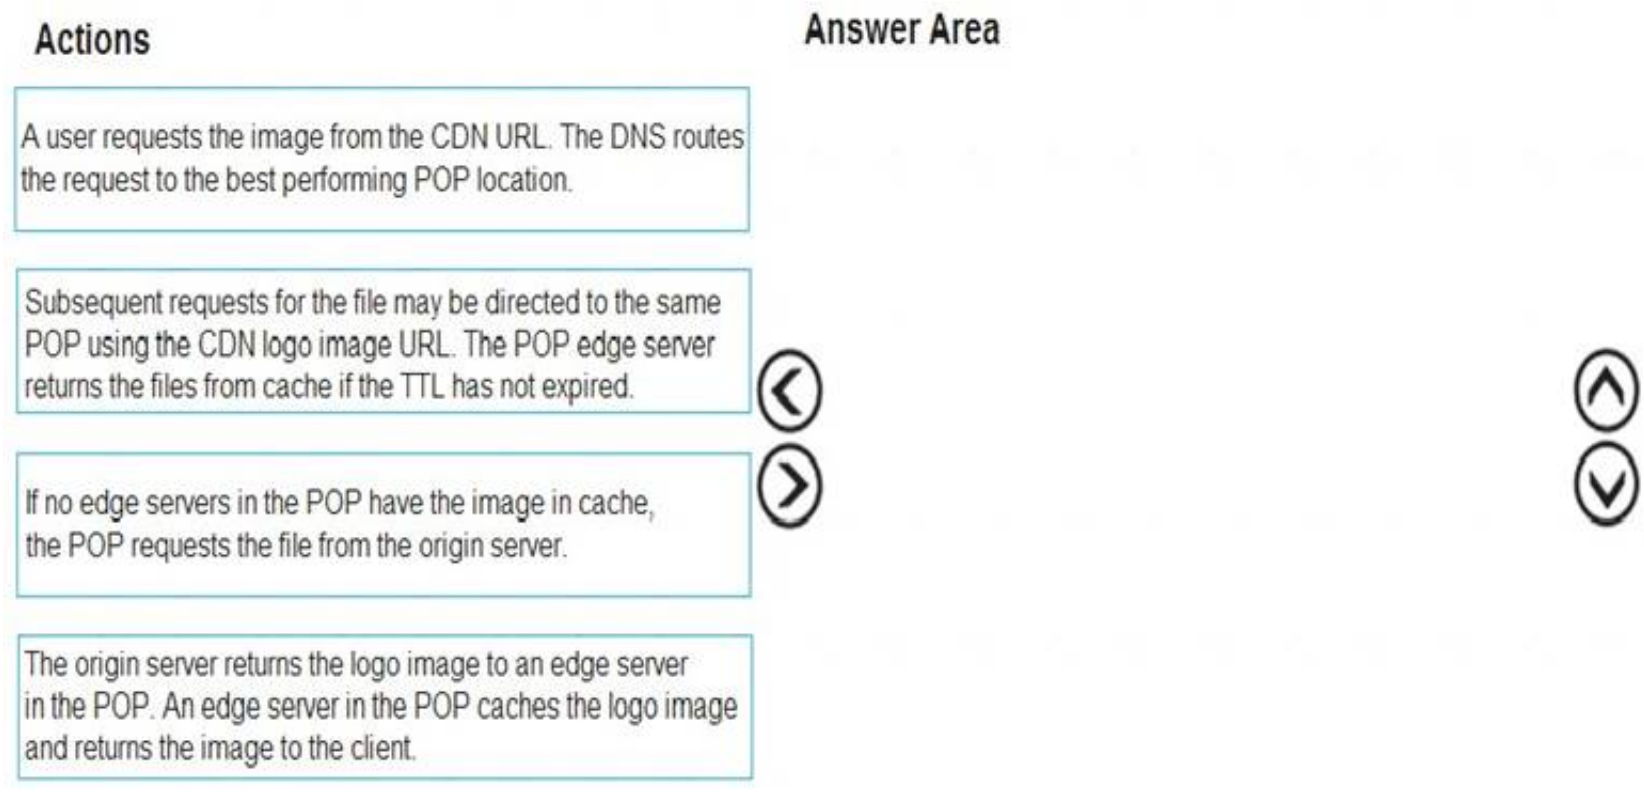

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: A user requests the image..
A user requests a file (also called an asset) by using a URL with a special domain name, such as <endpoint name>.azureedge.net. This name can be an endpoint hostname or a custom domain. The DNS routes the request to the best performing POP location, which is usually the POP that is geographically closest to the user.
Step 2: If no edge servers in the POP have the..
If no edge servers in the POP have the file in their cache, the POP requests the file from the origin server. The origin server can be an Azure Web App, Azure

Cloud Service, Azure Storage account, or any publicly accessible web server.
Step 3: The origin server returns the..
The origin server returns the file to an edge server in the POP.
An edge server in the POP caches the file and returns the file to the original requestor (Alice). The file remains cached on the edge server in the POP until the time-to-live (TTL) specified by its HTTP headers expires. If the origin server didn't specify a TTL, the default TTL is seven days.
Step 4: Subsequent requests for..
Additional users can then request the same file by using the same URL that the original user used, and can also be directed to the same POP.
If the TTL for the file hasn't expired, the POP edge server returns the file directly from the cache. This process results in a faster, more responsive user experience.
References:
https://docs.microsoft.com/en-us/azure/cdn/cdn-overview

NEW QUESTION 76
- (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 77
- (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 79
- (Exam Topic 3)
You have an application that provides weather forecasting data to external partners. You use Azure API Management to publish APIs.
You must change the behavior of the API to meet the following requirements:
• Support alternative input parameters.
• Remove formatting text from responses.
• Provide additional context to back-end services.
Which types of policies should you implement? To answer, drag the policy types to the correct scenarios. Each policy 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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 81
- (Exam Topic 3)
You are developing a software solution for an autonomous transportation system. The solution uses large data sets and Azure Batch processing to simulate navigation sets for entire fleets of vehicles.
You need to create compute nodes for the solution on Azure Batch. What should you do?

A. In Python, implement the class: TaskAddParameter
B. In Python, implement the class: JobAddParameter
C. In the Azure portal, create a Batch account
D. In a .NET method, call the method: BatchClient.PoolOperations.CreateJob

Answer: D

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.
Note:
Step 1: Create a pool of compute nodes. When you create a pool, you specify the number of compute nodes for the pool, their size, and the operating system. When each task in your job runs, it's assigned to execute on one of the nodes in your pool.
Step 2 : Create a job. A job manages a collection of tasks. You associate each job to a specific pool where that job's tasks will run.
Step 3: Add tasks to the job. Each task runs the application or script that you uploaded to process the data files it downloads from your Storage account. As each task completes, it can upload its output to Azure Storage.

NEW QUESTION 86
- (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.

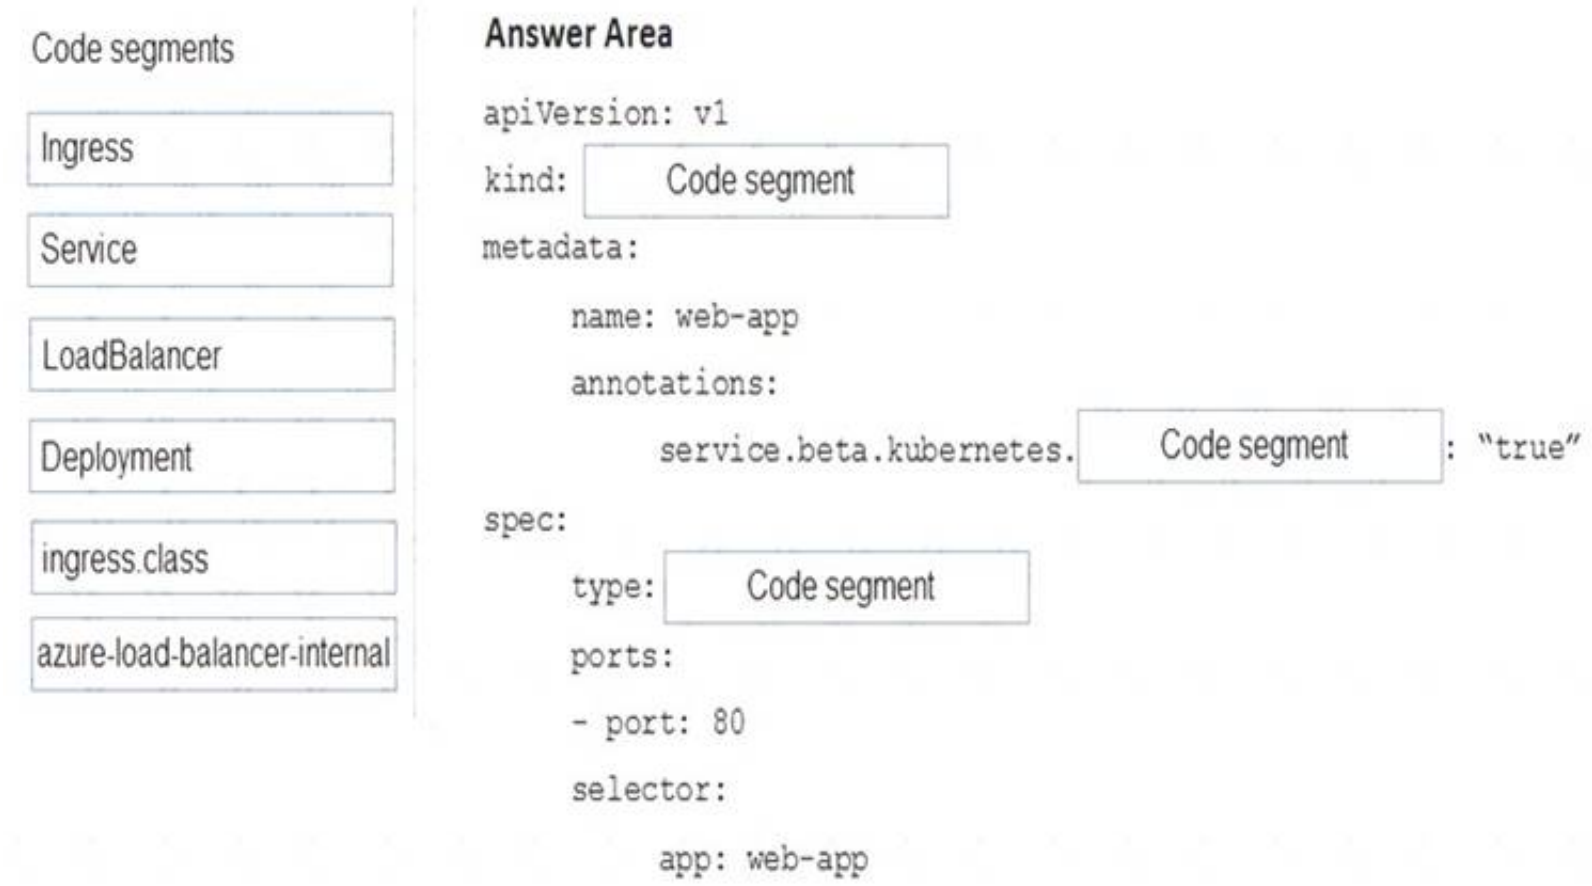

A. Mastered
B. Not Mastered

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 91
- (Exam Topic 3)
You have an Azure App Services Web App. Azure SQL Database instance. Azure Storage Account and an Azure Redis Cache instance in a resource group.
A developer must be able to publish code to the web app. You must grant the developer the Contribute role to the web app
You need to grant the role.
What two commands can you use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. New-AzureRmRoleAssignment
B. az role assignment create
C. az role definition create
D. New-AzureRmRoleDefinition

Answer: AB

Explanation:
References:
https://docs.microsoft.com/en-us/cli/azure/role/assignment?view=azure-cli-latest#az-role-assignment-create https://docs.microsoft.com/en-us/powershell/module/azurerm.resources/new-azurermroleassignment?view=azur

NEW QUESTION 94
- (Exam Topic 3)
A company uses Azure SQL Database to store data for an app. The data includes sensitive information.
You need to implement measures that allow only members of the managers group to see sensitive information. Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Include the managers group.
B. Exclude the managers group.
C. Exclude the administrators group.
D. Navigate to the following URL:

```
PUT https://management.azure.com/subscriptions/00000000-1111-2222-3333-444444444444
/resourceGroups/rg01/providers/Microsoft.Sql/servers/server01/databases/customers
/transparentDataEncryption/current?api-version=2014-04-01
```

E. Run the following Azure PowerShell command:

```
New-AzureRmSqlDatabaseDataMaskingRule -SchemaName "dbo" -TableName "customers" '
-ColumnName "ssn" -MaskingFunction "Default"
```

A. Option A
B. Option B
C. Option C
D. Option D
E. Option E

Answer: BE

Explanation:
Dynamic data masking helps prevent unauthorized access to sensitive data by enabling customers to designate how much of the sensitive data to reveal with minimal impact on the application layer.
SQL users excluded from masking - A set of SQL users or AAD identities that get unmasked data in the SQL query results.
Note: The New-AzureRmSqlDatabaseDataMaskingRule cmdlet creates a data masking rule for an Azure SQL database.
References:
https://docs.microsoft.com/en-us/powershell/module/azurerm.sql/new-azurermsqldatabasedatamaskingrule?view

NEW QUESTION 95
- (Exam Topic 3)

A company backs up all manufacturing data to Azure Blob Storage. Admins move blobs from hot storage to archive tier storage every month.
You must automatically move blocks to Archive tier after they have not been accessed for 180 days. The path for any item that is not archived must be placed in an existing queue. This operation must be performed automatically once a month. You set the value of TierAgeInDays to 180.
How should you configure the Logic App? To answer, drag the appropriate triggers or action blocks to the correct trigger or action slots. Each trigger or action block 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.

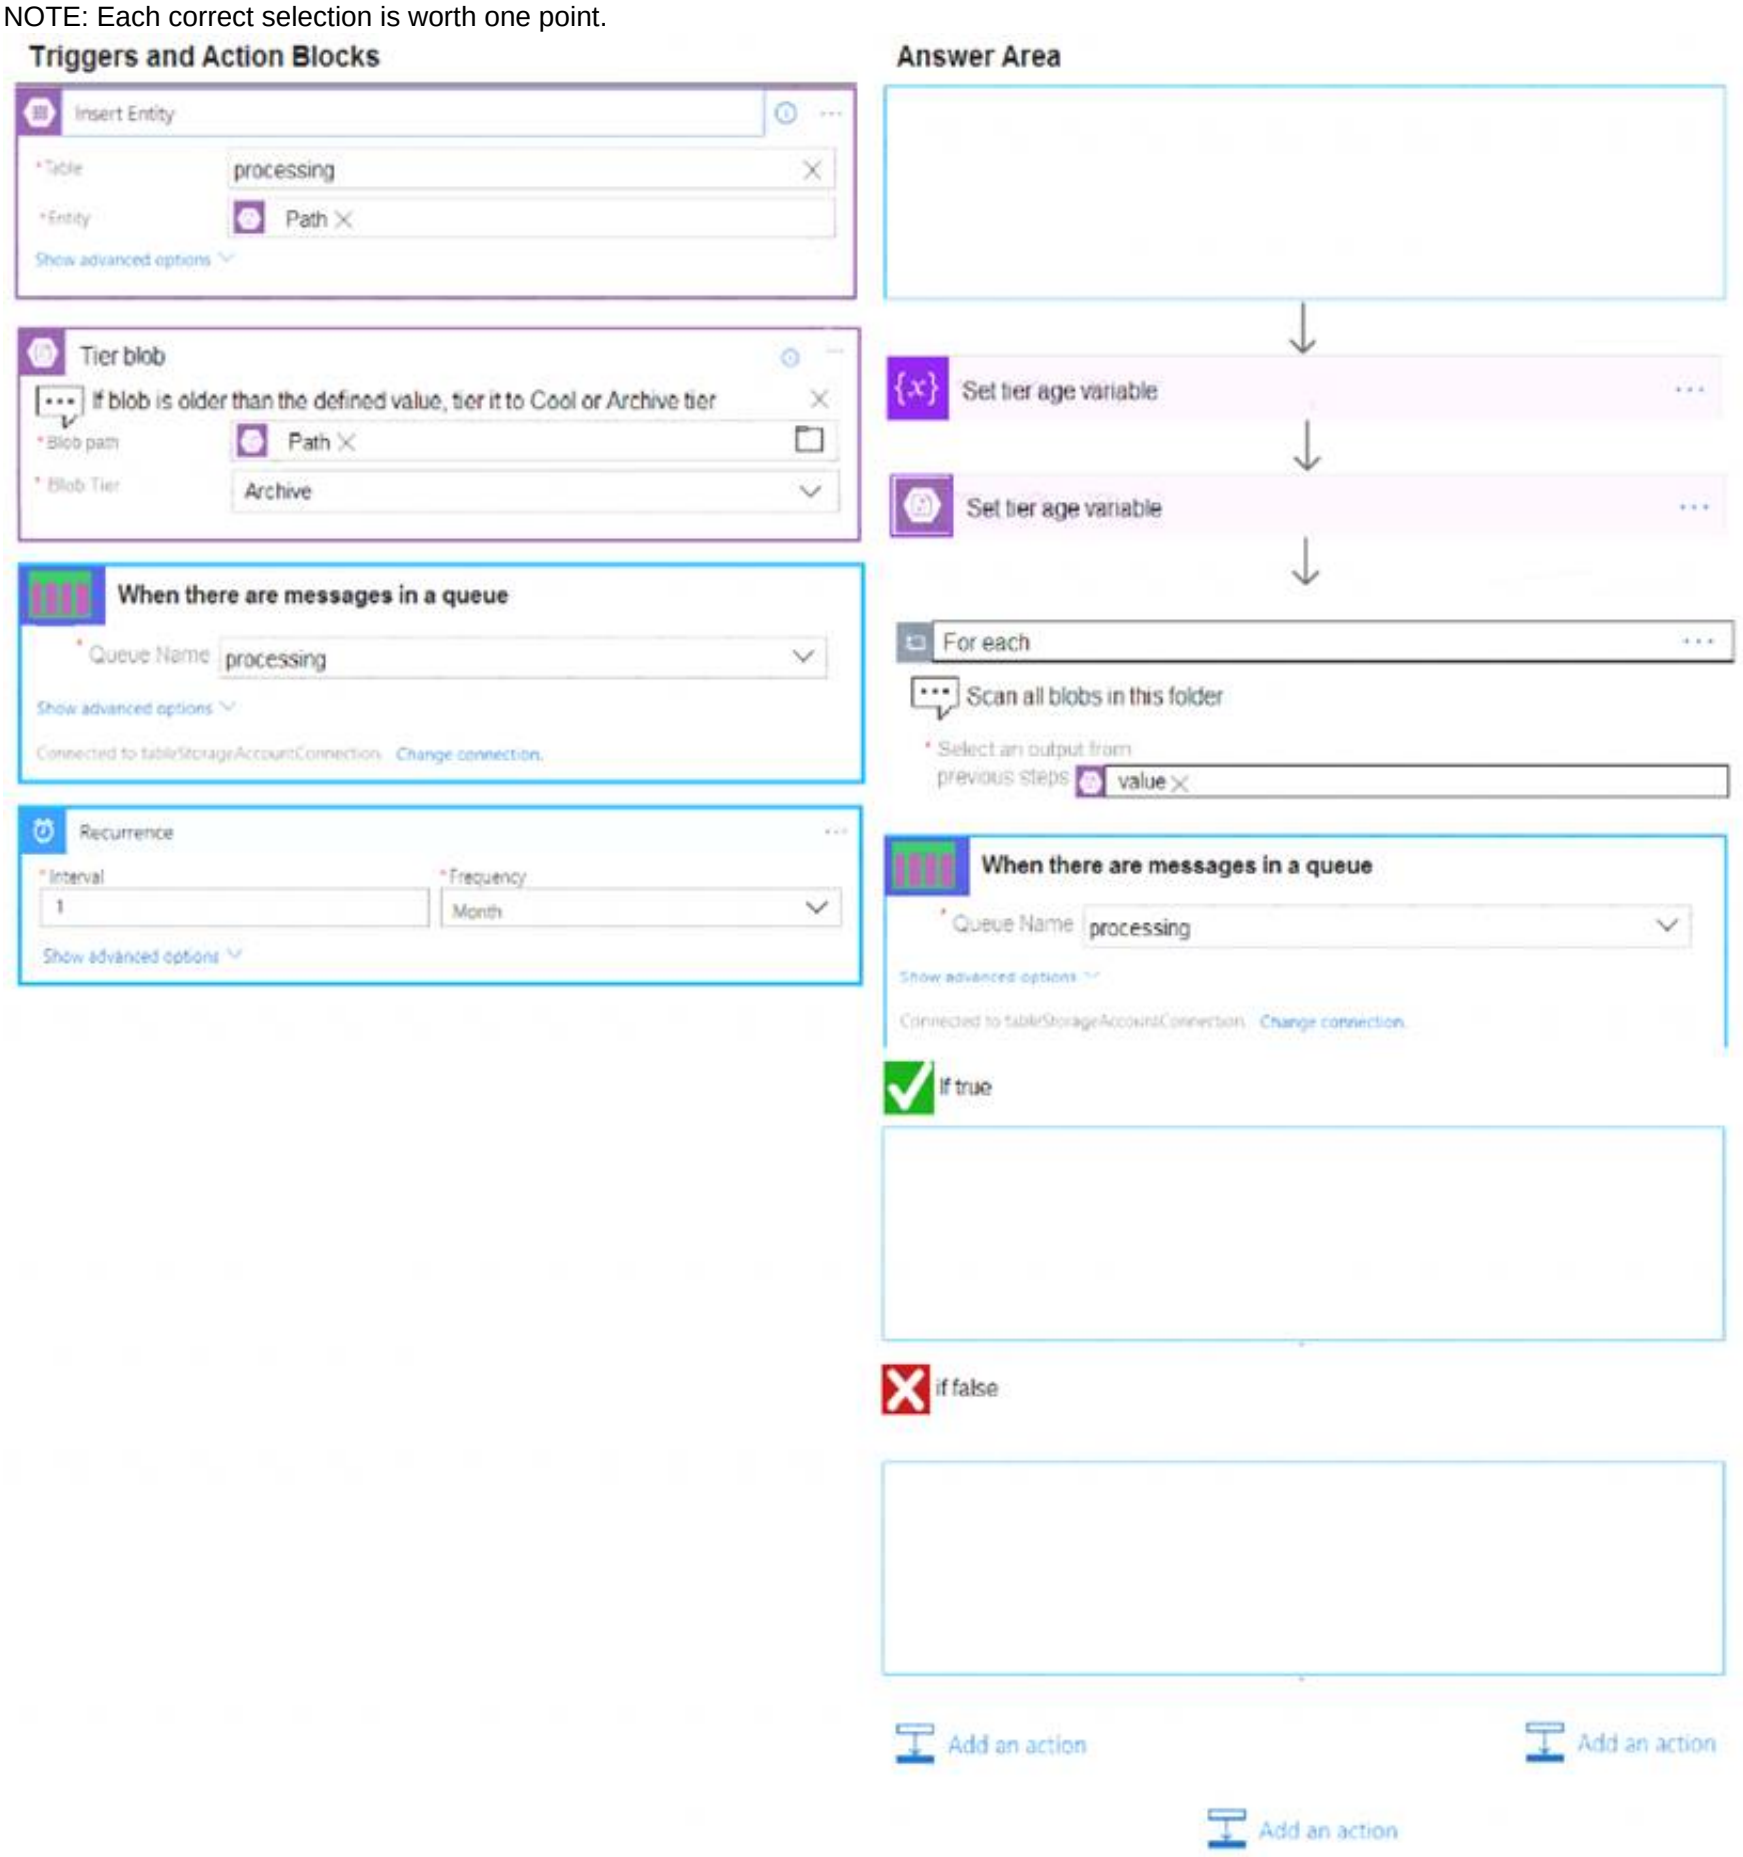

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Recurrence Box 2: Insert Entity
Box 3 (if true): Tier Blob Box 4: (if false):
Leave blank. References:
https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-perform-data-operations

NEW QUESTION 97
- (Exam Topic 3)
You are working for a company that designs mobile applications. They maintain a server where player records are assigned to their different games. The tracking system is new and in development.
The application uses Entity Framework to connect to an Azure Database. The database holds a Player table and Game table.
When adding a player, the code should insert a new player record, and add a relationship between an existing game record and the new player record.
The application will call CreatePlayerWithGame with the correct gameIdand the playerId to start the process. (Line numbers are included for reference only.)

```
01. namespace ContosoCradt
02. {
03.   public class PlayerDbContext : DbContext
04.   {
05.  public PlayerDbContext() : base ("name-dBConnString") { }
06.  public DbSet<Player> Players { get ; set ; }
07.  public DbSet<Game> Games { get ; set ]
08.  protected override void DmModelCreating{DBModelBuilder modelBuilder}
09.  {
10.   modelBuilder.Entity<Player>().MesMany(x => x.Games). WithMany (x => x Players);
11.  }
12.  }
13.   internal series class dbConfiguration : DbMigrationConfiguration<PlayerDbContext>
14.  {
15.  public dbConfiguration() . {AutomaticMigrationsEnabled = true ; }
16.  {
17.    public class mp
18.    {
19.  public void CreatePlayerWithGame(int playerId, int gameId) => AddPlayer(playerId, GetGame[gameId));
20.  public game GetGame(int gameId)
21.  {
22.    using (var db = new PlayerDbContext())
23.  {
24.     return db.Games.FirstOrDefault(x => x.GameId == gameId);
25.  }
26.   }
27.   public Player AddPlayer (int playerId, Game game)
28.   {
29.  using (var db = new PlayerDbContext())
30.  {
31.    var player = new Player
32.     {
33.      PlayerId = playerId,
34.      Games = new List <Game> {game },
35.  };
36.  db.Players.Add(player);
37.  db.SaveChanges();
38.  return player;
39.    }
40. }
41. public class Player
42. {
43.  public int PlayerId { get ; set; }
44.  public string PlayerName { get ; set; }
45.  public virtual List<Game> Games { get ; set; }
46.  }
47.  public class Game
48.  {
49.  public int GameIs { get ; set }
50.  public string Title { get ; set; }
51.   public string Platform { get ; set; ]
52.   public virtual List<Player> Players { get ; set; }
53.      }
54.   }
```

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 successfully insert a player record.	○	○
The code has a bug and will insert an additional copy of the Game record with a new Id.	○	○
The code has a bug and will insert the wrong gameId value.	○	○
There is a valid many-to-many relationship between Players and Games.	○	○

A. Mastered

B. Not Mastered

Answer: A

Explanation:
Many-to-many relationships without an entity class to represent the join table are not yet supported. However, you can represent a many-to-many relationship by including an entity class for the join table and mapping two separate one-to-many relationships.
protected override void OnModelCreating(ModelBuilder modelBuilder)
{
modelBuilder.Entity<PostTag>() HasKey(t => new { t.PostId, t.TagId }); modelBuilder.Entity<PostTag>() HasOne(pt => pt.Post)
WithMany(p => p.PostTags) HasForeignKey(pt => pt.PostId); modelBuilder.Entity<PostTag>() HasOne(pt => pt.Tag) WithMany(t => t.PostTags) HasForeignKey(pt => pt.TagId);
}
}

NEW QUESTION 102
- (Exam Topic 3)
You manage several existing Logic Apps.
You need to change definitions, add new logic, and optimize these apps on a regular basis.
What should you use? To answer, drag the appropriate tools to the correct functionalities. Each tool 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

Tools	Functionality	Tool
Logic Apps Designer	Edit B2B workflows	
Code View Editor	Edit definitions in JSON	
Enterprise Integration Pack	Visually and functionality	

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Enterprise Integration Pack
After you create an integration account that has partners and agreements, you are ready to create a business to business (B2B) workflow for your logic app with the Enterprise Integration Pack.
Box 2: Code View Editor
To work with logic app definitions in JSON, open the Code View editor when working in the Azure portal or in Visual Studio, or copy the definition into any editor that you want.
Box 3: Logical Apps Designer
You can build your logic apps visually with the Logic Apps Designer, which is available in the Azure portal through your browser and in Visual Studio.
References:
https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-enterprise-integration-b2b https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-overview

NEW QUESTION 105
- (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.

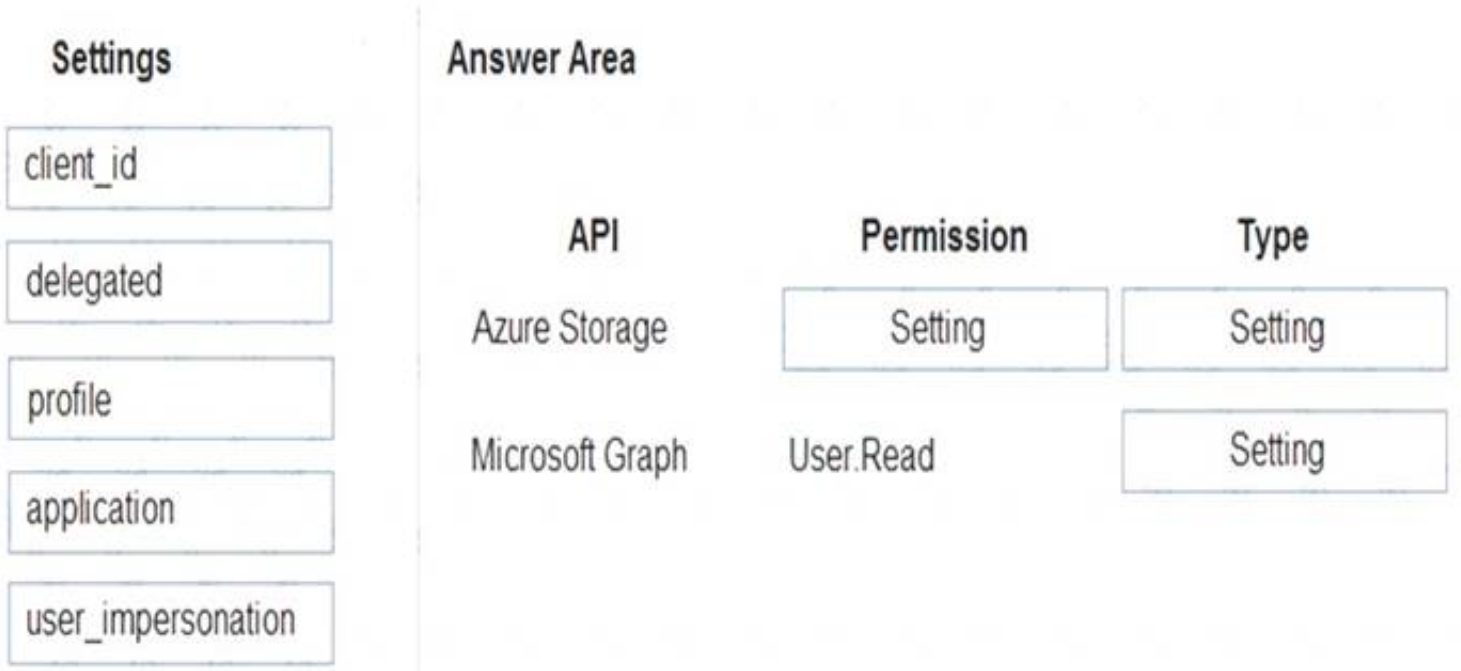

A. Mastered
B. Not Mastered

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 106
- (Exam Topic 3)
You develop an Azure web app. You monitor performance of the web app by using Application Insights. You need to ensure the cost for Application Insights does not exceed a preset budget. What should you do?

A. Implement ingestion sampling using the Azure portal.
B. Set a daily cap for the Application Insights instance.
C. Implement adaptive sampling using the Azure portal.
D. Implement adaptive sampling using the Application Insights SDK.
E. Implement ingestion sampling using the Application Insights SDK.

Answer: D

Explanation:
Sampling is an effective way to reduce charges and stay within your monthly quota.
You can set sampling manually, either in the portal on the Usage and estimated costs page; or in the ASP.NET SDK in the .config file; or in the Java SDK in the ApplicationInsights.xml file, to also reduce the network traffic.
Adaptive sampling is the default for the ASP.NET SDK. Adaptive sampling automatically adjusts to the volume of telemetry that your app sends. It operates automatically in the SDK in your web app so that telemetry traffic on the network is reduced.
References:
https://docs.microsoft.com/en-us/azure/azure-monitor/app/sampling

NEW QUESTION 110
- (Exam Topic 3)
You are working for Contoso, Ltd.
You define an API Policy object by using the following XML markup:

```
<set-variable name= "bodySize" value="@(context.Request.Headers["Content-Length"] [0])"/>
<choose>
 <when condition= "@(int.Parse(context.Variables.GetValueOrDefault<string> ("bodySize"))<512000)">
</when>
<otherwise>
     <rewrite-uri template= "/put"/>
     <set-backend-service base-url= "http://contoso.com/api/9.1/"/>
</otherwise>
</choose>
```

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 Area

Statement	Yes	No
The XML segment belongs in the <inbound> section of the policy.	○	○
If the body size is >256k, an error will occur.	○	○
If the request is http://contoso.com/api/9.2/, the policy will retain the higher version.	○	○

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Yes
Use the set-backend-service policy to redirect an incoming request to a different backend than the one specified in the API settings for that operation. Syntax: <set-backend-service base-url="base URL of the backend service" />
Box 2: No
The condition is on 512k, not on 256k. Box 3: No
The set-backend-service policy changes the backend service base URL of the incoming request to the one specified in the policy.
Reference:
https://docs.microsoft.com/en-us/azure/api-management/api-management-transformation-policies

NEW QUESTION 114
- (Exam Topic 3)
You are developing a ticket reservation system for an airline.
The storage solution for the application must meet the following requirements:

- Ensure at least 99.99% availability and provide low latency.
- Accept reservations event when localized network outages or other unforeseen failures occur.
- Process reservations in the exact sequence as reservations are submitted to minimize overbooking or selling the same seat to multiple travelers.
- Allow simultaneous and out-of-order reservations with a maximum five-second tolerance window. You provision a resource group named airlineResourceGroup in the Azure South-Central US region. You need to provision a SQL SPI Cosmos DB account to support the app.

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

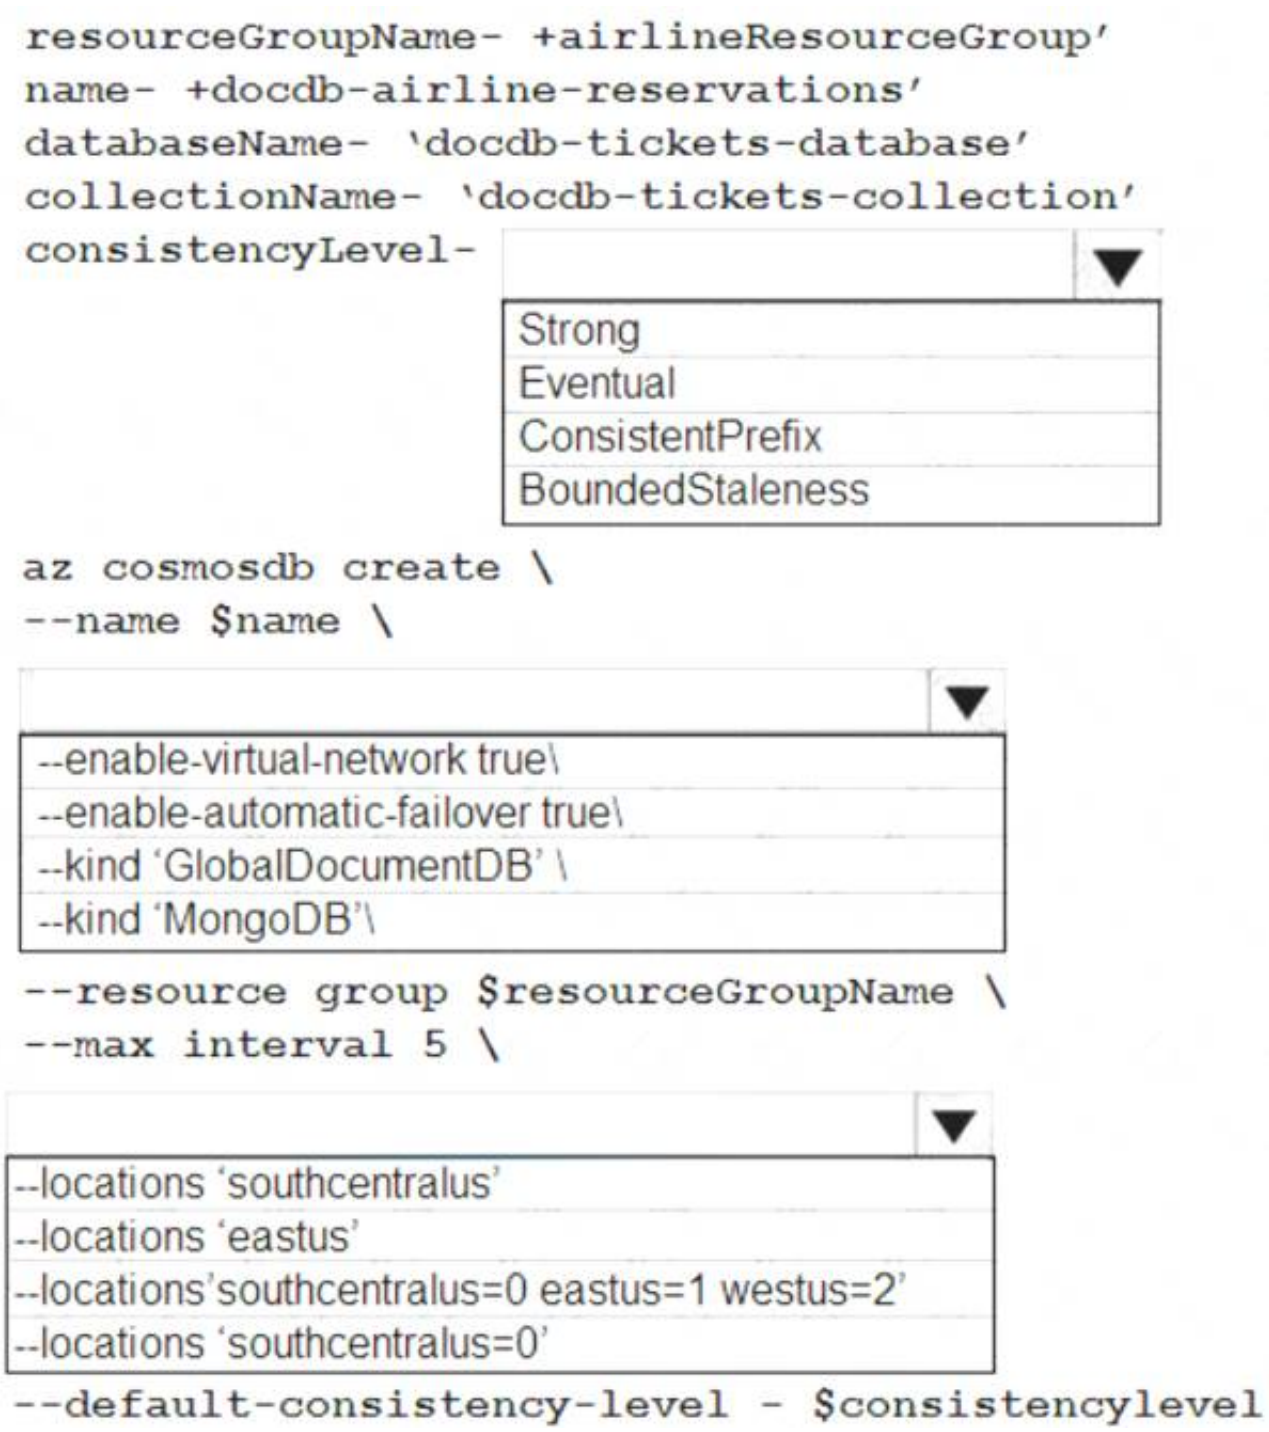

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: BoundedStaleness
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. In other words, 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 Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/cosmos-db/manage-with-cli.md

NEW QUESTION 116
- (Exam Topic 3)
You are developing an application. You have an Azure user account that has access to two subscriptions. You need to retrieve a storage account key secret from Azure Key Vault.
In which order should you arrange the PowerShell commands to develop the solution? To answer, move all commands from the list of commands to the answer area and arrange them in the correct order.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Get-AzSubscription
If you have multiple subscriptions, you might have to specify the one that was used to create your key vault. Enter the following to see the subscriptions for your account:
Get-AzSubscription
Step 2: Set-AzContext -SubscriptionId
To specify the subscription that's associated with the key vault you'll be logging, enter: Set-AzContext -SubscriptionId <subscriptionID>
Step 3: Get-AzStorageAccountKey You must get that storage account key.
Step 4: $secretvalue = ConvertTo-SecureString <storageAccountKey> -AsPlainText -Force
Set-AzKeyVaultSecret -VaultName <vaultName> -Name <secretName> -SecretValue $secretvalue After retrieving your secret (in this case, your storage account key), you must convert that key to a secure
string, and then create a secret with that value in your key vault.
Step 5: Get-AzKeyVaultSecret
Next, get the URI for the secret you created. You'll need this URI in a later step to call the key vault and retrieve your secret. Run the following PowerShell command and make note of the ID value, which is the secret's URI:
Get-AzKeyVaultSecret –VaultName <vaultName> Reference:
https://docs.microsoft.com/bs-latn-ba/Azure/key-vault/key-vault-key-rotation-log-monitoring

NEW QUESTION 118
- (Exam Topic 3)
You have a web app named MainApp. You are developing a triggered App Service background task by using the WebJobs SDK. This task automatically invokes a function code whenever any new data is received in a queue.
You need to configure the services.
Which service should you use for each scenario? To answer, drag the appropriate services to the correct scenarios. Each service 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.

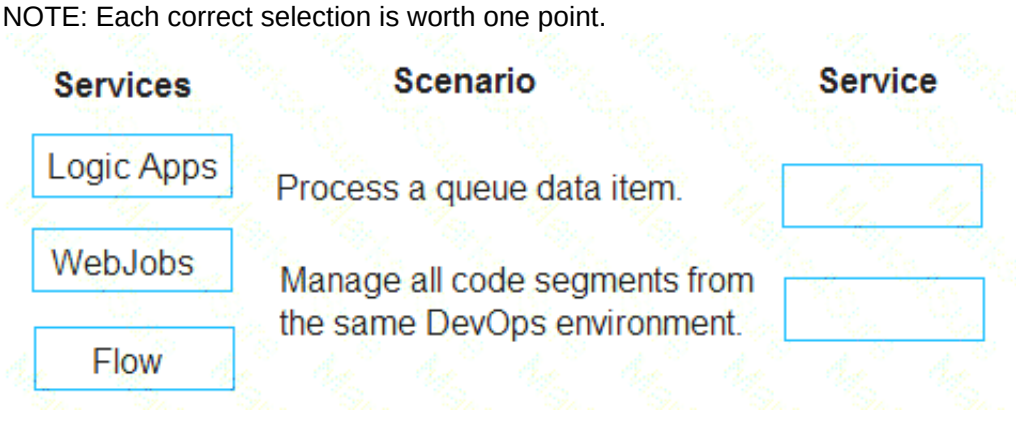

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: WebJobs
A WebJob is a simple way to set up a background job, which can process continuously or on a schedule. WebJobs differ from a cloud service as it gives you get less fine-grained control over your processing environment, making it a more true PaaS service.
Box 2: Flow

NEW QUESTION 120
- (Exam Topic 3)
You are developing Azure WebJobs.
You need to recommend a WebJob type for each scenario.
Which WebJob type should you recommend? To answer, drag the appropriate WebJob types to the correct scenarios. Each WebJob 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.

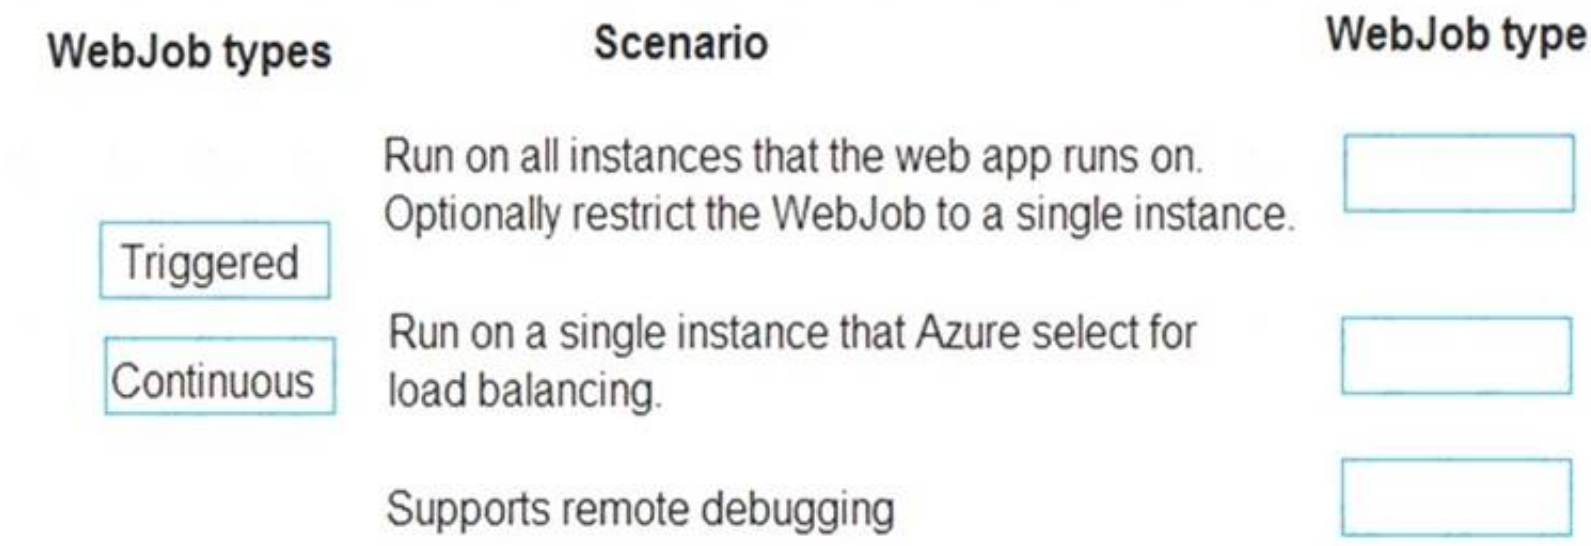

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Box 1: Continuous
Continuous runs on all instances that the web app runs on. You can optionally restrict the WebJob to a single instance.
Box 2: Triggered
Triggered runs on a single instance that Azure selects for load balancing. Box 3: Continuous
Continuous supports remote debugging. Note:
The following table describes the differences between continuous and triggered WebJobs.

Continuous	Triggered
Starts immediately when the WebJob is created. To keep the job from ending, the program or script typically does its work inside an endless loop. If the job does end, you can restart it.	Starts only when triggered manually or on a schedule.
Runs on all instances that the web app runs on. You can optionally restrict the WebJob to a single instance.	Runs on a single instance that Azure selects for load balancing.
Supports remote debugging.	Doesn't support remote debugging.

References:
https://docs.microsoft.com/en-us/azure/app-service/web-sites-create-web-jobs

NEW QUESTION 125
- (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.

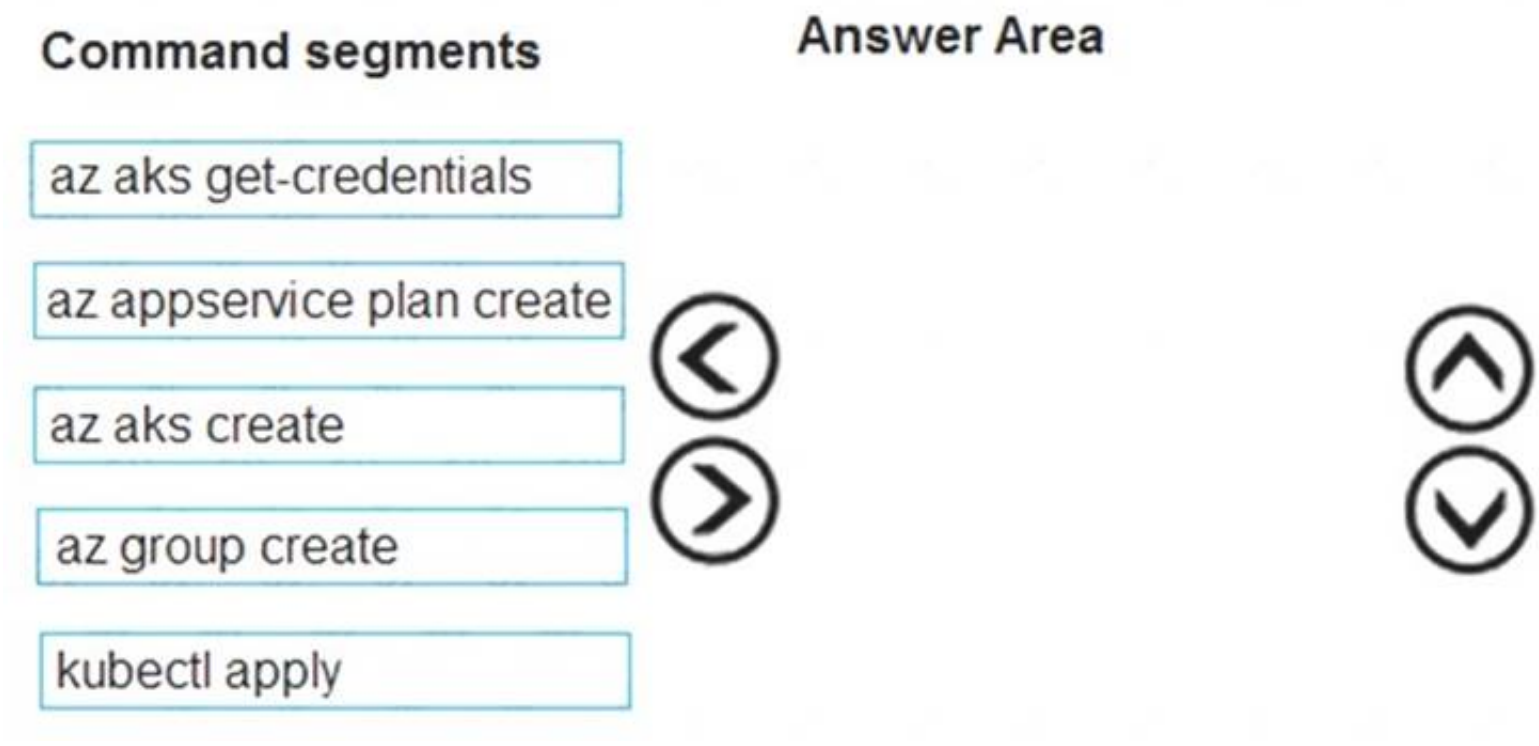

A. Mastered
B. Not Mastered

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 129
- (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 131
- (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.

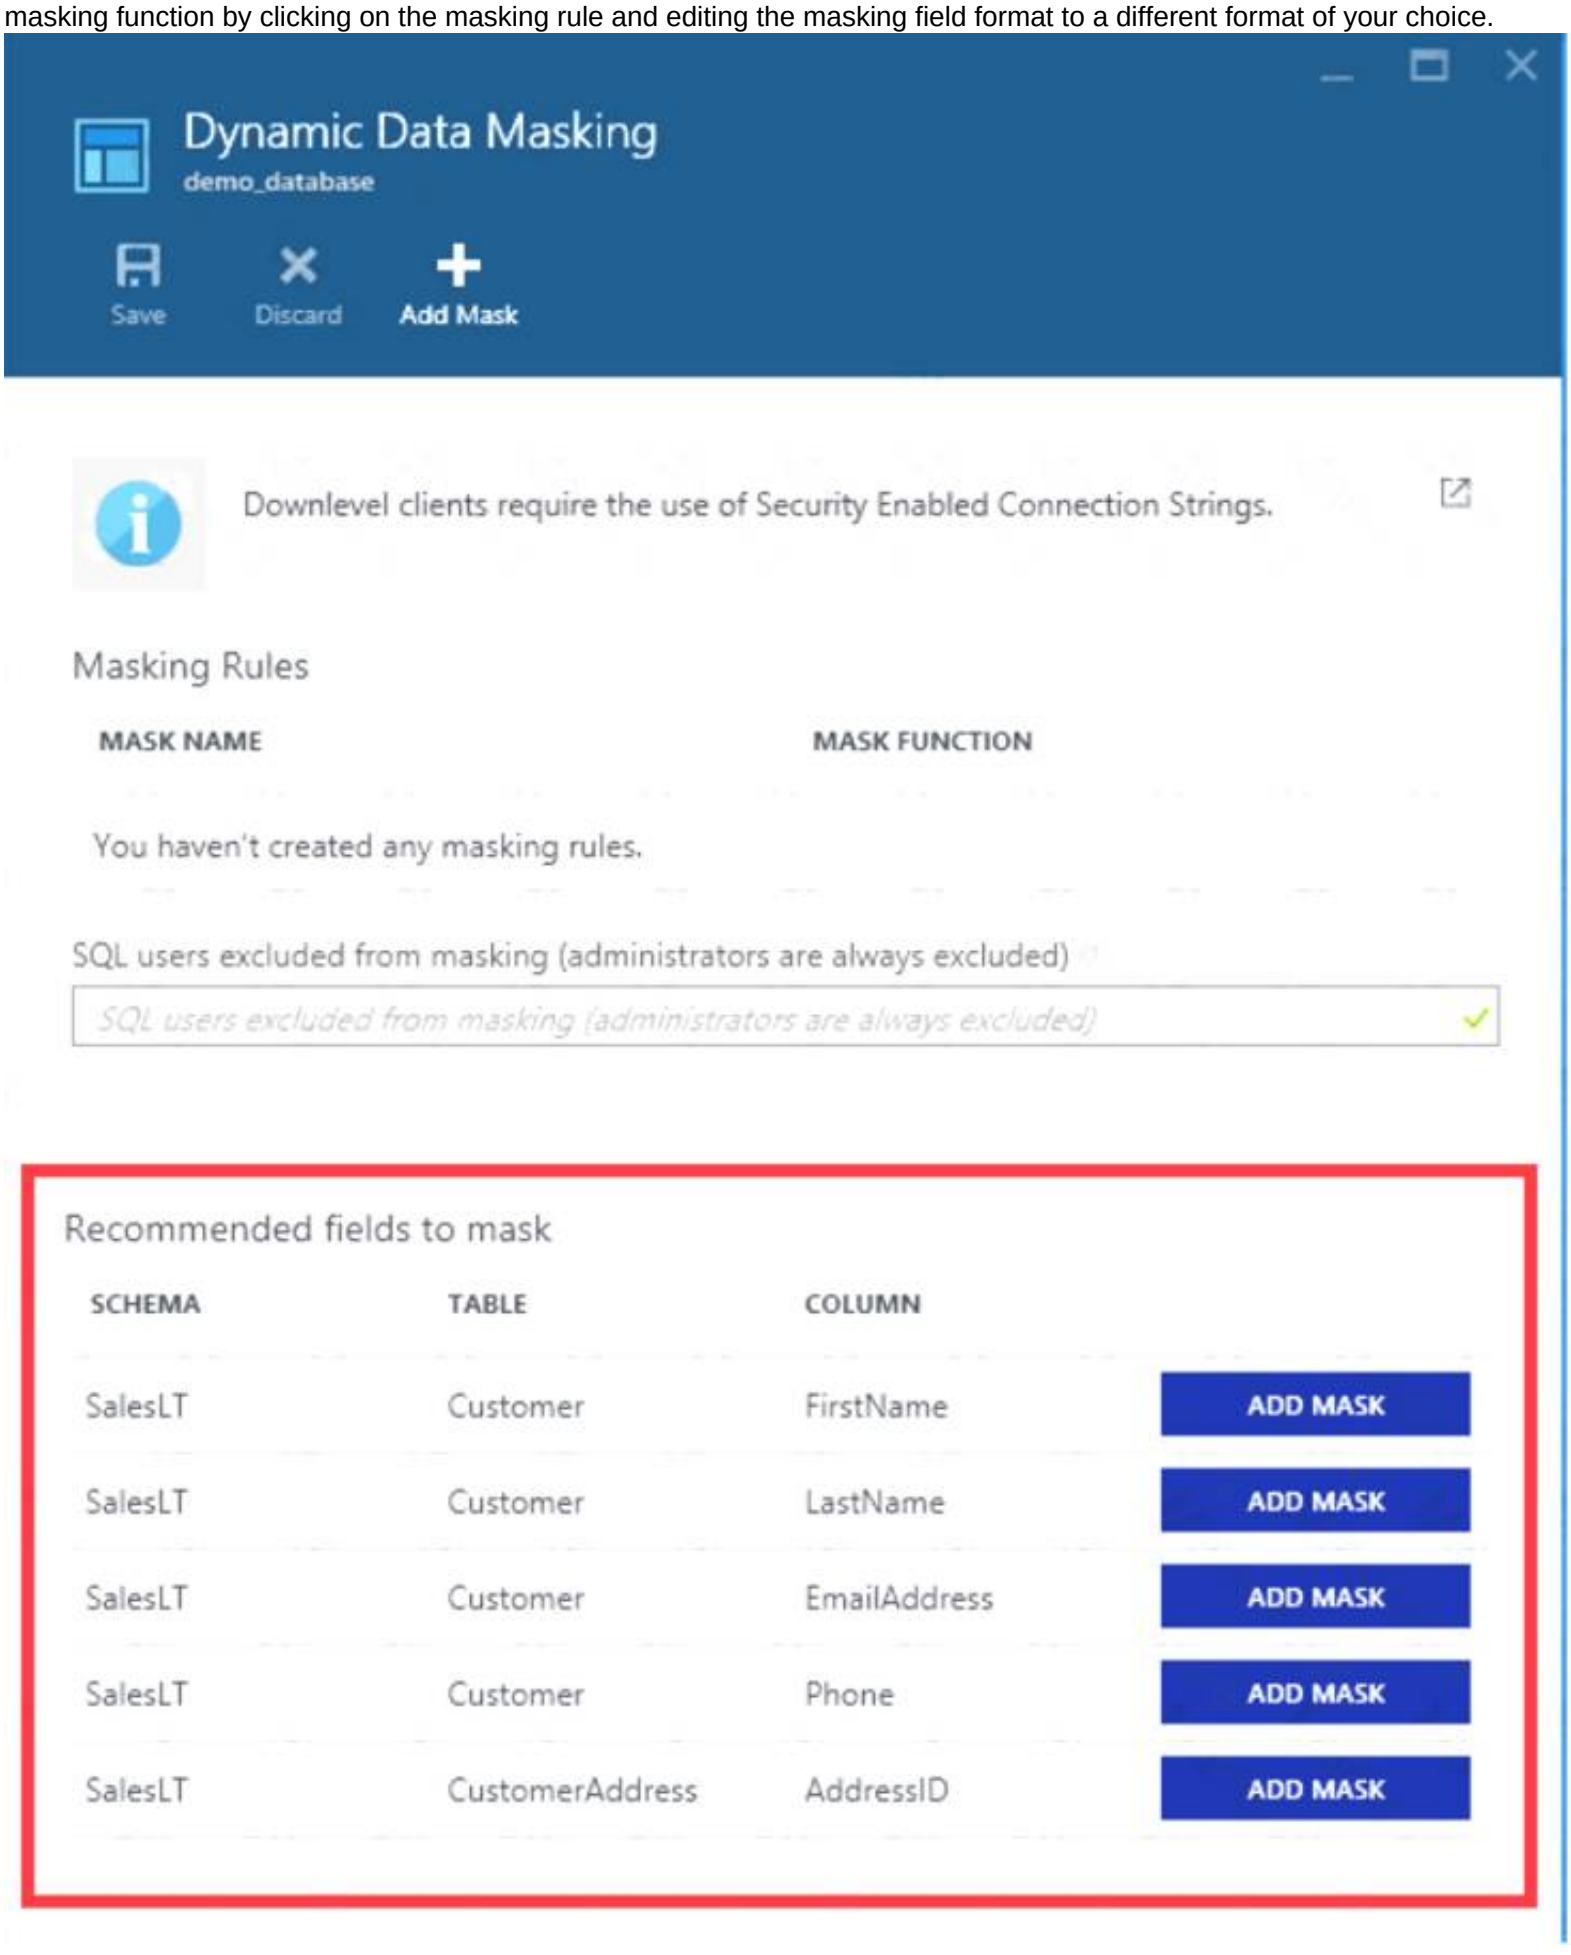

References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dynamic-data-masking-get-started-portal

NEW QUESTION 135
- (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 propert
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 139
- (Exam Topic 3)
You are developing an internal website for employees to view sensitive data. The website uses Azure Active Directory (AAD) for authentication. You need to implement multifactor authentication for the website.
What should you do? Each correct answer presents part of the solution. NOTE; Each correct selection is worth one point.

A. In Azure AD, create a new conditional access policy.
B. In Azure AD, enable application proxy.
C. Configure the website to use Azure AD B2C.
D. In Azure AD conditional access, enable the baseline policy.
E. Upgrade to Azure AD Premium.

Answer: AE

Explanation:
References:
https://docs.microsoft.com/en-us/azure/active-directory/authentication/howto-mfa-getstarted

NEW QUESTION 142
.....

Relate Links

100% Pass Your AZ-204 Exam with Exambible Prep Materials

https://www.exambible.com/AZ-204-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/